



OXFORD

Purpose in the Universe

*The Moral and Metaphysical Case for
Ananthropocentric Purposivism*

TIM MULGAN

Purpose in the Universe

*The moral and metaphysical
case for Ananthropocentric
Purposivism*

Tim Mulgan

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Tim Mulgan 2015

The moral rights of the author have been asserted

First Edition published in 2015

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2015938208

ISBN 978-0-19-964614-2

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

<i>Acknowledgements</i>	vii
1. Introduction	1
2. Meta-ethics	33
Part I. The Case against Atheism	
3. Cosmological Arguments	65
4. Teleological Arguments	99
5. Mysticism	130
6. Ontological Arguments	161
Part II. The Case against Benevolent Theism	
7. Arguments from Scale	193
8. Arguments from Evil	220
9. Religious Diversity	261
10. Immortality	280
Part III. Ananthropocentric Purposivist Morality	
11. A Dialogue	325
12. Human Well-being	343
13. Ananthropocentric Purposivist Moral Theory	373
<i>Bibliography</i>	417
<i>Index</i>	429

Acknowledgements

During the ten years I have been working on this book, I have accumulated more debts than I can hope to acknowledge here. Most of that time was spent in the Department of Philosophy at St Andrews, and I could not have asked for a more collegial academic environment. I am grateful to all my St Andrews colleagues, and especially to two successive heads of school—Peter Clark and Katherine Hawley. I am also grateful to my colleagues at the University of Auckland, where I have been based since 2012, and to the Princeton University Center for Human Values for their generous hospitality in 2010/11.

I have presented earlier versions of the ideas in this book in seminars, conferences, and public lectures at the universities of St Andrews, Auckland, Dundee, Edinburgh, Glasgow, Leeds, London, Otago, Oxford, Reading, Rennes, Rome, and Victoria. I am very grateful to audiences at all these venues for their comments and suggestions. For discussions of specific chapters or themes of the book, I am grateful to Nigel Biggar, Timothy Chappell, Peter Clark, Katherine Hawley, Melissa Lane, Janet McLean, John Perry, and Alan Torrance. I have also learnt much over many years from conversations with Sarah Broadie.

I am very grateful to the four anonymous readers for Oxford University Press for their detailed, forthright, and insightful comments; to Peter Momtchiloff for his advice and encouragement over several iterations of this project; and to Kim Richardson for copy-editing so efficiently.

For her constant love and companionship, for travelling to the ends of the earth, and for bringing purpose to my small corner of the universe, I am forever grateful to Janet McLean.

1

Introduction

Our universe is religiously ambiguous. It can be read in strikingly different ways. Two familiar readings are materialist atheism and the benevolent God of the Abrahamic faiths. In this book, I defend a less familiar reading, which I dub Ananthropocentric Purposivism (AP for short): there is a cosmic purpose, but human beings are irrelevant to it. AP borrows traditional theist arguments to defend a cosmic purpose, and then it borrows traditional atheist arguments to reject a human-centred purpose.

My academic background is not in theology or metaphysics or philosophy of religion, but in contemporary analytic moral philosophy. In this book, I explore the moral case for, and the moral implications of, AP. Contemporary philosophy typically begins with some privileged world view (theist or atheist), and then asks where (if at all) morality fits in. I begin with substantive normative commitments, and then ask what metaphysical picture best fits those commitments.

In this introductory chapter, I first outline AP in contrast to its two main rivals, atheism and benevolent theism. I then explore connections between morality and metaphysics, beginning with the key notion of religious ambiguity. After outlining my own moral commitments, and their role in the project, I offer some remarks on the moral impact of AP. The chapter closes with a summary of the rest of the book.

1.1 What is AP?

In 2003, I taught the existence of God in a large first-year metaphysics course, covering very familiar ground: cosmological, teleological, ontological, evil. I was struck by the following thought. Theist arguments, if they succeed, establish that there is a God *of some sort*. Atheist arguments, if they succeed, show that there is not a God *of one specific sort*. So why not cut and paste the positive and negative arguments, and thus end up with a God *of a different sort*?

This book is the result. I aim to introduce, defend, and explore the following underdeveloped alternative to both theism and atheism:

Ananthropocentric Purposivism (AP). The universe has a non-human-centred purpose—one to which human beings are completely irrelevant or incidental.

AP is a general idea that can be fleshed out in many different ways. For ease of exposition, I often focus on *theist AP*, where a personal creator gives the universe

its purpose. But AP could take other forms—perhaps modelled on John Leslie’s axiarchism, or on traditional Idealism or Platonism.¹ Indeed, a *personal* God may seem uncomfortably anthropocentric. Other things equal, AP is more comfortable with an impersonal cosmic purpose. I try to remain as agnostic as possible regarding the exact ontological nature of the cosmic purpose and its source. My primary interest is not in the metaphysical details of AP, but rather in its moral implications.

I contrast AP with two familiar positions. The first is traditional Western monotheism. The universe was created by a God who cares for individual human beings. Philosophers have offered many competing accounts of God down the centuries. Richard Swinburne’s definition captures the features that have become standard among analytic philosophers: ‘God is a necessarily existing person without a body who necessarily is eternal, perfectly free, omnipotent, omniscient, perfectly good, and the creator of all things.’² Unless I explicitly depart from it, I will be operating with this definition of God.

I dub this opposing view ‘benevolent theism’ (BT for short). The boundary between BT and AP is vague. BT says that God loves each individual human being; that human beings are an essential part of God’s plan for the cosmos; and that God created this cosmos (in part) because it would contain human beings. AP says that God does not love individual human beings; that God has no interest in the fate of humanity; and that the presence of human beings is a cosmic accident. Many intermediate positions are possible. Perhaps human beings are a subsidiary part of God’s plan for creation—less important than something else, but still not completely irrelevant; or perhaps God cares about human beings once they emerge, even though their existence played no role in God’s decision to create this particular cosmos. Some arguments for AP over BT are consistent with these intermediate positions, while others are not. Because my primary aim is to introduce AP into the philosophical landscape, I contrast two ‘pure’ positions: a benevolent God who cares for each human being and whose purpose for cosmic creation involves human beings; and an AP that denies both that God cares for human beings at all and that we played any part in God’s creative decision.

We are familiar with situations where my individual contribution or welfare is irrelevant. My vote makes no difference to the election, for instance, because my preferred candidate wins by a large margin. But here *aggregate* human impact is still relevant. *Our* votes together determine the election. AP goes further. Human beings are completely irrelevant to the cosmic purpose, to objective value, and to God’s plans, concerns, or reasons for creation. Our cosmic significance mirrors the electoral significance of mice, whose interests and opinions have no influence on human-centred elections.

¹ See, e.g., Leslie, *Universes*, and chapter 3 in this book.

² Swinburne, *The Existence of God*, p. 7.

However, AP does not insist on a complete separation between human beings and cosmic purpose. AP can admit that human beings resemble some aspect of the cosmic purpose; that we can understand that purpose to some degree; or that we can bring it into our lives in a meaningful way. The cosmic purpose can matter to us, even if we do not matter to it. These possibilities come to the fore in part III. Without them, AP would be much less morally interesting.

BT religions often contain both anthropocentric and theocentric strands. Much Christian theology, for instance, emphasizes the unknowability, the alienness, the non-human-ness of God.³ While this theology can come close to AP, actual theist religions typically contain some key doctrine that rules out AP. For Christianity, this is the Incarnation. The Christian God becomes a human being for the sake of other human beings. This is not something any AP God would do. AP cannot plausibly be offered as a Christian apologetic.⁴

I contrast AP with 'benevolent theism' rather than 'classical theism', both to include non-classical theisms such as process theology and also to emphasize the human-centred element. However, this terminology is idiosyncratic, and may look like a sleight of hand to conceal AP's shortcomings. One reader raises the following objection:

A cosmic purpose necessarily requires a divine person whose purpose it is. (Purposes require persons by definition, and only a divine person could give purpose to a cosmos.) The simplest imaginable divine person is the Omni-God of classical theism. *Ceteris paribus*, the simplest hypothesis is the most likely. Therefore, the classical theist Omni-God is the most a priori likely divine person. But this God, by definition, is benevolent to human beings. Therefore, AP is at a distinct disadvantage relative to classical theism, a disadvantage which is obscured by substituting 'benevolent theism' for 'classical theism'.⁵

As we shall see in chapter 3, the step from simplicity to probability is not straightforward. (How do we measure comparative simplicity? Is theism the simplest hypothesis or the most complex? And why expect reality to be simple in the first place?) But AP has two other replies. The first is that the claim that purposes *necessarily* require persons is a substantive thesis that must be defended, not a merely definitional truth. As a moral philosopher, my central interest in cosmic purpose lies in its connection to objective values. And it is certainly not obvious that the latter require persons. (Impersonal Platonism may be misguided, but it is not incomprehensible.) Perhaps purposivism will ultimately collapse into theism, but this remains to be seen.

³ One striking recent example is James Gustafson's two-volume *Ethics from a Theocentric Perspective*.

⁴ Theocentric Christians thus tend to downplay the second person of the Trinity. For instance, in 664 pages, Gustafson's *Ethics from a Theocentric Perspective* devotes only five pages to Christology and only one to Jesus. (I am grateful to Nigel Biggar and Mark Wynne for discussion of this issue.)

⁵ This paragraph paraphrases a much fuller discussion from an anonymous referee.

AP's second reply is more direct and much more central to our project. AP denies the *final* step from divine simplicity to divine benevolence. AP initially seems to be a *rival* for classical theist metaphysics or perfect being theology. I contend, instead, that AP can borrow from these traditions, by breaking the traditional link between perfection (even *moral* perfection) and concern for human beings. AP need not deny a divine creator. It can even admit an omnipotent, omniscient, or morally perfect God. Indeed, AP could borrow Swinburne's own definition in its entirety, by denying that a perfectly good being must care for human beings. Swinburne's definition is thus neutral between AP and BT. Swinburne himself reads BT into his definition, because he assumes that perfect goodness implies benevolence towards humans. But whether human beings are morally considerable from the perspective of a morally perfect divine person is the very question at issue between AP and BT. BT is one interpretation of classical theism: the result of combining perfect being theology with the substantive moral claim that human beings matter. By denying the latter, AP offers a rival *interpretation of classical theism*, not a rival to classical theism.

Similarly, BT and AP can offer competing interpretations of non-classical theist traditions such as finite theism or process theism. I return to these alternative traditions briefly in chapters 7 and 8. However, apart from a brief exploration of Leslie's axiarchism in chapters 3 and 4, I mostly consider BT and AP as competing interpretations of the classical theist Omni-God, who is omniscient, omnipotent, necessary, and morally perfect. Of course, controversy surrounds the coherence and compossibility of these divine attributes. (Can a morally perfect God be truly free? Can God create a stone that God cannot lift? And so on.) In so far as it is not committed to any Omni-God, these controversies are grist to AP's mill. But I largely set them aside in this book. I concentrate instead on the central question that separates AP and BT: does God care for us?

I contrast AP and BT with *atheism*. If we take atheism to be the denial of BT, then AP is atheist. If atheism is the denial that there is any God, then AP overlaps with atheism, because cosmic purpose can take non-theist forms. However, I shall take atheism to be the denial of all supernatural entities, divine beings, *and* cosmic purposes. Indeed, I go further. My atheist is also a naturalist who regards science as the model for all human epistemic enterprises and the final arbiter of what exists. (I discuss naturalism at greater length in chapters 2 to 5.) While this rules out some possible positions, it does cover most contemporary philosophical atheism, and leaves us with a simple threefold division. BT says the universe has a human-centred purpose; AP says it has a non-human-centred purpose; while atheism says it has no purpose.⁶

⁶ One problem case for my simple threefold distinction between AP, BT, and atheism is atheist religions, such as Buddhism, which combine supernatural commitments with the rejection of God. I briefly discuss Buddhist atheism, and its relationship to AP, in chapter 9.


The traditional debate between BT and atheism often tracks other philosophical disagreements in metaphysics, morality, and methodology. BT is often combined with belief in personal immortality, incompatibilist free will, dualism about mind and body, and robust moral realism; while atheism is typically combined with materialism, determinism, the denial of personal immortality, and a naturalist or anti-realist account of ethics. AP seeks a middle ground by separating the question of cosmic purpose from these other debates. In particular, AP separates cosmic purpose from the idea that human beings are metaphysically special.

AP is not unknown in the history of philosophy. (Similar themes are explored in Hinduism and Taoism, and in some less orthodox Western forms of Deism and Idealism.) But AP is definitely not a familiar position in contemporary Western philosophy. AP is worth exploring in part because it is a comparatively unexplored option. And there is also a surprisingly strong philosophical case for AP. Parts I and II develop that case, arguing that it is at least as strong as the case for either BT or atheism. Any philosopher who takes BT and atheism seriously should also take AP seriously.

AP is the conjunction of two claims:

- (1) The universe has a purpose; and
- (2) The universe does not have a human-centred purpose.

AP borrows the best anti-theist arguments from atheists and the best anti-atheist arguments from BT. BT arguments establish only (1), while atheist counter-arguments establish only (2). When atheism and BT are the only options on the table, these arguments succeed. But faced with a three-way choice between BT, atheism, and AP, they combine to establish AP.



Parts I and II explore many different arguments. No single argument is decisive and none is indispensable. Rather, each contributes to a *cumulative* case for AP. Part I develops the case for (1), arguing that AP can borrow the best traditional arguments for BT. Part II develops the case for (2), arguing that AP can borrow the best traditional arguments for atheism. The conclusion of parts I and II is that arguments leading to (1) are much stronger than those leading beyond to the negation of (2)—and vice versa on the atheist side. Theists and atheists should both

regard AP as their main rival. If BT and atheism are both worth exploring, then so is AP.

This will suffice to interest philosophers of religion in AP. But many moral philosophers have no interest in religion. The more ambitious task of this book is to convince those moral philosophers that, in addition to its metaphysical merits, AP is relevant *to them*. This is the task of part III. It is prefigured in section 1.5, and in the next chapter.

I claim that there is a philosophical case for AP that rivals BT and atheism. Ideally, I would defend this claim by presenting an actual case for AP that was as powerful as the best extant arguments for theism and atheism. Sadly, this task is probably beyond any single philosopher. It is certainly beyond me. The case for BT contains many arguments honed over many centuries. For most of its history, the relationship between God and morality was one of the central questions in Western philosophy and the philosophical elaboration of BT attracted the most brilliant thinkers of the age, while the past few decades have seen the best medieval arguments sharpened using the tools of contemporary analytic philosophy.

The extant philosophical case for atheism is less impressive. Within Western philosophy, atheism has moved too quickly from a blasphemy that dare not speak its name to a dominant world view that need offer no defence. Only within the past generation have atheist analytic philosophers confronted the need to justify themselves to their theist colleagues. Atheist apologetics is still in its infancy. Nonetheless, the philosophical resources available to contemporary atheism, while less formidable than those available to the theist, still far exceed anything any one individual proponent of AP could hope to construct in a single lifetime, let alone a single book (even one as unreasonably long as this).

Aside from sheer quality, the arguments for theism and atheism also display a *diversity* that AP cannot (yet) hope to match. The ‘case’ for BT is really a myriad of divergent but complementary approaches, rooted in distinct philosophical traditions, methods, and assumptions. Even confining themselves to the analytic mainstream, theists seeking justification can choose between the diverse approaches of Robert Adams, William Alston, William Lane Craig, Alvin Plantinga, Alexander Pruss, Richard Swinburne, Peter van Inwagen, or Linda Zagzebski, to name but a few.

I cannot offer a case for AP that is as rich, deep, or brilliant as those of its rivals. My goal is much more modest. I seek to illustrate AP’s potential by sketching the outlines of one particular cumulative case for AP, built on one initial set of pre-commitments and preoccupations, and emerging from one particular philosophical tradition. My hope is that those who prefer an alternative starting point may see how it too, in sympathetic hands, could ground an argument for AP.

I have described my case for AP as ‘cumulative’. Familiar cumulative arguments for BT start with a clearly defined thesis, and proceed by steadily increasing its probability. The classic exemplar is Richard Swinburne, who begins with a standard Christian definition of God, assigns a prior probability to the hypothesis that God

exists (using a priori appeals to relative simplicity), and then presents a series of interlocking inductive arguments designed to raise the probability of that hypothesis. The goal is to demonstrate that God probably exists.⁷

Things are different for AP. There is no agreed definition and thus no pre-existing hypothesis to be tested. AP begins as the vague claim that there is some non-human-centred cosmic purpose. Several key terms are under-defined (What counts as a purpose? What makes it cosmic? How exactly are humans excluded?), as are some key collateral concepts (Does a purpose require a person? Does AP need a God? Is God a perfect being? How does cosmic purpose relate to value, normativity, or morality?). Rather than settling these controversial questions by fiat at the outset, I progressively flesh out my own particular conception of AP through the course of the book. Each successive argument not only contributes to the case for AP; it also offers new insights into the content of AP. The view that emerges is one precisification: what AP might look like if one were persuaded by all the arguments in this book. Proponents of AP who are persuaded by other arguments, or by only a subset of those presented here, will endorse different interpretations.

One contrast is especially salient. *Simple AP* says that the universe was created by someone who does not care for human beings. Simple AP is easy to imagine and understand. Nor is it entirely implausible. For instance, if one combines a fine-tuning argument for theism with an argument from evil against *benevolent* theism, one will conclude that there is a creator God who is not benevolent. Many cultures have creation myths that instantiate Simple AP, where some distant creator sets the world in motion and then plays no further part in human affairs. (It is only lesser non-creator deities who concern themselves with the mortal realm.)

If I sought only to defend Simple AP, this book could be much shorter. Instead, I aim at *Normative AP*, where the non-human-centred cosmic purpose is a ground for objective values and external reasons that have normative significance for human beings. Normative AP clearly encounters difficulties that Simple AP avoids. Normative AP is much harder to grasp, understand, or imagine. Even if we grant Simple AP, we may still find Normative AP incredible. This is especially true if our Normative AP is *exclusive*: holding that non-human-centred cosmic purpose is the sole source of normativity for human beings. (Among other unpalatable conclusions, Exclusive Normative AP implies that human suffering has no objective significance at all.)

I will try to defend Exclusive Normative AP in this book. (Or, at least, to render it less implausible.) However, I am not committed to this extreme view. In part III, I explore more moderate variants of Normative AP, where cosmic purpose is only one source of values and reasons among others. On the other hand, despite its apparent difficulties, I *am* committed to Normative AP rather than Simple AP. The

⁷ For instance, Swinburne himself estimates the probability that God exists at somewhere between 0.2 and 0.8. (Swinburne, 'The Argument to God from Fine-Tuning Reassessed', p. 113.)

former is much more philosophically interesting, and more compelling, than the latter. Normative AP has resources that Simple AP lacks.

One might think the case for Simple AP must be stronger than the case for Normative AP, because the latter entails the former. But the real question is whether one can endorse Simple AP without also endorsing Normative AP. Can the proponent of Simple AP leave open the normative significance of cosmic purpose? I shall argue that she cannot. The best philosophical case for Simple AP already commits one to Normative AP. We cannot establish the existence of a cosmic purpose without also establishing its normative significance. We do not first prove AP using morally neutral arguments, and then assess its connection to values and reasons. Rather, if it arrives at all, AP comes already saturated with objective values.

The essential normativity of AP is the central theme of this book. This claim often strikes secular moral philosophers as absurd. But consider an analogous contrast between two interpretations of Benevolent Theism:

- *Simple (Morally Agnostic) BT*: God exists, but the normative significance of this fact for human beings is left open.
- *Normative BT*: God exists, and God's plans and purposes are a source of objective values and reasons with normative significance for human beings.

The dialectic is the same as for AP. On the one hand, Normative BT faces challenges, difficulties, and imaginative barriers that Simple BT avoids. Most notorious, of course, is the Euthyphro dilemma. Drawing on that dilemma, some atheists argue that, while Simple BT is imaginable (though under-supported by the evidence), Normative BT makes no sense. God could not have the right kind of normative significance. God's cosmic purpose might influence my response to moral facts about values and reasons, but it cannot *ground* those facts. (We return to Euthyphro in chapter 2.) On the other hand, Normative BT has resources that Simple BT lacks, because morality itself can now provide evidence that God exists.

Despite the extra burdens of Normative BT, no serious theist philosopher defends Simple BT. No doubt this is partly for doctrinal reasons. But the philosophical case for Normative BT is also much stronger and more interesting than the case for any morally agnostic theism. Arguments for BT work much better if they deliver a God who has an intimate connection to human morality.

1.2 Ambiguity: Religious and Philosophical

I have argued that, if BT and atheism are worth exploring, then so is AP. Of course, not everyone thinks that BT and atheism *are* both worth exploring. Some atheists find traditional theist arguments worthless, arguments from evil unanswerable, and the case for naturalism compelling. They conclude that there is nothing to be said for theism. Some theists find some argument for BT compelling and regard arguments from evil as worthless. They conclude that there is nothing to be said for atheism.

None of these people is likely to find this book very interesting. But many philosophers agree that knockdown arguments are rare anywhere in philosophy, and especially here. An emerging theme of the literature in philosophy of religion is that our universe is *religiously ambiguous*.⁸ BT and atheism are both reasonable interpretations of the available evidence. AP offers an alternative interpretation—a new way to read our cosmos.

In this book, I assume religious ambiguity. Given our limited understanding, radically different interpretations of the basic nature of the cosmos are equally reasonable: idealism versus realism, theism versus materialism, naturalism versus non-naturalism. Reason alone—scientific, philosophical, mathematical, or moral—cannot resolve metaphysical debate. Religious ambiguity is a non-eliminable feature of our current epistemic situation. It may even be intrinsic to the human condition *per se*.⁹ By highlighting the credentials of AP, I hope to strengthen the case for religious ambiguity. But my project is aimed squarely at those who already find this idea compelling.

Ambiguity is not peculiar to *religion*. It is a pervasive feature of all interesting *philosophical* questions. Philosophical debate is never settled by rational argument. Philosophy never answers its own questions. Consider the following philosophical questions, many of which will occupy us later in this book. Is genuine free will (the kind necessary for moral responsibility) compatible with determinism? Do human beings possess genuine free will? Is there a plausible reductionist naturalist account of knowledge or consciousness or mystical experience or modality or value or morality? If not, should we become sceptics about these domains or should we reject naturalism? Will the best moral theory take a deontological, consequentialist, or virtue ethical form? Is there any best moral theory? Are moral statements truth-apt? If so, are any of them true? If not, what on earth do they mean?

Philosophical ambiguity goes hand in hand with *methodological pluralism*. We should not presuppose that all philosophical arguments fit a common coin. In particular, arguments for the existence of God come in a bewildering array of styles. Many atheists try to reduce this rich variety to a single pattern: any argument for the existence of *x* must be an inference to the best explanation and all adequate explanations are scientific. Of course, some BTs do offer inferences to the best explanation. And some even present *scientific* explanations. But these are often the least successful BT arguments, and the least philosophically interesting. Not all good philosophical arguments are inferences to the best explanation. Of the classic arguments, only cosmological, teleological, and (some) moral arguments seek to *explain*

⁸ For references and further discussion, see chapters 5 and 9.

⁹ One intriguing question is whether religious ambiguity is also a necessary feature of the epistemic situation of *other* (non-human) non-divine agents, such as aliens or superintelligent machines. (The answer may turn on whether God has some particular reason to remain hidden from all created intelligences.) I briefly explore the significance of possible non-human agents in chapter 7, and hope to address it at greater length elsewhere.

anything. And—even more importantly—not all explanations are scientific. Indeed, both cosmological and fine-tuning arguments *start* from the claim that science cannot explain some striking fact. It would be odd if they then offered their own *scientific* explanations!

Methodological pluralism is especially prominent in contemporary moral philosophy. (I personally find it one of our discipline's most appealing features.) In its search for new methods, moral philosophy often engages more intimately with its own history than some other branches of contemporary philosophy.¹⁰ Contemporary moral philosophers constantly reinvigorate old ideas, approaches, and theories using modern analytic tools. Methodological pluralism leads us to take seriously questions and answers that were prominent in earlier philosophy, even if they are difficult to translate into contemporary analytic or scientific modes of thought. In particular, anyone exploring connections between God and morality is naturally drawn away from a present where both topics are philosophically marginal, and toward a past where their connection was perhaps *the* central philosophical question.¹¹

Several other kinds of pluralism feature in this book. Our discussion of mysticism in chapter 5 appeals to *doxastic pluralism*—the claim that there are several distinct, incommensurable, rational ways to form beliefs. Our argument from evil in chapter 8 appeals to *value pluralism*, where human values do not fit into a common scale, and are thus not fully commensurable. I devote much of chapter 9 to *religious pluralism*: one philosophical interpretation of the fact of religious diversity. Part III explores pluralism about *responses to value*: the thought that humans should respond to cosmic values in many different ways—including contemplation, study, and worship, as well as promotion.

A final pluralism concerns *individual intellectual commitments*. Many philosophical disagreements turn, not on purely factual or logical claims, but on divergent intuitions, considered judgements, basic orientations, philosophical sensibilities, or prejudices. Substantive philosophical disagreement often tracks some deeper disagreement about what is possible, valuable, plausible, simple, elegant, undeniable, or obvious.

Philosophers seek *decisive* intuitions that any acceptable theory must accommodate. But, in fact, most intuitions serve, not to establish or refute theories, but to *distinguish* them.¹² No intuition is uncontroversially decisive, if only because there is always a niche in the philosophical marketplace for the first person to reject it. An intuition that seems decisive to a theory's proponents often strikes its opponents as merely peculiar.

¹⁰ Consider the continuing influence, within normative ethics, of figures such as Aristotle, Hume, Mill, and Kant; and of historical works such as Schneewind, *The Invention of Autonomy*, or Irwin, *The Development of Ethics*.

¹¹ On the centrality of God and morality throughout medieval and early modern philosophy, see Schneewind, *The Invention of Autonomy*.

¹² Mulgan, *Future People*, pp. 2–4.

Some people are happy to accept that whatever science cannot explain is a *brute fact*. Science will eventually explain most things, and anything left over is simply inexplicable. Others feel, on the contrary, that the very existence of *any* universe (let alone one as remarkable as this) cries out for explanation, and that the search for such explanations is a central task of philosophy. If science cannot answer the basic questions of metaphysics, then we must look beyond science.

People similarly divide over the perennial philosophical programme that seeks to naturalize everything. Some take it for granted that human beings are *merely* physical animals who evolved via natural selection, and that *all* our features must be explicable in purely physical scientific terms. Our consciousness, rationality, freedom, and morality are nothing more than by-products of evolution. By definition, the best naturalist explanation must be adequate. Others think that reductionist projects inevitably leave out something that is vital to our nature; that a naturalized ethic is no ethic at all; and that there is something to the realm of normativity and value that cannot possibly be cashed out in naturalistic terms.

These are not disputes *within* science. They concern instead the limits and pretensions of science. Can science explain why there is a physical universe where life could evolve; why human beings can think; why torture is wrong; or why beauty is valuable? If not, must these things simply be rejected as illusions, or can they be accepted as brute facts, or do they instead cry out for explanation? The arguments that AP borrows from BT all appeal to the thought that there are some vital questions that science, operating in a purely atheist framework, simply cannot answer—and that these questions are worth asking.

AP thus responds to a common source of dissatisfaction with contemporary atheism. It also captures popular dissatisfaction with BT. Some people can accept that all the evil we see around us is part of the plan of a perfect loving God. They can believe and trust in such a God—and often see something impious or arrogant in human attempts to *explain* evil to our own satisfaction. But others find it obvious that the amount and distribution of evil in this world are far beyond what any *benevolent* God could permit. Indeed, many people find the very business of attempting to explain away or excuse the evils of the world tasteless or offensive. Not only don't they believe in God—they don't *want* there to be a God who would create a world such as ours. The arguments that AP borrows from atheism appeal to those who cannot imagine how any benevolent God could reconcile the evils of this world—those who find all extant theodicies unsatisfactory.

Atheism and BT each correspond to a natural package of pre-philosophical attitudes. Those who find science satisfying, and happily accept as brute fact whatever science leaves unexplained, often also find it impossible to believe in an evil-permitting God—partly because their naturalistic ethic leaves no room for anything that could outweigh the evils we see around us. This combination hangs together well. But so does the opposite combination. Those who can believe in a perfectly loving God who reconciles the evils of the world often also feel that both a naturalistic

ethic and a focus on the animal nature of human beings leave out something vital and that scientific explanations cannot be the whole story about the universe.

These pre-philosophical pictures explain why, here as elsewhere, philosophical argument so seldom results in anyone changing their mind. Philosophy is primarily apologetic rather than evangelical. Arguments serve to clarify the competing positions—bringing out the logical connections *within* each package.

The literature on *moral faith* offers two striking metaphors. One is the *leap of faith*: committing oneself to P despite one's realization that evidence and argument are insufficient to establish that P. This metaphor suggests one is initially agnostic, and then leaps to BT (or atheism). The recent literature often talks instead of *holding one's ground*: retaining one's commitment to P despite one's inability to establish that P.¹³ Here, one already believes in BT (or atheism) and merely seeks to rebut challenges to that view.

The distinction between these two metaphors matters, because it can be rational to stand firm where it would not be wise to leap. If I have built my life around a shared belief that P, this could be a compelling reason to hold fast to P, even if I could have formed an equally strong commitment to not-P. (Perhaps, at a pivotal time in my life, I happened to fall in with Protestants rather than Catholics, or with analytic logicians rather than postmodernists.)

The contemporary philosophical literature often takes the internal perspective of the religious believer, and then asks whether she has reason to abandon her faith. By contrast, I begin with an external agnostic perspective, and ask what reason there is to believe in BT or atheism or AP. This difference becomes especially significant in chapter 8, where it shifts our focus from *defences* of BT (accounts of God's possible reasons to permit evil) to *theodicies* (probable stories about God's actual reasons).

While the literature gestures at intermediate positions in logical space, atheist materialism and benevolent theism are the only *lived* philosophical positions in our time. AP is not a lived position. Belief in AP always requires a leap. But AP could still be the best way to stay put overall. If AP best explains the presence of *both* genuine value *and* genuine evil in the world, then those whose lives are built around the recognition of good and evil may need to posit AP if they are to resist the threat (or the lure) of nihilism.

The very fact that BT and atheism are so well suited to popular philosophical packages should make us suspicious. Perhaps they only dominate because of their attitudinal fit, and not from any philosophical superiority. As we shall soon see, proponents of AP diagnose a single moral failing in both traditional packages. They both rest on immodest premises, or on attitudes that overestimate the capacities and nature of human beings and give us an unwarranted role in the universe. Perhaps the traditional packages dominate precisely because, in their different ways, they serve

¹³ This is a prominent theme in the work of William Alston and Alvin Plantinga (see chapter 5 in this book).

our vanity—presenting humanity either as the culmination of the divine plan for creation or as the sole source of value in an otherwise meaningless cosmos.

On a more positive note, perhaps there is some other appealing combination of pre-philosophical attitudes that is represented by neither BT nor atheism. Surely one could feel that the very existence of the cosmos cries out for explanation; that the limits of science are not the limits of understanding; and that any naturalized ethic is unsatisfactory—while also finding all extant theodicies unconvincing. AP allows us to separate these different commitments and explore new combinations.

Unsurprisingly, the pre-philosophical attitudes I have just described are my own. While I can easily believe that the universe was designed—that it has a purpose—I cannot easily believe that it was designed by someone who cares for us. My aim in this book is simply to see where this plausible combination of pre-philosophical attitudes, which sits uneasily between BT and atheism, might lead. I do not claim that AP is the only way to accommodate some decisive intuition. But it does capture a reasonable and coherent picture of the cosmos and our place in it.

1.3 Morality and Metaphysics

AP is a metaphysical thesis: a claim about how the world is. In a book about metaphysics, many will feel that morality belongs only at the end. Perhaps a small epilogue will explore the moral implications of the metaphysical truth. But the metaphysical argument itself should be value-free.

This is perhaps the most common picture of relations between moral philosophy and metaphysics. Metaphysics is prior to morality. What we ought to do in the world depends on how the world is. This sounds fine in theory. In practice, however, it is less compelling. Once we recognize pervasive religious and philosophical ambiguity, we must abandon the idea that there is a single neutral objective account of ‘how things are’.

A second popular view reverses the priority between metaphysics and morality. A long BT tradition, associated most famously with Kant, holds that we have decisive practical reason to respond to religious ambiguity with a leap of faith. The basic questions of metaphysics cannot be answered by theoretical reason, and so we turn to practical reason. We can help ourselves to metaphysical claims that are necessary for morality—even if theoretical reason sheds no light on them whatever. (We consider Kant’s own leaps—freedom, God, and immortality—in chapter 10.)

Another response to religious ambiguity is to treat morality and metaphysics as completely independent domains. Moral philosophy seeks normative views that could be the subject of agreement among people with very different comprehensive metaphysical beliefs. (The most famous example, of course, is John Rawls’s political liberalism.) This strong independence claim is problematic even when confined to political justice. The search for a metaphysically neutral foundation for the whole of morality is even more ambitious, and hence even more problematic.

My own view is that the truth lies somewhere in-between these three positions. We cannot deduce our morality from our metaphysics or vice versa. But neither are the two completely independent. Sometimes a thesis in moral philosophy and a thesis in metaphysics are *mutually supporting*. While each has some independent plausibility, they also reinforce one another. Such mutual support is commonly sought within moral philosophy—and within metaphysics—so it is natural to seek it between the two. This book highlights one example of mutual support. It thus aims to bring contemporary moral theory and contemporary philosophy of religion into a closer dialogue. Both areas of philosophy have blossomed over the past thirty years. But they have largely done so in isolation. In particular, connections between consequentialist moral philosophy and contemporary philosophy of religion are under-explored. With some notable historical exceptions (such as William Paley), consequentialists have tended to be atheists (or at least agnostics), while BTs are invariably non-consequentialists. (Indeed, some of the most virulent recent opponents of consequentialism have been theists.) This book attempts to correct that imbalance. I therefore make no claims to provide a neutral or unbiased discussion.

Morality plays four roles in this book.

1. We might need AP to ground or explain moral facts. In the next chapter, I explore the BT argument that, without God, nothing could be good or evil, right or wrong; and I ask whether AP can borrow that argument.
2. Moral or evaluative claims often feature as premises in arguments for AP, even in arguments that seem purely metaphysical. Some of these premises are familiar, but others are surprising. The ubiquity of evaluative premises within metaphysical arguments is a central theme of this book and it recurs throughout parts I and II.
3. Religious ambiguity requires an ethic of belief. We need moral principles to guide us when non-normative considerations run out.
4. AP influences the *content* of morality. While AP supports some familiar moral ideals, it also pushes those ideals in a very unfamiliar direction. Adopting AP will not leave morality unchanged. Indeed, AP is so unfamiliar and so austere that one may worry that it will obliterate human morality altogether. This worry is addressed in part III.

All four roles depend on what our moral commitments actually are. My case for AP rests on my own moral views, some of which I outline in the next section. This may seem to drastically limit the book's relevance. Why should anyone else care? Why address only one person's moral commitments?

One reply is that this is my book. Where else could I start but from what I believe? As a moral philosopher, my strongest commitments are normative rather than logical or metaphysical. So that is where I begin. A full treatment of AP relative to every credible ethical view would simply be unmanageable.

If my moral commitments were idiosyncratic, the book would be of little interest. But, on the whole, they are not. While I sometimes appeal to very specific evaluative claims (about possible futures or components of human well-being, for instance), the central normative commitments invoked in this book are very general: the comparative value of suffering and freedom, the objectivity of value, moral realism, an austere picture of morality, and a broadly utilitarian ethic of belief. These views are not universal, but they are widespread.

This book develops only one possible case for AP driven by one specific starting point. But my broader aim is to demonstrate that anyone who takes morality seriously, and finds the existing alternatives unsatisfying, has good reason to consider AP. There are many possible routes to AP, just as there are many possible sources of dissatisfaction with both BT and atheism.

1.4 My Moral Commitments

One divide in contemporary philosophical discussion of morality is between those who study morality as an application of some other branch of philosophy, and those who explore our moral life from the inside. My approach is the latter. I begin with substantive moral questions, not with epistemology, philosophy of language, or metaphysics. In particular, I begin with substantive ethical commitments, rather than a commitment to some metaphysical world view. I take moral talk and moral life at face value and ask what metaphysical or epistemological picture does them justice, rather than requiring moral talk to fit some predetermined paradigm of legitimate philosophical discourse. This commitment to taking moral talk at face value may lead in surprising directions. At least, some have surprised me, and they have often dismayed my consequentialist colleagues.¹⁴

1.4.1 *Utilitarianism and suffering*

I operate within the utilitarian tradition, broadly construed. Any account of a tradition as rich and varied as utilitarianism is bound to be controversial. My aim here is not to offer a detailed historical exegesis, but rather to draw out some utilitarian themes. The defining feature of utilitarianism is that it measures everything (acts, rules, moral codes, social institutions, beliefs) in terms of the promotion of well-being. (Contemporary utilitarians often defend a broader *consequentialism* where morality is based on the promotion of values that may go beyond human well-being. This development is supported by AP, as we shall see in part III.)

¹⁴ This section draws freely on my own previous work, especially Mulgan, *Understanding Utilitarianism*; Mulgan, 'Mill for a Broken World'; Mulgan, 'Utilitarianism for a Broken World'; Mulgan, 'What is Good for the Distant Future?'; Mulgan, *Future People*; Mulgan, *Ethics for a Broken World*; and Mulgan, 'Ethics for Possible Futures'.

I begin with a very straightforward utilitarian commitment. Utilitarians emphasize the central ethical significance of suffering. As a result, they must take the problem of evil very seriously and cannot accept any traditional solution. In particular, utilitarians insist on the *comparative* importance of suffering. They reject the (as they see it) exaggerated notions of human freedom central to theodicy. Free will and self-development cannot justify the sufferings of this world. Faced with the bleak facts of human life, the utilitarian offers a bleak explanation. If there is a God, then God permits suffering because God does not care for individual human beings. This commitment plays a central role in chapter 8, the cornerstone of part II.¹⁵

The utilitarian emphasis on the significance of suffering is a double-edged sword. By linking objective value to a non-human-centred cosmic purpose, AP seems to imply that human suffering has no moral significance at all. Even the most intense agony is not bad at all, from the point of view of the universe. Utilitarians, along with everyone else, will wonder how anyone could endorse such a counter-intuitive claim.

I concede that AP's position on human suffering is very counter-intuitive. Perhaps it will ultimately prove too extreme for most utilitarians to accept. However, over the course of this book, I try to render this bullet easier to bite. One defence, of course, is comparative. All three competitors have counter-intuitive implications; it is only familiarity that makes BT and atheism *seem* more plausible. I also seek to reduce AP's moral counter-intuitiveness. In part III, I argue that AP can consistently agree that human suffering has great significance *for human beings*, even if it does not matter to God, and that this is sufficient to ground a recognizable (if very austere) human morality. The result is far from what we might initially expect, but it could represent a new reflective equilibrium: what we *would* think once we had adjusted our moral vision to a world governed by a non-human-centred God. Finally, as I'll now argue, utilitarianism has other commitments that pave the way to the transcendence of its own human-centred elements.

1.4.2 *A utilitarian ethic of belief*

Before outlining my other substantive moral commitments, I turn first to the ethics of belief. Moral philosophers should take the ethics of belief especially seriously. We naturally see every dispute as a question of what *to do*. So it is natural for us to see metaphysical disputes in this light. (And this should be especially natural for utilitarians, who bring all aspects of life before the tribunal of the utility principle.)

My utilitarian ethic of belief is built on two foundations: Mill's liberalism and Bentham's aversion to caprice. Straightforward application of Mill's famous utility principle to individual acts of *deciding what to believe* would yield a very simple principle: believe P if and only if believing P is what best maximizes human well-being. This is obviously too crude. I follow instead Mill's own liberal utilitarianism,

¹⁵ As we shall see in chapter 8, dialectical context is crucial here. It is as theodicy, not as defence, that utilitarians reject the standard BT story about free will.

which emphasizes the collective value of divisions of labour, and of each individual's following her own desires and values. Each individual life is, in Mill's words, part of an *experiment in living*.

I am especially interested in *moral experiments*. Some are practical. (Consider a landowner freeing his own slaves before emancipation, an early advocate of vegetarianism and animal rights, or a contemporary climate change activist.) But a moral experiment can also be speculative. In a work of fiction or philosophy, one might imagine a wider collective shift, and then offer that vision to one's fellow citizens as an inspiration for further experiments.¹⁶

This liberal utilitarian ethic of belief favours a variety of different moral experiments, a diversity of responses to religious and philosophical ambiguity. One response is *epistemicism*, where one only believes what is dictated by the evidence. The epistemicist response to ambiguity is agnosticism. This is one respectable response, but not the only one. There is room, in a pluralist Millian ethic of belief, for other responses.

Utilitarians will especially favour moral experiments that both promote well-being and also correct deficiencies in current commonsense morality. My second utilitarian foundation is a diagnosis of those deficiencies, based on Bentham's aversion to *caprice*. Utilitarians are suspicious of our natural human tendency to adopt interpretations of the world that favour ourselves or overestimate our own importance. They worry that current social structures and moral norms exist *because* they disproportionately serve the interests of the powerful. (After all, if you are not counting all interests equally, then you must be giving disproportionate weight to *someone*.)

This aversion to caprice inevitably takes a self-directed turn. The utilitarian thinker seeks out options than downplay her own interests—and also the broader interests of her own group, caste, class, nation, or even her own species. The history of utilitarianism is a constant tension between Mill's desire to create a private sphere safe from the incessant demands of morality, and the nagging Benthamite suspicion that this affluent safe haven is just another instance of indefensible caprice.

In the theoretical realm, caprice can arise both in our *epistemic* attitude to a proposition and in the proposition's *substantive* content. *Epistemic caprice* consists in overvaluing one's epistemic capacities or in placing undue faith in one's own cognitive resources. It is epistemic caprice to dismiss reports of mystical experience out of hand when one has never made any serious attempt to cultivate such experiences oneself; to assume that every question will eventually be answered by science, by philosophical argument, or by divine revelation (if one is, respectively, a scientist, a philosopher, or a theologian); or to take the success of empirical science to show that every significant question is amenable to the precision and tractability of

¹⁶ I attempt such a moral reimagining myself in *Ethics for a Broken World* and 'Ethics for Possible Futures'.

scientific explanation. Perhaps some very significant aspects of the cosmos are such that human beings can only glimpse them through a glass darkly.

Epistemic caprice is often combined with *substantive caprice*, where the favoured proposition has self-serving *content*. Consider the belief that one's own interests, perspective, or values carry more weight than those of others; or that there is something special about me, about my group, or about human beings in general. Of course, a belief in one's own superiority is not always capricious. You may have good reason to believe that you are better at philosophy than me or that human beings are more valuable than rocks. But suspicion is appropriate unless and until we have especially good reason to believe that we *are* special. Some cases are clear-cut: racism is capricious, while God's self-regard is not. Considerable room for reasonable disagreement lies between these extremes. One fertile source of contested cases is environmental philosophy. Do we have sufficient justification to place human interests above those of primates, or ecosystems, or species? Do human freedom, reason, or morality suffice to mark us out from the (merely) physical world around us?

Faced with religious ambiguity, or the ubiquitous fact that philosophical argument is never conclusive, we often have to take a leap of faith in one direction or another. My utilitarian ethic of belief cautions against self-aggrandizing leaps. We are naturally inclined to overestimate the case for our own significance. It does not follow that we are not significant, nor that it is irrational for us to believe so. An overestimated case can still be a sound one; and comforting self-appraisals can still be accurate. However, given our self-aggrandizing nature, we should always be open to discovering new and unexpected *kinds* of caprice and new sites of illicit self-aggrandizement. At the very least, any view that questions even our most cherished moral commitments on the grounds of caprice ought to receive a hearing. AP is one such view, and this book is an attempt to give it its day in court.

My aim, therefore, is to exaggerate the utilitarian suspicion of caprice, push it in unexpected directions, and follow where it leads. Many readers will feel that this stretches a legitimate ethical concern too far, especially when the collateral costs include many *other* fundamental moral commitments. I can only ask that such readers withhold judgement until the end of part III, and only reject AP once they have seen what positive moral vision (if any) it can build over the shattered ruins of our self-aggrandizing self-image.

If our total package of arguments leaves all three options on the table, then the choice between BT, atheism, and AP involves its own leap of faith. And here the least self-aggrandizing option, for human beings, is AP. Consider, first, the choice between AP and BT. Supposed we are convinced there is a God, but we are undecided between human-centred and non-human-centred interpretations of God's values, concerns, or purposes. Ex hypothesi, we have insufficient evidence or argument to support BT. (In particular, we lack compelling reason to believe that human beings possess a moral significance that would matter to God.) In this epistemic situation, it is simply caprice to leap to BT rather than AP. If we must make some leap of faith, then AP is

the only non-capricious option.¹⁷ (Of course, many theists will deny that this *is* our epistemic situation. But that is another matter!)

A leap to BT can thus seem capricious. What about atheism? Of course, most atheists don't accept that they have leaped *at all*. They indignantly insist that atheism is not a faith position, like Christianity or homeopathy, but merely a dispassionate inference from objective facts. But this is simply the denial of religious ambiguity. And the refusal to admit that one *has* leaped—the conflation of one's personal commitments with objective reality—is itself evidence of self-aggrandizing caprice.

If we acknowledge religious ambiguity, then atheism is a leap of faith. But is it a *capricious* leap? In stark contrast to BT's human-friendly cosmos, the atheist's world seems bleak, hostile, meaningless. What could be self-aggrandizing about the scientific picture of humans adrift in a vast uncaring universe? One answer lies in the dialectical fact that most arguments for atheism are actually only arguments against BT. Before one leaps to atheism, one must *first* leap to the *Human-Centred Conditional*: If there is a God, then God cares for human beings. (Otherwise, traditional 'atheist' arguments such as the argument from evil get us nowhere.) But this conditional is, if anything, *more* self-aggrandizing than BT itself. BT asserts that *God does care* for human beings, while the conditional insists that *any God must!*

To avoid caprice, the atheist must leap directly to a Godless universe, without recourse to the Human-Centred Conditional and in full recognition of the independent case for AP. But a leap to atheism *instead of AP* is a deliberate rejection of any non-human-centred source of objective values. The atheist insists that, if there is to be any value at all, then human beings must be its only source. This insistence itself is self-aggrandizing.

A Benthamite aversion to caprice thus cautions against any direct leap away from AP. It also operates *within* specific arguments, because many premises require their own leaps. For instance, BT responses to evil often need both personal immortality and incompatibilist freedom. For a defence, these need only be bare possibilities. But theodicy requires probability. At this point, evidence and argument run out. And AP objects that to insist that humans are metaphysically special in these ways is a kind of caprice.

Like any robust hermeneutic of suspicion, the utilitarian aversion to caprice eventually turns utilitarianism against itself. Does the utility principle itself give unjustified significance to *human* well-being? Utilitarians have long questioned the

¹⁷ This argument may look like double-counting. Doesn't caprice enter the equation twice? If our aversion to caprice has already been factored into our evaluation of the competing arguments for AP and BT, then why allow additional suspicion when we are choosing our leap of faith? AP has several possible answers here. One reply is that, if we want to ensure the avoidance of caprice, then we *should* be wary of it at every turn. A second reply is that, even if we prefer agnosticism, we still need an understanding of the alternatives between which we are undecided. This would suffice to motivate the full exploration of the resources of AP that this book seeks to inaugurate. Finally, while caprice *may* feature prominently in our case for religious ambiguity, it need not do so. In particular, perhaps caprice features in the case for AP only as a negative factor that rules out *leaps to BT* within the arguments for BT that AP borrows in part I.

anthropocentrism of privileging human beings *over other sentient animals*. But AP goes further. It denies that human beings matter *at all* from the point of view of the universe. Utilitarianism thus points to its own transcendence. The shift from utilitarianism to a broader consequentialism is a theme of recent work in the utilitarian tradition. Perhaps AP is just another, rather unexpected, step along that road. One obvious worry is that this is an *abandonment* of utilitarianism rather than a *reinterpretation*. This worry, which arose at the end of section 1.4.1 in relation to suffering, is addressed in part III.

Aversion to caprice plays a significant role throughout this book. While it arises within the utilitarian tradition, this Benthamite theme has broader resonance, as most ethical traditions caution against undue partiality or self-aggrandizement. Many ethics of belief will thus include some element that could play a similar role.

1.4.3 *Two pictures of morality*

Our ethic of belief complements a picture of morality that emerges from one strand within the utilitarian tradition. Much of my earlier work dealt with the demandingness objection to consequentialism: the complaint that consequentialism makes demands that are unreasonable, alienating, and psychologically impossible. Two contrasting reactions run through the voluminous literature on this topic.

Academic philosophers are well-paid, well-fed, affluent, upper middle-class professionals inhabiting a world of great inequality, need, and injustice. Any impartial moral theory will find our comfortable lives morally impermissible. Some philosophers find this situation deeply troubling; others do not. Some are distinctly uneasy living a life of luxury while innocent strangers starve; others feel no particular qualms about this. As a result, some regard demandingness as *the* central moral issue, while others dismiss it as a 'non-problem'.¹⁸

These two reactions to the demandingness debate highlight a broader divide between two pictures of morality, moral philosophy, and human nature, which I dub the *austere* and the *complacent*. On the complacent picture, our everyday commonsense morality is taken at face value. Our self-image is in order and there is certainly no need to give it up simply on the say-so of some philosopher. No plausible moral theory can demand more than we ordinarily expect. The task of moral philosophy is to explore our self-image, not to upset it.

On the austere picture, by contrast, morality is a source of demanding external reasons that compete with our various self-interested motivations. Human beings have a very strong tendency to confuse our self-interest with morality and to interpret moral norms and situations in ways that make life easy for ourselves. Our central moral task is thus to strive to be more impartial, to replace our parochial

¹⁸ For the former view, see Parfit, *On What Matters*, vol. 1, p. 501. The latter is harder to find in print, but has been the reaction of several prominent moral philosophers when I tell them I work on the demands of consequentialism.

viewpoints with the point of view of the universe, and to purge our moral thinking of self-serving delusions. The truth about our moral obligations and our place in the cosmos may turn out to be very alienating, demanding, or uncomfortable. Moral philosophy may be radically revisionary.

While the austere picture is especially associated with consequentialism, austere and complacent interpretations arise within most moral traditions. Consequentialists can be either extreme or moderate. Kantians can emphasize either Kant's pessimism about human nature and radical evil or his optimistic talk of human freedom and rationality. Rawlsian liberals can extend their egalitarianism across the globe, with radically unsettling results, or simply refuse to even speak of 'justice' outside the confines of their own affluent nation state. Christians, and other religious ethicists, can feel called to give all their wealth to the poor; but they can equally well regard their own affluence and others' destitution as morally unproblematic signs of God's favour.

Our utilitarian ethic of belief gives us solid reason to reject the complacent picture on the grounds of caprice. It is not hard to see why this picture appeals to affluent philosophers! In this project, I follow the austere picture. In a Millian spirit, I do not present the resulting austere morality as the only possibility, but as one credible moral experiment.

1.4.4 *Lessons from the future*

My substantive ethical commitments come from my own recent work on our obligations to future people. Utilitarians have long embraced *temporal impartiality*. Human welfare matters equally, no matter when it occurs. Given our potentially enormous impact on the welfare of future people, it is no surprise that, for the utilitarian, obligations to distant future people are of central moral concern. (By contrast, for non-utilitarians the future is usually at most an afterthought.)

In my recent book *Ethics for a Broken World*, I present another reason why utilitarians must focus on future people. Traditional ethical thought sets future people aside, in part because it assumes that they will be better off and that their interests do not conflict with our own. Faced with threats such as climate change, this optimism is no longer viable. We must confront the possibility of a broken future, where resources are insufficient to meet everyone's basic needs, where a chaotic climate makes life precarious, where each generation is worse off than the last, and where our affluent way of life is no longer an option. Given their commitment to promoting human welfare impartially, utilitarians must take this possibility especially seriously.

My most recent work imagines different possible futures and asks how future philosophers might reinterpret our own moral and political philosophy. I argue that reflection on possible futures transforms our moral thinking in many surprising ways. I thus draw on another utilitarian inheritance: Mill's belief in moral progress. This is not a naïve optimism about social progress, but rather a fallibilist admission

that future people will know more than ourselves about what is valuable and that future ethical inquiry might move in very surprising directions. We should be wary of projecting either optimistic empirical assumptions or controversial philosophical theories too far into the future. (This utilitarian openness to ethical revision comes to the fore in our discussions of moral naturalism, extraterrestrial life, and immortality in chapters 2, 7, and 10.)

For the present book, I draw one central lesson from the future: an adequate intergenerational morality must be founded on *objective values*. I have come to believe that, to make sense of our obligations to distant future people, and especially to give them adequate normative force, we need an objective list theory of well-being, together with objective values, and a metaphysically robust non-naturalist moral realism.

A full defence of this controversial claim would require a book of its own.¹⁹ I return to meta-ethics briefly in section 1.4.5 and at length in the next chapter. In this section, I briefly defend objectivism about well-being, value, and reasons. In each case, the best argument for objectivity is the inability of subjective accounts to do justice to our obligations to future people. (I return to these topics at greater length in part III, especially in chapter 12.)

I begin with well-being.²⁰ Contemporary debate contrasts three positions: *hedonism* (well-being is pleasure and the absence of pain); *preference theory* (well-being is getting what you want); and the *objective list theory (OLT)*, which offers a list of things that are good in themselves irrespective of the agent's attitude to them, such as knowledge, achievement, friendship, individuality, self-development, and so on.²¹

Objectivists argue that neither hedonism nor preference theory is satisfactory. Some pleasures are good, some bad, others are neutral. Some preferences improve your life, while others do not. Consider a child who wants to play in the sand rather than go to school. Clearly, we make his life go better if we send him to school. The challenge is to explain *why*. Education doesn't simply help people to satisfy their existing preferences. It also teaches them what to desire and which pleasures to seek. It is important to satisfy people's desires only *because* what they value is independently worthwhile. The objects are not valuable because they are desired—they are desired because they are valuable.

While debate between these three positions is ongoing, thinking about the future tips the balance in favour of OLT. Neither hedonism nor preference theory can capture our intergenerational obligations. Two familiar thought experiments help us to see why these subjective accounts fail.

¹⁹ I briefly defend the link between moral objectivism and intergenerational ethics in 'Ethics for Possible Futures' and 'What is Good for the Distant Future?'; and I plan to develop it at greater length elsewhere.

²⁰ For overviews of well-being, especially in the utilitarian tradition, see Griffin, *Well-being*; Parfit, *Reasons and Persons*, Appendix I; Mulgan, *Understanding Utilitarianism*, ch. 3; and the works cited in Mulgan, 'Consequentialism'.

²¹ Parfit, *Reasons and Persons*, Appendix I.

Nozick's Experience Machine. You are offered the chance to plug into an experience machine that perfectly simulates any possible human experience. Is it in your interests to do so?²²

The Virtual Future. In a broken future, people have abandoned the real world altogether and spend their entire lives plugged into an experience machine. The natural environment is so polluted and so resource-poor that people have little choice but to dream away their lives with no direct contact to any reality outside the machine. But this is all anyone has ever known and they find it perfectly satisfactory. If our present choices lead to this virtual future, have we done anything wrong?²³

Nozick's experience machine is often read as a decisive refutation of hedonism. Life in the machine is phenomenologically indistinguishable from the 'real thing'. If, as Nozick persuasively argues, it is a mistake to enter the machine, then there must be more to human flourishing than the quality of one's experiences. Our negative reaction to the virtual future supports this critique, as the hedonist must also find that world unobjectionable.

If Nozick refutes hedonism, our second tale rules out objectivism's other rival. The preference theorist cannot capture our unease about the imposition of any virtual future whose inhabitants are content with their lot. If we only look at individual preferences, then we cannot see what is wrong with avoiding our obligations to future people simply by manipulating their psychology or their environment so that they never want the good things we destroy.

By contrast, the objective list theory easily captures both Nozick's reaction to the experience machine and our reaction to the virtual future. If a connection to the natural world is intrinsically valuable, then human lives go better (and perhaps can only go well) when they instantiate that value. Some things matter, and it matters that people are connected to real values, not virtual ones. Even Peter Singer, the most prominent contemporary defender of preference utilitarianism, has recently acknowledged, on the basis of very similar examples, that we need a more objective account of well-being to make sense of our obligations to distant future people.²⁴

Unlike Nozick's original, and countless other fantastical tales, my virtual world is a *credible future*. This realization greatly strengthens the objectivist critique of hedonism and preference theory. In debates over well-being, as in many other philosophical topics, every theory stumbles over some ingenious imaginary case. Defenders of hedonism or preference theory can thus set aside Nozick's experience machine. No theory is perfect, after all.

We cannot demand that a theory of well-being must perfectly fit all our intuitions about imaginary cases. But we can reasonably insist that moral philosophy provide useful guidance about important actual decisions. An acceptable theory of well-being

²² Nozick, *Anarchy, State, and Utopia*, pp. 42–5.

²³ I discuss this example in Mulgan, 'Ethics for Possible Futures'.

²⁴ Singer, *Practical Ethics*, p. 244. See also de Lazari-Radek and Singer, *The Point of View of the Universe*, chs 8 and 9.

must help us to think clearly about our obligations regarding credible futures, especially when our present choices might harm future people.

Singer's own conversion is instructive here. As a practical ethicist, Singer focuses on first-order moral issues, such as abortion, our treatment of animals, or our obligations to the distant poor. His shift away from preference utilitarianism is driven by the failure of his own attempts to apply it to the newly urgent practical questions posed by climate change. The practical ethicist can sidestep the experience machine, but not the virtual future.

The further we look into the future, and the more that future might differ from our affluent present, the harder it is to believe that predictions about what people will (or might or could) desire have any real moral significance—or to believe that such predictions provide a solid foundation for morality.

The inadequacy of subjectivism goes deeper. Subjective stories about morality itself also cannot accommodate the future. Consider contractualism—the main contemporary rival to utilitarian intergenerational ethics. Reciprocity, sentiment, and mutual cooperation may provide good foundations for intra-generational ethics. But intergenerational contracts face two barriers: Parfit's non-identity problem and the impossibility of reciprocal interaction between present people and distant future people. How can we begin to imagine contracts, bargains, or cooperative schemes involving future people whose existence and identity depend upon what we decide and whose fate is entirely in our hands? By contrast, while utilitarians endlessly debate the precise details of our intergenerational obligations, they have no difficulty making sense of them. Obligations to future people are theoretically on a par with obligations to present people: both derive from the fact that our actions impact on the well-being of sentient beings. This doesn't prove that utilitarianism is superior all things considered, but it does significantly enhance its comparative appeal.

The need for objectivity has three further implications. First, once we acknowledge an objective account of human well-being, it is natural to posit objective values that are independent of human well-being. If knowledge makes my life go well, irrespective of my attitude to it, then knowledge must be good in itself. If it is good *for me* to achieve X, then X must be something that is independently worth achieving. Second, to avoid Parfit's non-identity problem, utilitarians themselves need *impersonal* judgements about the comparative value of possible futures. And these comparisons must also involve independent objective values. Finally, objectivity about well-being and impersonal values supports *externalism* about reasons. Internalists (such as Bernard Williams) insist that I only have a reason to X if X connects to my current motivations.²⁵ Externalists (such as Parfit) recognize reasons whose force is independent of my motivations.²⁶ OLT implies external prudential reasons, and independent impersonal objective values imply external moral reasons.

²⁵ Williams, 'Internal and External Reasons'.

²⁶ Parfit, *On What Matters*.

These brief remarks do not prove that there are objective values or external reasons. But they do explain why familiar utilitarian priorities commit one to such values. We urgently need a foundation for our obligations to distant future people, and objective values are, at present, the only game in town.

Objective list theory, independent impersonal values, and external reasons recur throughout this book. I defend these commitments at greater length in subsequent chapters, when they play specific roles in our detailed arguments. I now explore their general implications for our project.

One theme of this book is that objective value and cosmic purpose are mutually supporting. Part I explores the many crucial and unexpected roles that objective values can play in traditional arguments for BT. Here is one example. Objective value and cosmic purpose allow us to explain a number of puzzling general features of the cosmos, from the fact that there is something rather than nothing, to the fact that the universe is governed by precise elegant mathematical laws, to the fact that it is a place where conscious rational beings can emerge via processes of biological evolution. The atheist cannot explain any of these facts. He must regard them as brute facts, cosmic coincidences, just the way things happen to be. Without objective values, this brute fact response has some plausibility. If there is nothing special about the way things are, then why not admit that they just happen to be *this* way? (After all, things had to be *some* way.) But if things are an objectively special way, if the possibility that is realized is an unusually valuable one, then this does cry out for explanation.

Objective values thus support cosmic purpose. For its part, cosmic purpose supports objective values. For those suspicious of free-standing moral facts, cosmic purpose offers something to ground mind-independent values. Perhaps moral facts *are* facts about cosmic purpose. This brings us to meta-ethics.

1.4.5 Utilitarian meta-ethics

Utilitarians are split on meta-ethics. Utilitarianism is often associated with attempts to reconcile morality with naturalism—either by identifying moral properties with natural ones or by offering a non-cognitivist interpretation of moral talk.²⁷ This is a good fit for hedonists and preference utilitarians, whose accounts of well-being already look natural. But a second strand, going back to Henry Sidgwick and made especially prominent recently by Derek Parfit's *On What Matters*, combines consequentialist objective value with a commitment to non-natural sui generis moral facts. I myself favour this consequentialist non-naturalism, as it alone can ground the objective values that we need to underpin robust obligations to future people. Our utilitarian ethic of belief lends further support to a non-naturalist meta-ethic. Faced with meta-ethical uncertainty, utilitarians should favour options that enable them to

²⁷ The connection between utilitarianism and naturalism goes back to Bentham and Mill. Recent proponents include Jackson, *From Metaphysics to Ethics*, Hare, *Moral Thinking*, and Railton, 'Moral Realism'.

make sense of our obligations to future people. Our aversion to caprice will also play a role, as it prompts a suspicion of the cosy naturalist consensus that pervades contemporary meta-ethics. (I explore the connections between meta-ethics and AP at greater length in the next chapter.)

1.4.6 *An objective austere morality*

The moral commitments outlined in this section form a united whole, whose foundation is an austere picture of the demandingness and urgency of our obligations to future people. I claim that morality is objective in several distinct ways. Objective List Theory says that each person's well-being is largely a function of her connections to values that exist independently of her beliefs, wishes, projects, or pleasures. Objectivism about comparative evaluations says that these independent values include a ranking of different possible worlds. Externalism about reasons says that independent objective values provide the moral agent with reasons for action that are themselves independent of her desires, projects, or beliefs. Non-naturalist moral realism, explored at greater length in chapter 2, insists that there are facts about objective values and external reasons. It also claims that those facts are independent of the agent's own attitudes, distinct from any human moral practices, and not reducible to any natural facts.

My central moral claims are these: that my objectivist and externalist substantive claims are necessary components of any adequate intergenerational morality; that such a morality will strike human beings as very austere and demanding; and that only metaphysically robust moral realism can give us the motivation to follow such a demanding morality. This last claim is primarily psychological rather than logical or semantic or metaphysical. In philosophical debates about objectivism, its subjectivist opponents claim to borrow or mirror all the advantages of objectivism without its metaphysical extravagances. As it happens, I do not find subjectivism semantically credible. I regard objectivism as a much more natural interpretation of our everyday talk about well-being, value, and morality. (I defended this claim briefly in section 1.4.4, and I return to it in chapter 2.) But my central objection is motivational. Pleasures, preferences, internal reasons, and expressions of one's attitudes may prove sufficient to motivate commonsense morality's relatively undemanding obligations among contemporaries. But they cannot provide the impetus to sustain the sacrifices that utilitarianism demands in the face of a possibly broken future. For the hedonist, preference theorist, internalist, non-cognitivist, or naturalist, the sacrifices demanded by utilitarian morality must always seem imprudent, irrational, and extreme.

It is fair to criticize a rival metaphysical picture for failing to motivate people to save the world? Here, once again, we encounter reasonable disagreement. Secular liberals, naturalist philosophers, and moderate non-consequentialists invariably find this sort of criticism absurd. But many religious people—and many utilitarians—will find it both pertinent and decisive. My own position lies somewhere in-between. Perhaps no world view can motivate the moral demands of utilitarianism. Perhaps

this is too much to ask. But if one competing world view *can* ground, explain, and motivate the morality our future seems to demand, then this should count very strongly in favour of both that world view and the morality it supports. Nor is such support out of the question. After all, BT often does motivate extreme self-sacrifice in the service of its transcendent good, and AP can reasonably hope to borrow this motivational strength. At the very least, it is worth asking whether or not it can.

For the remainder of this book, I shall assume that morality *is* very objective, austere, and demanding, and ask what the world must be like if it is.

1.5 The Moral Implications of AP

The shift to non-human-centred cosmic purpose opens up new theoretical possibilities. On the one hand, AP differs significantly from BT. Theodicy and religious orthodoxy bring many commitments, both metaphysical and moral, which constrain the development of BT ethics. Freed from these constraints, AP can use the reality of cosmic purpose to ground normative theories hitherto neglected in BT ethics. On the other hand, the addition of cosmic purpose to an otherwise atheist world view gives moral philosophy new resources—new sources of information about what is valuable. If the universe exists because it is good, then learning how the universe *is* may tell us what is good. Another possible new source of moral insight is mystical experience, explored in chapter 5. In the final pages of his 1984 classic *Reasons and Persons*, Derek Parfit famously argued that non-religious ethics is in its infancy and it is not irrational to have high hopes for its future.²⁸ If so, shouldn't non-religious ethics be open to all possible sources of ethical insight—including those that arise within religion itself?

Any attempt to ground human morality in non-human-centred cosmic purpose faces obvious difficulties. The central challenge is to avoid *human* moral nihilism. If objective value is linked to a cosmic purpose that is indifferent to us, how does anything in our lives get any value at all? I address this challenge in part III. A second challenge is to defend human knowledge of morality. I explore this question at length in chapter 5. Existing moral resources (both BT and atheist) must be reimagined so that they offer genuine insight into the cosmic purpose, but seen through a human-tinted lens.

A final challenge is more existential. How do we bring a non-human-centred cosmic purpose into our lives? Why does the metaphysical dispute between atheism, BT, and AP matter? Suppose you shift from materialist atheism to cosmic purpose. Would that make any *practical* difference? (Or suppose you convert from BT to AP. Is this any different from converting to atheism? Can the new cosmic purpose play any of the moral roles that God once played? Can there be a non-human-centred

²⁸ Parfit, *Reasons and Persons*, p. 454.

religion?) These large questions will occupy us in part III, where I draw on a variety of models from both secular and religious ethics. Some roles for cosmic purpose are straightforward. If an accurate knowledge of one's place in the world is a component of a flourishing human life—as many atheists and theists alike agree—then we have good reason to explore any cosmic purpose. More ambitiously, cosmic purpose might offer a *model* of a valuable, creative life. Perhaps human lives can *resemble* value, even if they do not possess it. We can respond, perhaps in more or less valuable ways, to cosmic values that transcend us.

Any AP human morality will be strange, unfamiliar, austere, and demanding—more so than its atheist or theist competitors. But BT ethics also seems alien and extreme to secular eyes. And *any* moral realism is too demanding for the nihilist! This is only an objection to AP if we assume that morality should be comforting and familiar. But utilitarians, of all people, have no right to assume that.

1.6 Plan of the Book

My commitment to Normative AP structures this book. Our first task is to explore the moral motivation for AP, which sets the scene for all subsequent arguments. Accordingly, chapter 2 sets out and defends non-naturalist moral objectivism. It also develops a moral argument that we need cosmic purpose to ground objective values and external reasons. The chapter's key lesson is that AP and BT both agree that God (or cosmic purpose) has an ineliminable normative dimension. AP is therefore normative rather than morally agnostic, and it is closer to BT than to atheism.

Chapters 3 and 4 evaluate two very familiar arguments for BT: cosmological and teleological. I argue that AP can borrow the most plausible contemporary formulations of those arguments. In chapter 3, we learn more about both AP's normativity and its content. AP either presupposes or delivers comparative objective evaluations of possible states of affairs, and these include the judgement that the actual world is better than the empty world. We also encounter Leibniz's idea that God must create the best and therefore that this is the best of all possible worlds. AP is not wedded to Leibniz's maximalism. But it does breathe new life into it, for reasons that emerge more fully in chapters 7 and 8.

Chapter 4 surveys a range of teleological arguments, and concentrates on three prime candidates for AP to borrow. The most important of these is the fine-tuning argument, which starts from the observation that a range of cosmological constants appear to be remarkably fine-tuned to facilitate the emergence of life within the physical universe. In this chapter we learn that cosmic purpose involves general features of the cosmos, such as its friendliness to life, its governance by mathematical laws, and the fact that it can be understood by some of its inhabitants. These content-based lessons further reinforce Normative AP. Unless we agree that these features are objectively significant, we cannot hope to ground any successful fine-tuning argument for cosmic purpose.

Chapters 3 and 4 also teach us that some empirical questions are more significant to cosmic purpose (and, therefore, to human morality) than they might initially appear. Moral philosophers can no longer ignore cosmology! We also see that AP favours some empirical hypotheses over others, a lesson that is reinforced in chapter 7.

Cosmological and teleological arguments are explanatory. God is posited as an essential component of the best explanation of empirical facts about the existence of the physical universe or its friendliness to the natural emergence of life. Chapter 5 introduces a different style of argument, based on William Alston's defence of the internal rationality of Christian mystical doxastic practices. I contend that many philosophers make out-of-date assumptions about the nature of mystical experience, and this leads them to underestimate its moral significance. I then argue that AP has much to learn from BT mysticism. Chapter 5 also expands our knowledge of the content of cosmic purpose and the role of human experience in both morality and mysticism. We learn that the cosmic purpose is characterized by unity, transcendence, and non-self-centredness. Mystical experience seems to be non-moral. However, I argue that either mystical experience *is* moral experience (because moral facts *are* supernatural facts), or it only constitutes reliable *evidence* for cosmic purpose because it is linked to human moral improvement or insight. Any successful argument from mysticism must yield Normative AP rather than Simple AP. (Chapter 5 leaves one objection to mysticism unanswered: the threat of religious diversity. We pick up that challenge in chapter 9.)

Chapter 6 deals with the third classic BT argument: Anselm's notorious ontological argument. I argue that, read in its original religious and philosophical context, Anselm's *Proslogion* has many resources that the a-contextual readings of contemporary analytic philosophy tend to miss. This chapter explores the source of cosmic purpose, arguing that AP should posit a perfect being who gives the universe its cosmic purpose. This perfect being could, in principle, be either personal or impersonal. For convenience, I follow BT and speak of a perfect God. We learn that AP can legitimately borrow perfect being theology and the philosophers' Omni-God, and that all talk about a perfect being is ineliminably normative. Chapter 6 concludes that any successful ontological argument must link perfection to objective values via human mystical experience. (Chapters 5 and 6 thus support each other.)

In part II, we turn to arguments against a human-centred cosmic purpose. We suppose that the arguments of part I deliver a cosmic purpose, and we now seek further information about that purpose. Chapter 7 investigates a comparatively unfamiliar objection to BT, albeit one based on a very familiar intuition. The argument from scale objects that the vast universe discovered by science, where human beings play such a marginal role, is out of kilter with what we would expect from a God who was interested in us. After rejecting several unsatisfactory formulations, I conclude that the most plausible argument from scale concerns, not the size

of the cosmos per se, but rather the number, diversity, and sophistication of its inhabitants. Simply put: humans are too puny to be the centrepiece of so vast a cosmic canvas.

Chapter 7 also contains general lessons about God's creative reasons. I argue that the arguments of part I suggest a God who has impersonal consequentialist motivations. God responds to the comparative values of possible worlds and creates the best world. This divine consequentialism constrains BT's interpretation of God's benevolence. If God cares for finite creatures at all, God must create the best creatures. Finally, chapter 7 teaches us some surprising things about the cosmic significance of extraterrestrial life. AP expects God to create beings who are much better than humans and rejects any presumption that we are alone.

Chapter 7 concludes that, while God would care about superior beings, we have no good reason to believe that God cares for us. Chapters 8 to 10 complete the case for AP by arguing that we have reason to believe that God does *not* care for us. The main focus in these chapters is on controversial metaphysical claims that BT needs and AP avoids (such as libertarian freedom and personal immortality), and on BT's inability to adequately explain facts about evil and religious diversity. My arguments are often disjunctive. AP offers several related criticisms of BT and need only argue that at least one succeeds.

Chapter 8 examines the most popular atheist objection to BT: the argument from evil. Unsurprisingly, this is where our utilitarian commitments come to the fore. I present two arguments from evil that AP can borrow—based on animal suffering and the horrendous evils that humans inflict on one another. Much of the chapter is a sustained rebuttal of theodicies built around human freedom. I argue that BT must make very ambitious claims about actual human freedom—claims that AP can reasonably reject. The main lesson of chapter 8 is negative: God does not care about individual human beings at all. However, we also learn that AP seeks to remain agnostic regarding two related claims: that some things are good for humans and that human well-being has some objective value. We return to those topics in chapter 12.

Chapter 9 picks up where chapter 5 finished, asking whether the fact of religious diversity refutes BT. A separate chapter on religious diversity may seem redundant. Isn't this just another minor example of evil? (If BT can explain why God permits horrendous evil or vast animal suffering, surely it won't be troubled by a few religious disagreements?) I argue instead that religious diversity is a distinct challenge for BT, because it represents an unequal distribution of what BT itself regards as *the* central human good: right relationship to God. Chapter 9 concludes that AP can borrow atheist objections based on religious diversity, although we must then modify Chapter 5's claims about the significance of mysticism. AP can only borrow insights from competing mystical traditions when they overlap. Because mystics agree about abstract moral ideals rather than specific metaphysical or religious doctrines, this reinforces Chapter 5's emphasis on the moral dimension of mysticism.

makes no sense in a purely materialist universe, then perhaps we must move beyond materialism. Under BT or AP, value is no longer an isolated anomaly. It is built into the very fabric of the cosmos. This need not commit us to God, let alone to any specific religion. But it does get us surprisingly close to cosmic purpose.

This moral argument is not sufficient on its own. Nor is it essential to AP's case against atheism. One could find the arguments of part I persuasive without agreeing that we need cosmic purpose to explain objective values. (Indeed, as we'll see, the arguments of chapters 3 and 4 might be *most* appealing to those who *can* countenance independent non-natural values.) Nonetheless, a moral argument can form part of the cumulative case for AP.

2.1 Two Approaches to Meta-ethics

As I explained in chapter 1, my starting point is normative ethics, not metaphysics. I begin with the moral commitments set out in section 1.4, especially regarding religious ambiguity and objective values, and then ask what meta-ethical position we need to make sense of those commitments. This is *not* the standard approach among contemporary meta-ethicists, who typically start with some already privileged meta-physical picture of the world, and then ask where (if at all) our moral talk might fit into that picture.

Most meta-ethicists begin with a presumption of *global naturalism*, where 'the only kinds of things whose existence we ought to countenance are things that fit into a unified scientific framework'.¹ Contemporary meta-ethics is thus driven by a dilemma that Frank Jackson dubs the 'location problem'.² When we say that things are good or bad, right or wrong, forbidden or permissible, we make claims that go beyond a mere description of human practices. Our moral talk seems to demand moral *facts*. But where, in the naturalistic world, are moral facts to be found? And where do they get their authority over the agent and her desires?

The global naturalist has only three broad meta-ethical options:

1. Moral statements do not express propositions. (Non-cognitivism)
2. Moral statements express propositions that are reducible to propositions about natural properties. (Moral naturalism)
3. Moral statements are all false. (Moral nihilism)

For those who wish to avoid moral nihilism, the only alternatives are non-cognitivism and moral naturalism. But both are semantically implausible. It is *prima facie* obvious that, when we use moral language, we do mean to make claims that can be true or false, and we do not take ourselves to be talking about purely

¹ Joyce, *The Evolution of Morality*, p. 145. Other definitions of naturalism are explored in chapters 3, 4, and 5.

² Jackson, *From Metaphysics to Ethics*, pp. 113–38.

ambiguity. I seek a plausible account of the meaning of moral talk. From this perspective, semantic considerations have priority over ontological ones. The semantic debate between cognitivists and non-cognitivists is empirical. Their competing claims concern the current function of moral talk in actual human languages. And the evidence points against non-cognitivism. Huemer puts the point well.

The non-cognitivist makes a claim about actual meanings, but all of the reasonably direct and objective evidence, from the way we use ethical language, points towards cognitivism.... Evaluative statements act in every way like factual claims.⁵

For a taste of the difficulties here, consider how the non-cognitivist might analyse the following utterances:⁶

- If pleasure is good, then chocolate is good.
- It is better to achieve something worthwhile than to fulfil your desires.
- I wonder whether I did the right thing.
- If murder is wrong, then I did the right thing.
- Did you try to do the right thing?
- It is true that pain is bad.
- Is euthanasia wrong?
- Do the right thing!
- Some things are valuable.
- Morality demands great sacrifices on behalf of future generations.
- If God commands X, then X is required.
- If God does not exist, then nothing is good.

No doubt some non-cognitivists can accommodate some of these linguistic phenomena. But it is hard to see how *any* non-cognitivist could account for them all *as easily as the cognitivist can*. (As we'll see, non-cognitivists are especially troubled by moral objectivity, a difficulty they share with moral naturalists.)

The semantic evidence places the non-cognitivist on the back foot. The non-cognitivist saves moral talk, but only by changing the subject. Such reinterpretation is necessarily a second-best option. But without a prior commitment to global naturalism, we have no need to settle for second best. Given its semantic inadequacies, non-cognitivism is best seen as *revisionary* rather than descriptive. It is not an account of what moral talk *has meant*, but rather instead a suggested reinterpretation whereby moral talk remains useful despite the non-existence of moral facts. (Just as religious talk might be radically reinterpreted by someone who no longer believes in the existence of God.) Revisionary non-cognitivism is a *form* of moral nihilism, and not an alternative to it. (We return to moral nihilism in section 2.4.)

The unspoken atheism of contemporary analytic philosophy is nowhere more evident than here. Many non-cognitivists reject moral facts as logically or

⁵ Huemer, *Ethical Intuitionism*, p. 25.

⁶ Adapted from Huemer, *Ethical Intuitionism*, p. 23.

metaphysically incoherent. Indeed, the *obvious* impossibility of moral facts is often seen as semantic evidence for non-cognitivism. Any informed person can see that there could be no moral facts. So no sensible moral talker could ever mean to refer to such facts. As an argument about actual human moral talk, this is unconvincing—for the simple reason that most human moral talkers are not global naturalists. In particular, many are theists. If there is a God, then there might well be moral facts. The incoherence of moral facts is not obvious, *unless* the non-existence of God is also obvious. In a religiously ambiguous universe, moral facts are not especially anomalous. Non-cognitivism has a place in atheist apologetics. But as an ontologically neutral account of moral semantics, it is a non-starter.

Two features of our project count against non-cognitivism. The first is our rejection of its presumption of global naturalism. The second is our commitment to moral objectivism and an austere demanding morality. Moral talk is puzzling because it combines factual and motivational elements. Moral statements both express our moral beliefs and motivate us to act. Non-cognitivism claims to capture the motivational aspect of our moral talk better than cognitivism, because it connects a person's moral utterances directly to her motivations. This makes sense if we presuppose subjectivism about morality and internalism about reasons. But what if the most urgent moral talk explicitly aims to lead us beyond any actual human motivation?

In chapter 1, I introduced a package of moral commitments with several distinct objective elements: an objective list theory of well-being, mind-independent comparative evaluations of possible states of affairs, external reasons, and a very demanding consequentialist moral theory. By contrast, non-cognitivism seems more at home with a subjective complacent morality, where the link between my present attitudes and my moral obligations is straightforward and there is no need for moral motivations beyond my actual current desires. In that undemanding context, perhaps non-cognitivists can talk the moral talk (albeit somewhat artificially). But when it comes to expressing the demands of the future, non-cognitivism comes up short. Non-cognitivist paraphrases may suffice if we only want to use moral propositions as premises in abstract reasoning. But they lack the external normative push that this project's substantive moral commitments demand. Or so it seems to me.

My primary complaint is not that non-cognitivism cannot construct intersubjective standards to evaluate the correctness of moral utterances, but that it cannot do justice to the existence of mind-independent intrinsically motivating objective values and the external reasons they generate. This motivational deficiency goes hand in hand with non-cognitivism's ontological narrowness. Contemporary non-cognitivism claims to capture (or mimic) the objectivity of realist moral talk. As I suggested earlier in this section, this seems implausible if we include the everyday moral talk of *theist* moral realists. But even if non-cognitivism could mimic the logical and semantic features of realist moral talk, it would surely lack the urgent normative push provided by real

If 'good' means 'N', then no native speaker would ask this question. (Just as no one who understands 'unmarried man' and 'bachelor' can wonder whether this particular unmarried man is a bachelor.) But this question *is* open. It always makes perfect sense, given *any* natural description of an action or event, to ask whether it is good. This shows, says Moore, that no natural definition captures the meaning of 'good'.

Moore's OQA is semantic, and therefore empirical. The definitional moral naturalist makes claims about the current role of moral talk, which Moore seeks to rebut. Moore appeals to his own linguistic intuitions and those of his readers. While this does make the OQA 'intuitionist' in one sense, it is *not* intuitionist in a more familiar (and much more controversial) sense. The OQA appeals to linguistic intuitions, not moral ones. Its conclusion is not that there *are* moral facts, or that human beings *do* have intuitive access to those facts. While Moore himself did hold these further positions, the OQA itself merely concludes that moral terms do not designate natural properties. (We return to ethical intuitionism in section 2.5.)

The OQA is on firm ground against any specific reductive analysis. Consider a utilitarian claim that 'right' means 'maximizes human welfare'.⁹ Native speakers can see that 'right' and 'maximizes human welfare' are not synonymous. It is absurd to accuse non-utilitarians of misusing moral language, or failing to speak English properly. And it is surely significant that the OQA has succeeded against every other extant reductive analysis.

But is Moore entitled to the more ambitious claim that no natural definition is *possible*? Can't the naturalist simply reply that a successful analysis will be forthcoming in the future? Why not treat the OQA as a spur to further refinement rather than a knockdown objection to an entire research programme?

A more ambitious OQA argument could simply be *inductive*. Every past natural analysis has failed the open question test. Therefore, every future analysis will fail. Given the number of analyses philosophers have produced, this induction would carry some weight. (It is certainly more compelling now than it was in 1903.) But, as with any induction from past failure, it is not decisive. Fortunately, there is a stronger Moorean argument available, based on a direct appeal to linguistic intuition. When native speakers reflect on the failures of particular analyses, as revealed by the original OQA, they can *see* that no such analysis could ever succeed. The OQA helps us to realize that moral talk is not a shorthand for some very complex natural talk, but instead refers to distinct moral properties that cannot be analysed in natural terms. As we shall soon see, this doesn't show that moral properties are not natural properties. But it does strongly suggest that moral terms are not synonymous with natural terms. Like non-cognitivism, definitional moral naturalism is a second-best option. And it is too soon in our exploration of meta-ethics to settle for

⁹ Utilitarians need not adopt this definitional claim. There are non-naturalist utilitarians. (Indeed, this book is proof of that!) There is always a distinction between the *meta-ethical* claim that rightness *is* X, and the substantive normative ethical claim that whatever is X is right.

second best. (In addition, the arguments deployed against non-definitional naturalism in the next section would also defeat definitional naturalism.)

2.3.2 *Non-definitional moral naturalism*

Many moral naturalists agree that the OQA defeats *definitional* naturalism. Moral talk is not analytically reducible to natural talk. Moral terms and natural terms are not synonymous. But this does not mean that moral properties are not, in fact, identical to natural properties.

According to non-definitional moral naturalism, the identity between a moral fact and a natural fact is discovered a posteriori. The model is the analysis of natural kind terms offered by Kripke and Putnam.¹⁰ Water is H₂O. This is not a definitional truth. ‘Water’ does not *mean* the same as ‘H₂O’. Similarly, ‘good’ does not *mean* the same as any natural property term. Moore’s open questions are not semantic questions, but substantive questions in chemistry or moral philosophy. To discover that water is H₂O, we examine the *substance* water—not the *word* ‘water’. Similarly, Richard Boyd, a prominent recent proponent of non-definitional moral naturalism, advocates the following procedure for moral terms.¹¹ We first list the features associated with our idea of goodness, rightness, or any other basic moral property. ‘Good’ refers to whatever property *actually* has these features. Assuming that all properties are natural, we then examine the natural world, and (hopefully) find the property that has those features. Call it N. We conclude that ‘good’ refers to N. Like many moral naturalists, Boyd is a utilitarian. He conjectures that, if we were to carry this project to fruition, we would find that goodness *is* maximum human well-being.

The most obvious weakness of the analogy between water and goodness is that the identity with H₂O has already been established beyond reasonable chemical doubt, as have many other scientific accounts of natural kinds. There is no moral equivalent. Non-definitional moral naturalists gesture at such accounts, and some claim to have produced them. But no one has actually produced an identity that has anything like the authority of ‘water is H₂O’.

Early chemists wondered whether water was H₂O. But this question is no longer *scientifically* open. No informed person could now wonder whether this sample of normal water is H₂O. Anyone who does wonder merely reveals his own ignorance. No moral question has ever been closed like this. There is no putative moral identity whose questioning would cast doubt on the speaker’s basic competence—whether linguistic or empirical. ‘How can you possibly doubt that this action which maximizes human well-being is right? Haven’t you read Mill/Hare/Boyd/Railton/Jackson? Don’t you know anything?’ Unlike the equivalent response in chemistry, this is just intellectual bullying.

¹⁰ Kripke, *Naming and Necessity*; Putnam, ‘The Meaning of “Meaning”’.

¹¹ Boyd, ‘How to be a Moral Realist’; Boyd, ‘Finite Beings, Finite Goods’. (Boyd calls his position ‘ontological’ moral naturalism.)

image

not

available

3.5 Explaining S

If S is not a brute fact, then S has an explanation. But what could possibly explain why there is anything at all? I borrow a useful notion from Parfit.³³ A *selector* is a feature that explains why x exists. Its utility for nail-hitting is a selector for my hammer, because its having this feature explains why my hammer exists. A selector for the universe is a feature of the universe that explains why it exists. To posit a selector just is to deny that S is a brute fact.

Parfit's notion of a selector is perhaps best illustrated using God. Anyone teaching the cosmological argument encounters the question: 'Yes, but who made God?' Any answer posits a selector for God. If that selector had to be a person or process who made God, then BT would be doomed to regress. Some BTs, like Swinburne, deny that God's existence has any further explanation. God is the ultimate brute fact. But other BTs do offer explanations for God. They say things like this: 'God is a perfect being, and God exists because God is a perfect being.' Here, *being perfect* is a selector for God.

Any theory that explains S is superior to one that does not *other things being equal*. But other things are seldom equal. Any explanation for S (especially one that invokes or implies a cosmic purpose) has metaphysical costs. Even if we agree that God would explain S (and perhaps that nothing else could), we might still balk at such metaphysical extravagance.

Opponents of BT often appeal to *parsimony*. We should not posit entities beyond necessity or seek explanations where none is available. Parsimony is a good principle. We should *seek* the least extravagant hypothesis. But parsimony is not the only explanatory virtue. It competes with *explanatory power*. In our present context, any hypothesis has explanatory power if it offers a better explanation of S than the null hypothesis that S is a brute fact.

The balance between parsimony and explanatory power depends on the urgency of the ultimate question. If we endorse PSR, then we may be willing to pay a very extravagant metaphysical price to explain S. At the other extreme, if our interest in S is mere idle curiosity, then we may privilege parsimony. As we have seen, our commitments in this project place us somewhere in-between. The ultimate question is significant but not inescapable. So we must balance parsimony and explanatory power. The most parsimonious explanation of S would invoke only familiar entities and modes of explanation. If a purposeless explanation for S were available—or even likely to be forthcoming—then the cosmological argument would collapse. But we have already seen that familiar explanations fail. S is so different from any fact that science has hitherto explained that there is no reason to expect any scientific explanation—and especially not to expect a *purposeless* one. Given our commitments, it is not immodest

³³ Parfit, 'Why Does the Universe Exist?'; Parfit, 'Why Anything? Why This?'

to ask our question, and not unreasonable to explore answers that are both non-scientific and purposive.

Science and everyday life provide our background of familiar entities and familiar modes of explanation. If we cannot explain S using only familiar entities and modes of explanation, then every explanation carries a metaphysical cost. We can evaluate the metaphysical costs of competing explanations of S along two dimensions: new ontological commitments and new modes of explanation. I will explore two purposive explanations of S: BT and axiarchism.

The BT explanation of S is familiar. God created the universe. Recall Swinburne's definition: 'God is a necessarily existing person without a body who necessarily is eternal, perfectly free, omnipotent, omniscient, perfectly good, and the creator of all things.'³⁴ God stands outside the physical universe, and God clearly has the power to bring it into existence. The main question for BT is why God would choose to create. Here, BT explanations fall into two broad camps. The *intellectualist* says that it is better for a physical universe to exist than not. This fact is God's reason to create. Being omniscient, morally perfect, and omnipotent, God will know, be motivated by, and act on this reason. Therefore, God will create. We know God has this reason, because we can discover the objective values that provide God's moral reasons. On the *voluntarist* view, by contrast, God's creative act is itself the source of values and reasons in our cosmos. God simply chooses to create a physical universe. There is no further explanation.

As I said in chapter 2, while I am committed to objective values, I aim to remain neutral between voluntarism and intellectualism. When it comes to constructing a human morality, objective values that flow from God's decision to create are just as good as ontologically independent values.

BT introduces one new entity (God) and one new mode of explanation (divine creation). Whether BT introduces a new *kind* of entity or explanation depends on our background assumptions about human persons and objective values. Is God another person like us, or something entirely new? Are we already committed to objective values, or are these also new?

Axiarchism is much less familiar than BT. The axiarchic explanation of S is very simple.

1. The physical universe is better than nothing.
2. The physical universe exists because of (1).

For BT, (1) might capture God's reason to create. In axiarchism, comparative value is itself the *direct* reason why the universe exists. The physical universe exists *because it is good* period. Axiarchism can be regarded as either introducing a new mode of explanation using familiar entities (objective values), or as introducing a radically

³⁴ Swinburne, *The Existence of God*, p. 7.

new kind of entity: directly efficacious values. Like the intellectualist, the axiarchist must defend *independent* objective moral values. (By contrast, the voluntarist offers to explain the emergence of value itself. Voluntarism thus provides BT with distinctive resources that both intellectualism and axiarchism lack.)

We thus have three explanations for S: axiarchism, intellectualist BT, and voluntarist BT. Are there alternative explanations?

On one level, there are innumerable alternatives. As Hume famously pointed out, an explanation of the cosmos based on a person or persons of less than infinite power would seem more familiar.³⁵ Moving beyond persons, any number of impersonal creators might explain S. And each possible creator offers both intellectualist and voluntarist explanations, as they may either respond to external objective values, or create on the basis of their own whims or internal reasons. Finally, Parfit notes that goodness competes with a myriad of *other* direct selectors. Just as the cosmos might exist because it is good, it might also exist because it is governed by beautiful mathematical laws, or because it is 'ontologically maximal'.³⁶

Fortunately, we can reasonably treat our three purposive explanations as representative of this much larger set of possible explanations. Any explanation of S is either direct or indirect. Axiarchism is a direct explanation: the cosmos has property P, and this is why it exists. There is no intermediary linking goodness and existence. There is nothing more to be said about *why* the best possible world exists. An indirect explanation posits some mechanism, intermediary force, or creative agent to bridge the gap between goodness (or some other property) and existence. We might say that a direct explanation involves a *selector* (in Parfit's sense), while indirect explanation involves a *creator*.

As our primary interest is in cosmic purpose, we focus on two questions: (1) Does the explanation yield a purpose? (2) Can AP borrow that purpose? It may seem that AP cannot remain too agnostic. We can imagine selectors without purpose, and creators without value. Surely these possible explanations threaten AP. But things are not so simple. Our vocabulary of 'purposes' might itself be an illegitimately anthropocentric leap. Selection needs no purpose in the sense that *we* have purposes, just as it needs no personal creator *who is like us*. 'Purpose' is a metaphor, as is 'selector'. The *real* question is whether our explanation of S warrants talk of values and reasons. I argue in section 3.6.4 that we can legitimately assimilate all direct explanations to axiarchism, because any selector will ground objective values; while section 3.7.2 argues that AP can borrow any plausible creator. In part III, I argue that any explanation involving objective values will yield something practically equivalent to a cosmic purpose, as will any creator who is both (a) powerful enough to create a physical universe and (b) not obviously morally deficient. Therefore, any explanation for S yields cosmic purpose.

³⁵ Hume, *Dialogues Concerning Natural Religion*.

³⁶ Parfit, 'Why Anything? Why This?'

Index

- Adams, M. 234
- Adams, R. [41](#), 56–7, 180, 202–6, 226–7, 343, 373, 399, 407, 410
- afterlife 280–322
- and AP 280–1, 299, 303, 305–7, 309, 340–2
 - and atheism 280, 291, 295–6, 299, 307, 309–10, 317, 340, 342
 - and BT 281–90
 - animals 233, 281, 283, 290, 309, 311–13
 - evidence for 290–3
 - and infinity 211, 252, 282, 283, 308, 312, 315, 342; *see also* infinite utility; infinite value
 - moral argument 281, 304, 307–22, 340
 - and personal identity 293–303
- Alston, W. 130, 131, 135–51, 155, 156–7, 165, 182–3, 187, 261, 266, 271, 277–8, 331–2
- ambiguity
- philosophical (defined [9](#)) [9–13](#), [17](#), [18](#), 50, 163, 269, 306, 353
 - religious (defined [9](#)) [8–9](#), [13](#), [14](#), [17–19](#), [37](#), 42–3, 50, 56, 59, 69, 78, 163, 218, 266, 269, 273–4, 309
- ananthropocentric purposivism (AP)
- and afterlife 280–1, 299, 303, 305–7, 309, 340–2
 - and argument from
 - evil 88, 220–5, 337–40, 364, 367
 - scale 195–7, 200, 213–19, 328–31
 - and atheist consequentialism 374–403
 - and BT ethics 392, 404–15
 - and Buddhism 267–8, 279
 - and Christian mystical practice 157–60
 - contrast with BT [2–4](#), [11](#), 60–1, 215–17, 329
 - and cosmic features argument 108–9
 - and cosmological argument [82](#), 83, 88, 91, 95–7, 325–6
 - defined [1–5](#)
 - and divine command ethics 405–8
 - and divine motivation theory 412–14
 - and extraterrestrials 200, 208, 211, 213–19, 231, 236, 269, 329–31, 339, 365, 395–6
 - and fine-tuning argument 128–9, 326–8
 - and God [3–4](#), [7](#), [18](#), 61–2, 96–7, 104, 108, 157–9, 187–8, 299–300, 320, 357–60, 367, 404–5, 410
 - and human features argument 104–5
 - and human well-being [16](#), 212–13, 224, 251–2, 258, 338–9, 343–72, 375, 385–6, 388, 392–3
 - leap of faith [18–19](#), 309, 345, 347, 351–4, 367
 - and the meaning of life 309, 316–22, 340–2, 361, 381, 385, 388, 394
 - and meta-ethics 59–62, 155, 307, 368–9, 382
 - moral argument 33–4, 60–2, 85, 348, 368–9
 - moral implications of [27–8](#), 156, 203, 243, 311, 334, 337–8, 378
 - and moral knowledge 61–2, 156, 158, 279
 - and moral theory 373–415
 - and mysticism 131–2, 138–40, 156–60, 183, 187–9, 279, 331–3, 414–16
 - normative (defined [7](#)) [7–8](#), [28–9](#), 79, 98, 105, 109, 113, 131, 159, 163, 189, 196, 309, 318, 345
 - and objective list theory 344, 360–6, 370–2, 381
 - and ontological argument 174–5, 181, 183, 187–9, 267, 334–7
 - and religious diversity 146, 156, 183, 186, 264–8, 277–9, 319, 331–3, 414
 - and religious pluralism 276–7, 279
 - and teleological arguments 100–1
 - and value-based morality 278–80
- Anselm 161–76, 179–89, 334–7, 343, 373, 406, 408–9
- on divine incomprehensibility 168–70, 173, 179
 - on the fool 168–72, 180, 335
 - on moral psychology 406, 408–9
 - and Neoplatonism 165–6, 172, 174, 177, 180, 406
 - as rational mystic 170–4
 - on thinking and understanding 168–70, 335
- argument from evil [8](#), [10–12](#), [16](#), [19](#), 87–8, 96, 154, 159, 188, 193–4, 198, 201–2, 216, 218, 220–60, 261–3, 269, 270–2, 281–90, 296, 302, 306, 325, 327, 329, 338–9, 341, 357, 364, 366, 367, 369
- animal suffering 225–33
- and AP 88, 220–5, 337–40, 364, 367
 - and atheism 220, 337
 - and BT 221–5
 - and free will 233–60
 - and religious diversity 214, 262–3
- argument from scale 31, 91, 97, 128, 193–219, 221, 261, 268, 328–31, 334, 340, 353, 357, 378
- and AP 195–7, 200, 213–19, 328–31
 - and atheism 193, 195–200, 205–6, 210, 212–15, 330
 - and BT 194–6, 198