


OXFORD

a luxury of the understanding

on the value of true belief

ALLAN HAZLETT



EVERYTHING IS GOING TO BE ALRIGHT

A Luxury of the Understanding

This page intentionally left blank

A Luxury of the Understanding

On the Value of True Belief

Allan Hazlett

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

© Allan Hazlett 2013

The moral rights of the author have been asserted

First Edition published in 2013

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: 2013938358

ISBN 978-0-19-967480-0

Printed and bound in Great Britain 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.

I wasn't sure whether to go back to that, whether, if I insisted, he'd end up telling me what had happened or what he'd found out, I knew he would tell me something, however partial or erroneous, but it's easy to want to know nothing when you still don't know, once you do, you've no choice, he was right, it's better to know about things, but only once you do know them (and I still didn't).

—Javier Marias, *A Heart So White*

This page intentionally left blank

Acknowledgements

Even the sympathetic reader will note that this book leaves myriad questions unanswered, even unconsidered; and most of you will have many other complaints. Revisions, amendments, corrections, additions, codas, qualifications, and clarifications could have gone on indefinitely. There are some vexing omissions: a discussion of contemporary accounts of intellectual virtue would have been appealing; an historical genealogy of the idea of the value of truth would have been illuminating; an inquiry into the nature of belief would have been another book. In any event, before acknowledging my numerous debts to others, I should like to acknowledge these failures (among others) of my own.

I developed the approach to philosophy applied in this book, and the broader ethical picture that it supports and with which it coheres, with Simon Feldman, over the last decade. For being a great interlocutor and friend, I owe him thanks.

The Epistemology Reading Group at the University of Edinburgh read a manuscript of this book in the Fall of 2012. The feedback I received from members of that Group has saved me from many errors; for their help I am most grateful.

I should like to thank, for conversations and valuable feedback: Mark James Adams, Cathy Bach, Anne Baril, Tim Bates, Michael Brady, Evan Butts, Dave Chambers, Matthew Chrisman, Bryan Frances, John Greco, Stephen Grimm, Brian Hazlett, Daniel Howard-Snyder, Xingming Hu, Sebastian Köhler, Michael Lynch, Robin McKenna, Duncan Pritchard, Mike Ridge, Bill Rowley, Nishi Shah, Ted Sider, Ernie Sosa, Asbjørn Steglich-Petersen, Valerie Tiberius, Nick Treanor, Lani Watson, Ralph Wedgwood, Lee Whittington, and Jared Woodard, as well as Peter Momtchiloff and my readers at Oxford University Press. I'll also include here the Graduate Center Bar, in Providence, and the Brass Monkey, in Edinburgh.

In May 2008 research for this book was supported by a Faculty Research Grant from Fordham University (who provided cash) and by the Arché Centre at the University of St. Andrews (who provided an office), and in April 2012 by a Small Research Grant from the British Academy.

I presented material for this book in 2008 at the APA Central Division meeting in Chicago, at the Arché Centre in St. Andrews, at Durham University, and at the University of Arizona; in 2009 at a conference on *Epistemic Goodness* at the University of Oklahoma, and at a workshop on *The Normativity of Belief and Epistemic Agency* at the Instituto de Investigaciones Filosóficas at the Universidad Nacional Autónoma de México; in 2010 at the University of Edinburgh, at the University of Glasgow, at Rochester University, and at East Carolina University; in 2011 at the winter meeting

VIII ACKNOWLEDGEMENTS

of the Scots Philosophical Association, in Edinburgh; and in 2012 at Edinburgh's Departmental Symposium. My sincere thanks to all those audiences for their questions and objections, and for their hospitality.

AH

Edinburgh, December 2012

Prologue

Before making any move ... he threw back his head and gave a long trombone-blast of anarchistic laughter. It was all so wonderful, even if it did go wrong, and it wouldn't.

—Kingsley Amis, *Lucky Jim*

Imagine that you have just fallen in love. Your love is extraordinary. It's come unexpectedly and without warning, but it's the most brilliantly dizzying feeling you've ever experienced. You've been transformed, but more than that: your world has been transformed, for now everything seems different. Things, places, events, ideas—their existence, their very nature, is colored by their relationship to you and your beloved. You look at each other, embarrassed, and the meaning of the song playing on the radio changes forever. You'll never visit Logan Airport again without thinking of her, the way her eyes lit up when she came into the arrivals hall. The significance of entire cities shifts: he was born in Lexington, went to college in Chicago, and fell in love with you in Baton Rouge. But more than this: you've not only been transformed, you've been transubstantiated: you aren't just yourself in a new relationship, you're a new self, essentially defined by that relationship. You've become a relational being. Unused to this existence, you imagine that it's not really happening. Your feet feel light and the sidewalk feels elastic. Can other people tell how strange you've become? Your behavior must give you away: you drive from one coast to another for a surprise rendezvous that is both spontaneous and inevitable; you cringe at a cynical joke that you would have loved—before your strange metamorphosis. You don't know who you are; your beliefs and values seem up for grabs: you're a vegetarian but her roasted beef ribs somehow seem appealing; you've never liked pets but you find yourself wanting to play with his dog. You're in love, and you realize that the supreme good in the entire universe is your beloved's smile. You keep saying to yourself: don't let this end, don't let this end, don't let this miraculous, incredible, perfect state of affairs ever come to an end.

But your love is ordinary. Like most loves, it will probably come to an end. Your beloved will stop loving you, and you will stop loving your beloved. And your relationship with your beloved will probably be severed along the way. It might be an amicable and mutually agreeable breakup; it might be a messy divorce. Someone will cheat on someone—you hope it doesn't come to that, but it probably will. You do some research on other people's relationships, hoping that it's some alien demographic that's driving the divorce rate up. Imagine that it isn't: you learn that most romantic relationships, of the sort that you are in, are impermanent. This seems like good evidence that your

relationship is likely to end; suppose that's right. And now imagine further that your relationship *is* likely to end—in whatever sense future events are likely or unlikely.

What should you believe, when it comes to the proposition that your relationship is likely to end? We stipulated that that proposition is true, and that you have good evidence that it is true. You hope that your relationship won't end; you want nothing more than for it not to end; you *need* it not to end. But the question remains: what should you believe?

This book is motivated by that question. It won't answer it, and indeed it won't address it directly. A more basic question needs to be addressed first, a question about the value of true belief, and that question will be the topic of most of what follows. But I hope most of all to convince you that the answer to the present question—about what you should believe, when you're in love—is not at all obvious. It's profoundly un-obvious what you should believe, in the situation described above.

You might think it's obvious that you *should not* believe that your relationship is likely to end. One reason you might think this would be that you think something like this: in the situation described, you have evidence that you and your beloved are an exception to the statistical regularities that you researched. In effect, on this proposal, you don't have reason to think that most romantic relationships, of the sort that you are in, are impermanent, because there is something exceptional about the relationship you are in—it's not actually of the relevant "sort." There is something right about this idea. When you are in love, it sometimes seems like your relationship will never end—it can be hard, or impossible, to imagine such a thing. Inconceivability is a kind of evidence for impossibility, so perhaps the lover's inability to imagine her relationship ending is a kind of evidence that it won't. But this seems like a bad argument. For presumably all relationships feel this way at first, including those that end, and which comprise the majority of relationships. Better: presumably there is no reliable connection between the feeling of romantic necessity, as we might call it, and actual relationship success. Perhaps this is wrong, but it seems plausible, and you've been given no reason to think that it's true, in the situation described above. If this is right, then the present proposal, in defense of the view that you obviously should not believe that your relationship is likely to end, won't work.

A different kind of reason that you might think it's obvious that you should not believe that your relationship is likely to end is that you think such a belief would be really unpleasant, in the situation described, and that you should therefore avoid forming such a belief, either by suspending judgment or by believing that your relationship is not likely to end. There is something right about this idea, too. But as articulated, this proposal won't work either. The reason is that it seems to rely on the principle that you should never do something that will be really unpleasant. But this is implausible, since there are lots of situations in which the best thing to do is really unpleasant. Health-preserving medical treatments can be really unpleasant, for example, but we should not conclude, from that fact alone, that you should not undergo health-preserving medical treatments. Romantic relationships, to take

another example, can be really unpleasant, but we should not conclude, from that fact alone, that you should avoid romantic relationships. So the present proposal, as articulated, won't work.

You might reply that the relevant belief—that your relationship is likely to end—isn't merely really unpleasant, but is also bad vis-à-vis its potential consequences on the relationship, and so (plausibly) on your general wellbeing. This idea is important, and we'll be looking closely at the relationship between true belief and wellbeing in what follows. But, as we'll see, that relationship is fantastically complicated, and far too complicated for it to be *obvious* that you should not believe that your relationship is likely to end, in the situation described above.

You might think it's obvious that you *should* believe that your relationship is likely to end. One reason you might think this would be that you think that not believing this truth is likely to have detrimental consequences on your relationship, and so (plausibly) on your general wellbeing. Again, we'll be looking closely at the relationship between true belief and wellbeing below, but the relationship is too complicated for anything about it to be obvious.

A different kind of reason that you might think it's obvious that you should believe that your relationship is likely to end is that you think that true belief is valuable for its own sake. The truth—that your relationship is likely to end—is staring you in the face; you've got very good evidence in favor of this proposition. You should therefore believe that proposition, so the argument goes, and neither suspend judgment nor believe its negation. There is something right about this idea, as well, but what exactly is right about it is difficult to articulate. This is why we'll soon turn our attention to the question of the value of true belief.

In what follows I'll argue for two theses. The first is that true belief is at most *sometimes* valuable vis-à-vis the wellbeing of the believer. The second is that, although true belief is always "epistemically" valuable, this is a trivial consequence of the definition of "epistemic," and we should be anti-realists about "epistemic normativity." The best way to capture the stance towards the value of true belief that is adopted in this book is to say that it is a *skeptical* stance. It is most certainly skeptical of the idea that the value of true belief is *obvious*.

In my view, true belief is valuable only to the extent that, and in the sense that, some people value true belief. Our love of truth—the contingent fact that some people love true belief—is all there is to the value of true belief. I won't be able to provide a comprehensive defense of this view here, but I will try to convince you that some common and initially plausible positions, which favor the value of true belief, should be rejected. This should at least make plain the appeal of the view just mentioned.

This page intentionally left blank

Contents

1. Two Ancient Ideas	I
 Part I. The Eudaimonic Value of True Belief	
2. Greatness of Mind	39
3. Partiality and Charity	87
4. True Belief as a Non-Ideal Good	109
 Part II. The Epistemic Value of True Belief	
5. The Problem of the Source of Epistemic Normativity	131
6. Humean Approaches	155
7. Darwinian Approaches	180
8. Kantian Approaches	206
9. Anti-Realism about Epistemic Normativity	242
Epilogue	273
<i>Bibliography</i>	279
<i>Index</i>	297

This page intentionally left blank

1

Two Ancient Ideas

Methinks I am like a man, who having struck on many shoals, and having narrowly escap'd ship-wreck in passing a small firth, has yet the temerity to put out to sea in the same leaky weather-beaten vessel, and even carries his ambition so far as to think of compassing the globe under these disadvantageous circumstances.

—Hume, *Treatise of Human Nature*

This book is a critical study of the value of true belief, an examination of the idea that true belief has value.¹ This chapter will introduce two ideas that will be with us throughout. The first idea, which is evaluative, is that true belief is better than false belief. The second idea, which is metaphysical, is that there is something in the nature or essence of human beings, or in the nature or essence of our beliefs, that makes us, or our beliefs, directed at truth.

Although a few important conclusions will be drawn in this chapter, most big issues will be introduced here and examined in detail in later chapters, which are divided into two parts: Part I concerns what I will call the “eudaimonic” value of true belief (Chapters 2–4); Part II concerns the “epistemic” value of true belief (Chapters 5–9). Our task in this chapter will be to get a sense of the issues that are at stake in those later chapters, and to draw some distinctions that will prove essential later on.

The philosophical questions at stake here are complex and not well understood. It would be insensible for anyone to claim to know their answers. My aim is to establish the virtues and vices of various views, which is not to say that my approach isn't partisan: I argue for a particular answer to the question of the value of true belief. On my view, although it's plausible that true belief is sometimes valuable, it's not the case that true belief is always valuable (Chapter 4), nor is it the case that true belief is normally valuable (Chapters 2–3). And I shall defend anti-realism about the “epistemic” value of true belief (Chapter 9). Overall, these views are best described as a form of skepticism about the value of true belief. But all my arguments are based on assumptions; your

¹ On the title of the book, see the Epilogue. “The understanding” refers to the intellectual faculty (as in: “a heart unfortified, a mind impatient, an understanding simple and unschool'd”), not to the species of knowledge (as in: “give it an understanding, but no tongue”).

assumptions might be different and you might come to different conclusions. When we have finished, however, we will be able to see various assumptions in play, and we will be able to see the motivations for, and commitments of, the various views from which we can choose.

If we can say anything about philosophers in general, it's that they love truth. Someone who purported to be a philosopher, but lacked a passion for truth, would be a fraud. If a philosopher defends either of the two ideas mentioned above, therefore, we should be suspicious, just as we are always rightly suspicious when someone from some particular group insists that anyone outside that group is not "really" happy, or not doing what "nature" intended, or whatever. There is a temptation for human beings to want to say that their way of life, at least if they are happy with it, is the *only* way of life that could make anyone happy. David Hume describes this tendency in his essay of 1741 on "The Sceptic":

Almost every one has a predominant inclination, to which his other desires and affections submit, and which governs him, though, perhaps, with some intervals, through the whole course of his life. It is difficult for him to apprehend, that any thing, which appears totally indifferent to him, can ever give enjoyment to any person, or can possess charms, which altogether escape his observation. His own pursuits are always, in his account, the most engaging: The objects of his passion, the most valuable: And the road, which he pursues, the only one that leads to happiness. (Hume 1985, p. 160)

But, as Hume observes, such thinking, though common even among philosophers, is absurd. It forgets the "vast variety of inclinations and pursuits among our species" (ibid.). That the two ideas we'll discuss exhibit this common pattern may explain their popularity, but it only diminishes their credibility. Hume's skeptical interest in individual differences was not new—such interest can be traced back to the Ancient skeptics, who marveled at the diversity of human opinions, values and preferences, practices, laws, and customs, and whose work inspired a renaissance of skeptical thinking in the Early Modern period²—and my inquiry in this book, which is skeptical of the value of true belief, similarly emphasizes difference as against universality.

1.1 The eudaimonic value of knowledge

We shall begin by asking after the value of knowledge, which we'll first articulate as "eudaimonic"—i.e. value vis-à-vis wellbeing.

1.1.1 *Socrates on the unexamined life*

Among the most well-known slogans in the history of philosophy is "The unexamined life is not worth living," articulated by Socrates in Plato's *Apology* (38a).³ Socrates

² See, for example, Sextus Empiricus, *Outlines of Pyrrhonism*, Book I, 79–90 and 145–61, or Book III, section xxiii, and Michel de Montaigne, "An Apology for Raymond Sebond."

³ For English quotations from Plato I am using G.M.A. Grube's translations of the *Apology*, *Meno*, and *Republic*, in Plato 1997.

has been asked to defend his philosophical practice against the charge that it's impious and socially corrupting. He responds by maintaining that philosophy is "the greatest good for a man," and that "the unexamined life is not worth living for men." Call this the **Socratic maxim**.

Life is said to be not worth living if one fails to *exetazein* life; the word means to examine well or to inquire into. Inquiry is intentional activity with a certain goal: knowledge about the thing examined. Someone who inquires about x seeks knowledge about x . Why should we think, then, that intentional activity aimed at knowledge of life (setting aside for the moment what knowledge "of life" amounts to) is so valuable that life is not worth living without it?

The Socratic maxim suggests that knowledge (of life) is extremely valuable. This suggestion is based on the idea that inquiry about x is valuable only if knowledge about x is valuable. And this idea, in turn, seems well supported on the assumption that aiming at x is valuable only if x is valuable. But this is not always the case. Suppose my physician prescribes the following regimen: I shall stretch the muscles in my lower back by trying to touch the floor while standing up. Aiming at touching the floor is good, since it leads me to stretch, but touching the floor (we can easily imagine) is not good. We can even imagine a variant on this case in which aiming at x is valuable while x is disvaluable: if my physician gets me to stretch by placing an appealing cupcake on the floor, which cupcake would actually be quite unhealthy for me were I to get my hands on it. Aiming at acquiring the cupcake is good, but acquiring the cupcake is bad. These cases show not only that aiming at x can be valuable while x is not valuable, but also that the value of aiming at x need not be explained in terms of the value of x . The goodness of touching the floor is not what explains the goodness of trying to touch the floor; it's the goodness of stretching that explains the goodness of trying to touch the floor.

It's possible that the relationship between inquiry and knowledge, when it comes to their value, is akin to the relationship between trying to touch the floor and touching the floor, when it comes to their value. But note well that this conclusion is not implied by the skeptical thought that inquiry will never terminate with the acquisition of knowledge. That is consistent with knowledge being valuable, and with the value of knowledge explaining the value of inquiry. There are at least two (compatible) possibilities here. The first is that inquiry is good because it gets us *closer* to knowledge, which is good. So even though knowledge cannot be acquired, the value of knowledge explains why seeking knowledge is valuable. The second is that inquiry is good because it *represents* knowledge, which is good, as being good. Even though knowledge cannot be acquired, by seeking knowledge the inquirer treats knowledge as something valuable, and since knowledge is valuable, inquiry is therefore valuable in the sense that it is an apt response to the value of knowledge.

And for this reason, saying that the value of knowledge explains the value of inquiry does not imply that inquiry is merely of instrumental value. Something has **instrumental value** iff it is valuable for the sake of something else, i.e. if it is valuable as a

means to some (wholly distinct) end;⁴ something has **final value** iff it is valuable for its own sake. The value of knowledge could explain the value of inquiry, even if inquiry never leads to knowledge (because knowledge is impossible), and therefore even if inquiry has no instrumental value, because it is not a means to the end of knowledge.

Here I'll pursue the idea that the Socratic maxim suggests that knowledge (of life) is extremely valuable, noting only that this does not imply that knowledge (of life) is possible nor that inquiry is merely instrumentally valuable. But you may opt for an alternative interpretation, on which the relationship between inquiry and knowledge, when it comes to their value, is akin to the relationship between trying to touch the floor and touching the floor, when it comes to their value.

So the Socratic maxim suggests that knowledge (of life) is extremely valuable. On the most obvious reading of this, the value in question here is value vis-à-vis wellbeing. It's the value of knowledge vis-à-vis wellbeing that explains why the unexamined life is not worth living (or "not to be lived," an equally good translation). This sort of value—which we'll call **eudaimonic value**⁵—will be our concern in Part I of this book. Eudaimonic value, for S, is value vis-à-vis the wellbeing of S. Eudaimonic value, for a person, concerns what is good and bad for that person, i.e. her wellbeing. "**Wellbeing**" is used here in a broad sense, such that it is a name for welfare, the good life (on one disambiguation, see below), quality of life, happiness—in the sense of a happy life, rather than the feeling of being happy (Kraut 1979, p. 179), and not necessarily in the sense of a life of "contentment, enjoyment, or pleasure" (Foot 2001, p. 85)—or for "living well and faring well" (Aristotle, *Nicomachean Ethics*, 1095a15).⁶ Just as "everything" is the uninformative answer to the question of what there is, "wellbeing" is the uninformative answer to the question "What makes someone's life go best?" (Parfit 1984, pp. 493–502), or, more exactly, "What makes a life a good one for the person who lives it?" (Scanlon 1993, p. 185) Given this broad sense of "wellbeing," there are various theories of wellbeing, including hedonist theories, desire-fulfillment theories, life-satisfaction theories, and essentialist theories—more on which below (§1.1.6). "Eudaimonic value," therefore, does not necessarily refer to value vis-à-vis Aristotelian *eudaimonia*, nor to value vis-à-vis *eudaimonia* according to any particular theory, nor to value vis-a-vis wellbeing, on an essentialist theory of wellbeing. "Eudaimonic value"

⁴ Note the requirement that the end be wholly distinct from the means. What I will call "constitutive value" (§1.1.5) is a species of final value, but some would call it instrumental value where the means is "internal" to the end.

⁵ From the Greek word *eudaimonia*: "prosperity, happiness" (Liddell and Scott, *Greek-English Lexicon*, abridged). This terminology reflects consideration. "Wellbeing" doesn't work grammatically: there's no adjectival and adverbial forms. Eudaimonic value is the same as what some would call "prudential value," but "prudential" suggests expediency and self-interest, which we should not build in to our conception of wellbeing. "Ethical value" captures the idea that we are concerned here with the question of how to live a good life, but is too suggestive of morality, or of the view that we ought to pursue or promote (our own or others') wellbeing (on which we'll remain neutral). "Pragmatic value" suggests philosophical pragmatism, and, again, expediency and self-interest.

⁶ For English quotations from Aristotle I am using W.D. Ross' translations of the *Nicomachean Ethics* and the *Metaphysics* in Aristotle 1984.

refers, by definition, to value vis-a-vis wellbeing, whatever wellbeing is. In other words, “eudaimonic” is to be understood formally, not materially.⁷

Wellbeing is the uninformative, i.e. trivial, answer to the question of what makes someone’s life go best for her. It is *not* a trivial answer to the **ethical** question of what she ought to pursue or promote (cf. Sumner 2002, pp. 33–4, Tiberius 2008, pp. 9–15), although one might defend that non-trivial ethical view.

Furthermore, eudaimonic value is conceptually distinct from **moral** value, moral virtue, and from other moral notions (cf. Sumner 1996, pp. 24–5). Eudaimonic value concerns the wellbeing of a person, and the concept leaves open the possibility that something might be morally good for a person but eudaimonically bad for her. This possibility is plausibly realized in cases of morally virtuous self-sacrifice. However, the concept also leaves open the possibility of causal and constitutive connections between morality and wellbeing. This possibility is plausibly realized in as much as moral virtue is a constituent of wellbeing. The view that morality is among the constituents of wellbeing represents an appealing *via media* between, on the one hand, the view that being moral is irrelevant to a person’s wellbeing, and, on the other hand, the view that being moral never involves a sacrifice of wellbeing. In what follows I will assume that falling short of moral virtue is sometimes better for a person than not (§2.9.1), but that moral virtue is partially constitutive of wellbeing (§3.3.2). We’ll look at the social value of true belief, below (§4.4).

I assume no particular theory of wellbeing. I will, however, make some controversial assumptions about wellbeing. When I say that I assume no particular theory of wellbeing, I don’t mean that what I’ll say will be compatible with any theory of wellbeing. You might object that little can be said about wellbeing without first settling on a theory of wellbeing. I disagree. It’s wrong, in general, that particular claims about *x* can’t be evaluated until a philosophical account of *x* has been given. There are two reasons for this. First, so long as a concept is one that we already possess, prior to philosophical theorizing, then however vague, ambiguous, and in need of clarification that concept is, there will still often be claims, employing that concept, that are obviously true. You don’t need to settle on a theory of wellbeing to know that your life would go worse if, while everything else remains the same, you were to be tortured every day for the rest of your life. Second, and more importantly, we must consider how the imagined methodology would go. How shall we evaluate, for example, various theories of wellbeing? The only possible way to do this is to consider their particular implications about wellbeing, and to consider the plausibility of those particular implications. But

⁷ This usage is appealing, in part, because it easily allows us to give a natural interpretation of the disagreement between (for example) defenders of desire-fulfillment theories of wellbeing and Aristotelian essentialists: the two camps are disagreeing about the nature of wellbeing (cf. Kraut 1979, and also MacIntyre 2007, p. 181–2, on “virtue”). This isn’t to say that there aren’t other ways to articulate the disagreement (e.g. we might treat both camps as agreeing that wellbeing is desire-fulfillment, but disagreeing about whether wellbeing is to be pursued). But understanding this as a disagreement about wellbeing is the most natural interpretation of the debate.

this means that, in at least one sense, we can evaluate particular claims about wellbeing without having antecedently settled on a theory of wellbeing. This is not to say that we might not learn something about particular cases by seeing what our preferred theory says about them. But this means that it's not misguided to consider particular claims about wellbeing without first adopting a specific theory of wellbeing. (Compare other areas of philosophy.)

Our discussion of the eudaimonic value of knowledge, and later true belief, will require a modal notion of eudaimonic betterness:

Definition: The proposition that *p* is better for *S* than the proposition that *q* iff either *p* and were it the case that *q* then *S* would be worse off, or *q* and were it the case that *p* then *S* would be better off.

To say that the proposition that *p* is better for *S* than the proposition that *q* is to say that the fact that *p* makes (or would make) *S*'s life go better than it would go were it the case that *q* (or better than it actually goes given the fact that *q*). To say that the proposition that *p* is better for *S* than the proposition that *q* is a way of saying that the proposition that *p* is good for *S*, that the proposition that *p* contributes to *S*'s wellbeing, at least compared to the proposition that *q*, and that the proposition that *q* is bad for *S*, that the proposition that *q* detracts from *S*'s wellbeing, at least compared to the proposition that *p*. This is what I shall mean below when I speak of one thing being "better" than another; I'll not mention the subject *S* when she is implied by the context.

Returning to the Socratic maxim: why should we think that knowledge of life has eudaimonic value? Is it the case, for example, that knowledge of life is better than ignorance of life? Here we must make a decision: does the eudaimonic value of knowledge of life derive from the eudaimonic value of knowledge in general, or from something about knowledge of life in particular?

It seems clear, at least when it comes to the idea that the unexamined life is not worth living, that Socrates holds the latter view. Examination of life seeks a particular species of knowledge, and it's the pursuit of this particular species of knowledge that Socrates maintains is necessary for life being worth living. For it's the practice of *philosophical* examination that is at issue in the *Apology*. What Socrates defiantly refuses to give up is inquiry into philosophical questions in ethics and metaphysics—the subject matter of his inquiries as depicted in Plato's dialogues. Knowledge about these matters, therefore, is eudaimonically valuable, in such a way that examination of these matters is necessary for one's life being worth living. Consider the maxim in context:

[I]f I say that it is the greatest good for a man to discuss virtue every day and those other things about which you hear me conversing and testing myself and others, for the unexamined life is not worth living, you will believe me even less. (op. cit.)

The Socratic maxim, on this interpretation, says that knowledge of answers to philosophical questions in ethics and metaphysics, which we might call "wisdom" (although this is out of step with contemporary usage), is eudaimonically valuable such that life is not worth living unless one seeks it. Is this plausible?

I do not think it is, but because my primary concern is with the value of knowledge in general, I will be brief here. The idea that the lives of non-philosophers are not worth living (whether this is down to their lack of eudaimonic value, or whether it's down to something else) seems both elitist, in condemning the great majority of human lives to "not worth living" status, and suspiciously self-serving, in as much as this idea about the worthlessness of non-philosophical lives is put forward exclusively by philosophers. Even granting that philosophical inquiry is eudaimonically valuable, the Socratic maxim seems to ignore all the other eudaimonically valuable things that might make a life go well enough to be worth living: scientific inquiry, the appreciation or pursuit of athletic excellence, friendship, living with a family, the appreciation or creation of music, literature, cuisine, and the other forms of art, romantic relationships, the pleasures of drink and sex, and so on. Even this list seems rarefied; where are the suburban values of owning a home and having a career? If Aaron and Maria excel at their well-paying and interesting jobs, maintain a pleasant house where they raise their happy and precocious children, enjoy home-cooked food and going out to minor-league baseball games with the kids, but never give a moment's thought to the big philosophical questions, could anyone in good faith say that their lives are not worth living?

It will not help to say that Socrates is employing an ancient conception of wellbeing, one that is foreign to contemporary common sense, unless this is just to offer an historical explanation of why he got things wrong. Nor will it help to say that Socrates is appealing to a notion of perfect human excellence that is a rare and elite thing, acquired only by a select few. Consider the context of the *Apology*: Socrates has been threatened with death, and in the passage in question he is explaining why he would not cease to practice philosophy even if ordered to do so by the law, on pain of death. Socrates effectively says to the jury, "Give me philosophy, or give me death!" and the jury chooses to give him the latter. In the context of that story, the Socratic maxim is most naturally interpreted as meaning that death is superior to a life without philosophical examination, i.e. that the unexamined life is not worth living, not (merely) that the unexamined life falls short of perfect human excellence.⁸

Socrates defends his maxim on the grounds that philosophy is "the greatest good for men." If not for this fact, we might interpret his claim in a relative way, on which Socrates is saying that, given *his own* commitment to philosophy, he would rather die than give up the pursuit of wisdom. In other words, we might interpret him as saying that the unexamined life is not worth living *for Socrates*. On this reading, Socrates would be heroically affirming his love of wisdom, without making any claim about the value of philosophy for human beings in general. It can certainly be the case that someone cares about something so much that she would rather die than give up its pursuit, and we should have no objection to caring about wisdom in that way, nor to Plato's depicting his teacher as caring about wisdom in that way.

⁸ Cf. the distinction between flourishing and excellence (§1.1.6).

1.1.2 *The eudaimonic ideal of knowledge, implausibly formulated*

Let's turn our attention to the idea that knowledge in general is eudaimonically valuable. When we speak of the eudaimonic value of knowledge (or true belief), we'll mean the value of knowledge vis-à-vis the wellbeing of the knower (or believer). Let's say that **S knows whether p** iff either p and S knows that p or $\neg p$ and S knows that $\neg p$. We'll take **ignorance about whether p** to be the logical negation of knowing whether p. Consider the following thesis, which affirms the eudaimonic value of knowledge in general:

Implausibly formulated eudaimonic ideal of knowledge: For any subject S and proposition that p, knowing whether p is better for S than being ignorant about whether p.

In other words, for any subject S and proposition that p, either S knows whether p and were she ignorant about whether p she'd be worse off, or S is ignorant about whether p and were she to know whether p she'd be better off.⁹ As the name suggests, this claim is not plausible:

Karen is scheduled to compete in a tennis match against an opponent that, unbeknownst to her, is very much her superior on the court. Given an anxious disposition, she will perform terribly and almost certainly lose if she is made aware of her opponent's talent. Karen's coach has seen the opposing player in practice, but knows that Karen stands a chance only if she believes that she and her opponent are evenly matched. So he tells Karen that she and her opponent are evenly matched, and she believes this. She begins the match confident and optimistic.

Karen would *not* be better off were she knowledgeable about whether her opponent is vastly superior to her, so knowing is *not* better for her than being ignorant.¹⁰ So the implausibly formulated eudaimonic ideal of knowledge is false.

I'll have much more to say in defense of the existence of species of eudaimonically valuable false belief. We'll consider evidence from social psychology that supports the idea that false beliefs, resulting from "self-enhancement bias," are often eudaimonically valuable (Chapter 2), and I'll argue that false beliefs about other people are often eudaimonically valuable (Chapter 3). And we'll consider some other candidates for "functional false belief," below (§7.4.1).

There is an important objection to our counterexample that can be addressed here. You might argue that Karen is irrational or vicious in some way, and that this is what puts her in the unfortunate position where ignorance is better for her than knowledge. What would be best for Karen, so the argument goes, is not to falsely believe that she

⁹ As formulated the eudaimonic ideal of knowledge leaves open the question of whether ignorance in the form of false belief is worse than ignorance in the form of lack of belief, or ignorance in the form of suspended judgment. Few would want to say that a lack of true belief about uninteresting or insignificant truths is disvaluable, vis-à-vis wellbeing (cf. §§4.2–4.3). My focus in this book will be on ignorance in the form of false belief.

¹⁰ Note that this doesn't imply that Karen's coach is morally justified in lying to Karen.

is evenly matched with her opponent, but to abandon her nervous disposition, and develop the strength of character needed to perform well despite knowing that she is the inferior player. But given our modal definition of betterness (§1.1.1), this objection is a non sequitur. The implausibly formulated eudaimonic ideal of knowledge entails that, when you are ignorant, you would be better off were you knowledgeable. It does not merely entail that, when you are ignorant *and* you are perfectly rational or perfectly virtuous, then you would be better off were you knowledgeable.

The eudaimonic ideal might be reformulated, with a restriction placed on the domain of quantification, such that it applies only to the perfectly rational or perfectly virtuous. This might yield what Gavin Lawrence (1993) calls an “ideally circumstanced...ideal” or a **utopian ideal** (p. 8). We won’t consider that idea further. One reason why such a reformulation will be set aside is that we seek to articulate principles of wellbeing that might provide a certain kind of **guidance** for living well that applies to us, as we actually and presently are. If I know that the proposition that *p* would be better for me than the alternatives, then I’ve got a nice bit of advice: try to make it the case that *p*. But if all I know is that the proposition that *p* would be better for my perfectly rational or perfectly virtuous counterpart, then I’ve got nothing to go on, given the fact that I know I’m neither perfectly rational (nor even especially rational) nor perfectly virtuous (nor even especially virtuous).¹¹

This is not to say that a utopian ideal cannot provide guidance of a different kind. A description of my perfectly rational or perfectly virtuous counterpart would provide me with something to aspire to and work towards. Guidance would come in the form of the injunction to try to make myself like my perfectly rational or perfectly virtuous counterpart as possible. But I will require a different sort of guidance so long as my efforts to become perfectly rational or perfectly virtuous are not immediately successful. Imagine that I suffer from intemperance, and that I am on my way home for Thanksgiving dinner, when I realize that I ought to stop at the pub to wish some old friends a happy holiday. However, I have excellent reason to think that entering the pub will precipitate a bender that will land me in the drunk tank and ruin Thanksgiving for my family. I should obviously not go to the pub. But my perfectly virtuous counterpart would be better off going to the pub, to say hello and have a quick pint, before temperately heading home. I, however, am too intemperate to have any use for the advice “do what your perfectly virtuous counterpart would do.”

Ideals should provide guidance. But does that mean that the eudaimonic ideal of knowledge commits us to the view that we can *choose* whether or not we are knowledgeable or ignorant about some question? That suggests a kind of voluntary control over our beliefs that many have thought impossible (cf. §8.4.1). But we are not committed to this by insisting that the eudaimonic ideal of knowledge provide guidance. The eudaimonic ideal of knowledge tells us what is, or would, be good for us.

¹¹ As Valerie Tiberius (2008) puts a related point, “[g]iven that we are not, nor ever will be, ideally or perfectly rational, it is not obviously helpful to be told that we should choose whatever we would choose if we were” (p. 7).

Guidance will be implied by the ideal in the form of the injunction to try to make it the case that *p*, when the proposition that *p* is better for me than the alternatives. How might you be guided by the eudaimonic ideal of knowledge? The implied injunction is to make it the case that you are knowledgeable. How can you be guided by this? By inquiring carefully, by reflecting critically, by engaging in the practices of the intellectually virtuous with an aim to habitually acquiring those virtues yourself, by being vigilant when it comes to your own prejudices and biases, by keeping company with the intellectually virtuous and avoiding the unreliable and the incurious, and so on.

There are other objections you might have to the putative counterexample. These will be addressed below (see, in particular, §§2.9.3–2.9.4 and §§2.10.6–2.10.7). For now, we'll reformulate the eudaimonic ideal of knowledge so as to avoid it.

1.1.3 *The eudaimonic ideal of knowledge*

If you are sympathetic with the eudaimonic value of knowledge, how might you respond to the putative counterexample (§1.1.2)? First, you might say that knowledge always has **prima facie** eudaimonic value (Lynch 2004, pp. 46–57, 2009a, pp. 225–8, cf. Kvanvig 2003, p. 93). “Prima facie” means something like “at first glance” or “on its face.” Something that is prima facie valuable may turn out to have no value at all, if that's what further examination reveals. Knowledge doesn't *always* have prima facie value—that's just what the counterexample shows. So perhaps knowledge has prima facie value, in the sense that our default assumption is that knowledge has eudaimonic value. This claim is quite weak. Knowledge might have prima facie eudaimonic value, in this sense, but rarely turn out to be valuable, since appearances might be misleading. Fool's gold, for example, always has prima facie economic value. The philosopher sympathetic to the eudaimonic value of knowledge should not be satisfied by this. But those who say “prima facie” often mean “pro tanto”; we'll return to that more promising idea below (§1.1.4).

Second, you might say that knowledge is better than ignorance “**other things being equal**” (Finnis 1980, p. 72, Lynch 2004, p. 47, p. 54, p. 144). This claim is also weak. The eudaimonic ideal might be true, other things being equal, and yet it might turn out as a matter of fact that things are rarely equal. Imagine that I defend the eudaimonic value of taking cocaine on the grounds that, other things being equal, it's always better to take cocaine. You can't object to this that cocaine is unhealthy, or that cocaine is addictive, or that it destroys you psychologically, since I only claim that cocaine is good *other things being equal*. What I claim is that if you have two people who are the *same* in terms of health, addiction, and so on, then the one who takes cocaine is better off than the one who doesn't. This is clearly an inadequate defense of the eudaimonic value of taking cocaine. Or, at best, it shows us very little about the eudaimonic value of taking cocaine. The philosopher sympathetic to the eudaimonic ideal of knowledge should not be satisfied by this.¹²

¹² The idea that *x* is good, other things being equal, might be offered by way of articulating the idea that *x* has *intrinsic* value; the idea being that the goodness of *x* does not depend on anything other than *x* itself. So

Third, you might say that:

- Knowledge is **normally** better than ignorance.
- Knowledge is **generally** better than ignorance.
- Knowledge is **typically** better than ignorance.
- **For the most part**, knowledge is better than ignorance.¹³
- Knowledge is better than ignorance, save in **exceptional** cases.

Thus Michael Lynch (2004) says that cases of bad true belief are “exceptions that prove the rule” (p. 46), namely, the “general rule that it is good to believe what is true and only what is true” (p. 48). Linda Zagzebski says that knowing the truth is valuable “most of the time (1996, p. 200)”, that “true belief is usually good for us” (2003b, p. 23), and that “sometimes caring about knowledge...can be over-ridden by other things that we care about, but my conjecture is that this does not happen very often among self-reflective persons” (2004, p. 372), and Richard Foley (1987) says that false belief is practically rational only in “funny” situations, but not in “relatively normal situations,” and that while there are “exceptions,” it is “highly improbable that this will happen frequently” (p. 224; see also Foley 1993, p. 27).

Claims about what is normal can be tricky to evaluate. Part of the reason is that our judgments about what is normal tend to manifest our normative commitments, even when there is agreement about statistics. Suppose that 1% of offshore oil drilling operations result in oil spillage. Oil company executives will insist that this is simply the “exception that proves the rule,” namely, the rule that offshore drilling is generally safe. Environmentalists will vehemently object to this characterization; the spills are “all too typical” of offshore drilling. What the fan calls exceptions are just what the critic calls counterexamples. The oil companies call oil spills “exceptions to the rule” because they support continued offshore drilling, and the environmentalists call them “all too typical” because they oppose continued offshore drilling. Our normative commitments, therefore, can influence what cases we take to be normal. Defenders of the eudaimonic value of knowledge often appeal to perception of our immediate environment in defense of their view, where ignorance is presumed to be fantastically dangerous. In my critique of the eudaimonic value of knowledge I’ll appeal to cases involving other species of ignorance: ignorance of self (Chapter 2) and ignorance of other people (Chapter 3). If perception is taken to be the normal case, then these counterexamples will appear abnormal, mere “exceptions to the rule.” But this is special pleading, on behalf of the cases that support the eudaimonic value of knowledge, at the expense of the cases that threaten it. There is a worry here, that when we describe cases in which ignorance is better than knowledge as abnormal we may be doing nothing more than emphasizing our commitment to the eudaimonic value of knowledge.

you might say that knowledge has intrinsic value. We’ll consider the idea that true belief has intrinsic value, below (§1.5).

¹³ Cf. Aristotle: “We must be content...to indicate the truth roughly and in outline, and in speaking about things which are only for the most part true” (1094b19–23).

Someone skeptical about the eudaimonic value of knowledge, who wanted to get to the bottom of things, should be worried about all this. And if there is genuine controversy about the eudaimonic value of knowledge, little progress will be made by trading intuitions about what is, or isn't, abnormal.

Despite these worries, we'll employ an articulation of the eudaimonic value of knowledge in these terms. I'll borrow some language from Rosalind Hursthouse (1999), who articulates the idea that "the virtues, for the most part, benefit the possessor" (p. 173). On her view, this is falsified neither by cases of virtuous people faring badly nor by cases of wicked people flourishing. The claim that the virtues benefit the possessor is the claim that the virtues are "one's only reliable bet as far as a flourishing life is concerned" (p. 174). Hursthouse's claim is that no "regimen," other than virtue, "will serve one better" (p. 174), which leaves open the possibility that the "regimen" is not an especially reliable bet, just the most reliable of the options. (Consider here Hursthouse's comparison of the virtue theorist's "regimen" with a doctor's prescription: a medical treatment might be the best option, and still likely, perhaps even almost certain, to fail.) This would be falsified only by a "clearly identifiable pattern" of wicked people flourishing (pp. 173–4) or by a "clearly identifiable pattern" of virtuous people faring badly. All cases in which wicked people flourish or in which virtuous people fare badly, however, must be seen as abnormal, on this view.

Given our comparative notion of betterness (§1.1.1), we'll adopt a comparative version of Hursthouse's idea. Consider:

Eudaimonic ideal of knowledge: For any subject *S* and proposition that *p*, knowing whether *p* is normally better for *S* than being ignorant about whether *p*. In other words, there is no clearly identifiable pattern of cases in which ignorance is better than knowledge. Thus, for any subject *S* and proposition that *p*, knowing whether *p* is a more reliable bet, when it comes to wellbeing, than being ignorant about whether *p*.

So the basic idea is that knowledge is normally better than ignorance, and what we mean by that is that there is no clearly identifiable pattern of cases in which ignorance is better than knowledge, and we take this to imply that knowledge is a more reliable bet, when it comes to wellbeing, than ignorance. (The formulation is neutral as to whether the claim is necessary or contingent.) This isn't falsified merely by cases in which ignorance is better than knowledge, as those cases can, and must, be seen as abnormal: as exceptions to the rule, as strange or one-off cases. Take, for example, Foley's (1987) "extreme example" in which believing some falsehood is necessary to save the world. Such cases are familiar in epistemology. The eudaimonic ideal of knowledge is not falsified by "extreme" cases in which ignorance is better than knowledge; it is only falsified by a clearly identifiable pattern of cases in which ignorance is better than knowledge (cf. McKay and Dennett 2009, p. 498).¹⁴

¹⁴ Compare Linda Zagzebski's (1996) approach to the value of the virtues (which she does not take to be eudaimonic value, p. 82). "[T]he sense in which virtue makes its possessor good," for Zagzebski, is that

Our worries, articulated above, remain. As Hursthouse (1999) points out, controversy about whether the virtues benefit their possessors is controversy over “beliefs... about...human nature and the way human life works,” and such beliefs are not straightforwardly beliefs about “empirical facts...accessible from ‘the neutral point of view,’” nor are they straightforwardly “evaluative beliefs” (p. 189). They are, she concludes, “ethical but non-evaluative beliefs about human nature and how human life goes” (ibid.). The idea that knowledge is normally better than ignorance is of exactly the same kind; and controversy about it is likewise problematic.

However, we can assume from the start that the eudaimonic ideal of knowledge is not verified if it turns out that in the statistical majority of cases, knowledge is better than ignorance (whatever exactly it would mean for there to be a statistical majority of “cases of knowledge”). On the statistical majority of roads, the more reliable bet, when it comes to avoiding a crash, is to drive on the right. But we should not say that, for all roads, driving on the right is the more reliable bet, when it comes to avoiding a crash, than driving on the left. The right thing to say is that *sometimes* driving on the right is the more reliable bet, and *sometimes* driving on the left is the more reliable bet. When you’re in the UK, for example, driving on the left is the more reliable bet. There are exceptions to this rule, as when another car has stopped in the left lane and you need to swerve into the right lane to avoid it. *That* is an abnormal circumstance. But the whole business of driving in the UK is not an abnormal circumstance. It is a clearly identifiable pattern of cases in which driving on the left is the more reliable bet, when it comes to avoiding a crash.

Our counterexample (§1.1.2) didn’t suggest that knowledge is never better than ignorance (cf. Lynch 2004, pp. 46–51). Below (Chapters 2 and 3), I’ll describe species of eudaimonically valuable false belief. I’ll argue that these constitute clearly identifiable patterns of cases in which false belief is better than true belief. By appealing to empirical psychology (Chapter 2), we’ll be able to move beyond the assessment of “ethical but non-evaluative beliefs,” which exist somewhere between the empirical and the non-empirical, in the direction of the empirical. And this is the best we can

“[a]nyone who has it is closer to reaching a high level of excellence than one who lacks it, other things being equal” (p. 95). Virtue in combination with vice can make someone less excellent; for all Zagzebski says, this may even be typical of virtue, when combined with vice. The sense in which it is good to possess some particular virtue, even when this makes you less excellent than you would otherwise be (because of some vice that you possess), is that possession of that virtue is an essential constituent of “a high level of excellence,” and in that sense possessing the virtue makes you “closer” to “a high level of excellence.” For this reason “virtue...invariably makes its possessor closer to a high level of admirability” (p. 101). My worry about this approach is that the supposed value of the virtues does not seem to be a value worth seeking. Suppose that a cheeseburger with ketchup would be the best meal for me to eat, but that I can only afford soup. It would not be good, in any way, to add ketchup to my soup, even though ketchup is an essential constituent of the best meal for me to eat, and in that sense adding ketchup will make my soup “closer” to the best meal. Note that Zagzebski also maintains something analogous to Hursthouse’s claim about the eudaimonic value of the virtues: being virtuous “*usually* results in an actual increase in a person’s overall moral worth” (p. 95, my emphasis) and “nothing is a virtue unless it benefits both the possessor and others in the *typical case*” (p. 100, my emphasis).

do, if we want to make progress. Our best bet for identifying patterns, when it comes to “how human life works,” is empirical psychology.

1.1.4 *The pro tanto eudaimonic value of knowledge*

There is an alternative, and more modest, way of articulating the eudaimonic value of knowledge. Rather than saying that knowledge is *normally* better than ignorance, we might say that knowledge *always* has pro tanto eudaimonic value (DePaul 2010, p. 114). Something valuable x has **pro tanto** value iff the value of x can be trumped by other values. On the present proposal, although knowledge always has eudaimonic value, its eudaimonic value can be trumped by the value of other things. This claim is relatively weak, in the sense that it might turn out that cases in which knowledge is better than ignorance are few and far between, since its eudaimonic value might be trumped more often than not. Below (Chapter 4) we’ll consider the idea that true belief always has pro tanto eudaimonic value.

We can now distinguish between two ways of articulating the idea that x has eudaimonic value. On the view that x is a **eudaimonic ideal**, x is normally better than alternatives. On the view that x is a **non-ideal eudaimonic good**, x is (sometimes or always) eudaimonically valuable, but x is not a eudaimonic ideal. One natural way of articulating the idea that knowledge is a non-ideal eudaimonic good is by saying that knowledge always has pro tanto eudaimonic value. An **ideal approach to the eudaimonic value of x** is one that maintains that x is a eudaimonic ideal; a **non-ideal approach to the eudaimonic value of x** is one that maintains that x is a non-ideal eudaimonic good.

1.1.5 *Instrumental vs. constitutive eudaimonic value*

Consider the idea that the only reliable way for someone to get what she wants is to make informed, in other words knowledgeable, decisions. You might defend the eudaimonic value of knowledge by appeal to this. This would be an **instrumental value approach** to the eudaimonic value of knowledge, on which knowledge is said to have **instrumental eudaimonic value**.

On such an approach, knowledge and wellbeing are conceived of as distinct (though causally connected) entities. One might argue, by contrast, that knowledge and wellbeing are not wholly distinct, as knowledge partially constitutes wellbeing. Some necessary connections are constitutive connections, when it is a necessary truth that all F s are G because being G (at least) partially constitutes being F , i.e. because (at least) part of what it is to be F is to be G . And some valuable things seem valuable in virtue of being parts of valuable wholes. Consider the aesthetic value of the parts of an aesthetically valuable painting. The aesthetic value of the painting depends, at least in part, on its beauty, and the painting’s parts do not cause the painting to be beautiful, but rather constitute its beauty. Because of this, they are (so the argument might go) valuable in virtue of being parts of a valuable whole, namely, the painting. Something similar is

often said of friendship: friendship does not merely cause a person's life to go well, but rather, or in addition, partially constitutes her good life. We have the following notion of a certain species of value:

Definition: Something x has **constitutive value** when x has final value in virtue of the fact that x (at least) partially constitutes some finally valuable whole.^{15,16}

Given this, we can articulate a notion of final eudaimonic value, namely, constitutive eudaimonic value:

Definition: Something x has **constitutive eudaimonic value** when x has final eudaimonic value in virtue of the fact that x (at least) partially constitutes wellbeing.

Moral virtue, friendship, and knowledge have all been said to have constitutive eudaimonic value. In the sequel, I'll use "constitutive value" to mean constitutive eudaimonic value. Something has constitutive value when it is eudaimonically valuable **in virtue of** being (at least) partially constitutive of wellbeing. Given this, I shall assume that something has constitutive value only if it **contributes** to the goodness of the whole of which it is a part, by partially **explaining** the goodness of the whole of which it is a part. There may be things that are, in some sense, parts of my life, and that life may be a good human life, but those things make no contribution to the goodness of my life.

Given this conception of constitutive value, you might propose a **constitutive value approach** to the eudaimonic value of knowledge, on which knowledge (sometimes or always) has constitutive value. Constitutive value is a species of final eudaimonic value, although it may not be the only species of final eudaimonic value; here it will be treated as the alternative to instrumental eudaimonic value. (On the intrinsic value of knowledge, see §1.5.) The constitutive value of knowledge is suggested by the familiar Aristotelian idea that human beings are essentially rational beings. Thus knowledge is said to be a "basic form of good" (Finnis 1980, p. 65), "being able to ... think, and reason" is described as a "central human functional capability" (Nussbaum 2000, p. 78), "intelligence" is said to be a "primary good" (Rawls 1971, p. 62), and "the pursuit of knowledge" is said to be a "human good," suited "for an important if not a central place in our life" (ibid. p. 425).

If knowledge has constitutive value, it's because knowledge is *partially* constitutive of wellbeing. Being knowledgeable is not all there is to living well. Consider someone

¹⁵ Although we define constitutive value as a species of final value, this assumption won't make a difference for what follows. You might object to this definition, on the grounds that constitutive value is conditional and extrinsic value; cf. our definitions of final and instrumental value (§ 1.1.1). Cf. Kvanvig's (2003) distinction between the "external" and "internal" value of knowledge.

¹⁶ The notion of constitutive value is not to be confused with the notion of a constitutive standard of correctness (§8.1.2).

knowledgeable to the highest possible degree whose life consists in nothing more than being pointlessly tortured all day, every day. Such a life, given some obvious assumptions about the story, goes poorly, at least to some extent, and therefore there is more to wellbeing than knowledge.

Can we say more about the specifics of the constitutive connection between knowledge and wellbeing? For example, might this idea be used to defend the eudaimonic ideal of knowledge?

A difficulty arises here because the claim that knowledge has constitutive value is ambiguous. Consider the fact that something might be partially constitutive of wellbeing, in one sense, even though that thing rarely has eudaimonic value. This is because of the nature of constitution. Think of the way in which eggs partially constitute Hollandaise sauce. It's only under certain idiosyncratic conditions that a given egg is part of some portion of Hollandaise sauce. The great majority of eggs are not part of any portion of Hollandaise sauce. For all we have said, knowledge may partially constitute wellbeing in a formally analogous way: only under certain uncommon conditions does a given piece of knowledge have any eudaimonic value at all.¹⁷ Alternatively, one might take the claim that knowledge partially constitutes wellbeing to mean that knowledge *per se* has eudaimonic value.¹⁸ But this too is compatible with rejecting the eudaimonic ideal of knowledge. For one might think that the value of knowledge, although present in every case of knowledge, is not strong enough (as it were) to make it the case that knowledge is always better than ignorance. Indeed, even if knowledge *per se* has constitutive value, it might turn out that knowledge is better than ignorance only under certain idiosyncratic conditions.

We can now articulate an important distinction between constitutive value approaches to the eudaimonic value of knowledge (cf. §1.1.4). **Ideal constitutive value approaches** maintain that, in virtue of the constitutive value of knowledge, knowledge is normally better than ignorance. **Non-ideal constitutive value approaches** maintain that knowledge (sometimes or always) has constitutive value, but that it's not the case that knowledge is normally better than ignorance.

Saying that knowledge (always) has constitutive value does not force us to look at our counterexample to the implausibly formulated eudaimonic ideal of knowledge any differently than we were originally inclined to look at it (§1.1.2). That original thought was something like: ignorance is better for Karen than knowledge. That

¹⁷ This just raises the question of *how* knowledge constitutes wellbeing. Is it like eggs and Hollandaise sauce (a necessary ingredient, in the right proportion), or is it like walnuts and bran muffins (not an essential ingredient, but constitutive in some cases); is it like mass and weight (adding more mass gets you more weight), but not like eggs and Hollandaise (adding more eggs won't get you more Hollandaise)? Some of these issues are explored below (§4.3).

¹⁸ Things that are F are *per se* valuable iff things that are F are valuable "as such" or in virtue of being instances of F. A corollary of this is that if things that are F are *per se* valuable then things that are F are always valuable: if instances of F are *per se* valuable, then all instances of F must be valuable, since if they weren't, it couldn't be the case that instances of F are valuable in virtue of being instances of F—their being instances of *x* wouldn't explain their having value. Note that to say that things that are F are *per se* valuable is not the same as to say that things that are F have final value (Sosa 2001, p. 51).

knowledge has constitutive value is orthogonal to the question of whether or not that thought is correct. Likewise, that knowledge has constitutive value does not conflict with denying the eudaimonic ideal of knowledge (§1.1.3).

However, one might understand the appeal to constitutive value in this way: that we were right in our treatment of the counterexample, in as much as we were right that ignorance can sometimes have good consequences vis-à-vis wellbeing, and that knowledge can sometimes have bad consequences vis-à-vis wellbeing. What we were wrong about, so the argument goes, was the thought that this is enough to show that ignorance is sometimes better than knowledge. Our description of the case was fine; it was our evaluation that was flawed. The constitutive value of knowledge is such that, even though Karen will likely lose her match a result of knowing about her chances, this would be better for her than to live in ignorance. After all, the unexamined life is not worth living! So the argument goes, we mistakenly took Karen to be better off ignorant than she would be were she knowledgeable because we mistakenly ignored the constitutive value of knowledge, focusing on the eudaimonic value of success in one's projects. Expanding our conception of wellbeing so as to include knowledge as a constituent allows us to see that Karen would be better off were she to possess the knowledge that she lacks.

There is a fair and important point being made here, which is that the critic of the eudaimonic value of knowledge should not confine her attention to instrumental value approaches. It's possible that knowledge has final eudaimonic value, in virtue of being partially constitutive of wellbeing. My point is only that one could embrace this idea while rejecting (either formulation of) the eudaimonic ideal of knowledge. Recognizing the constitutive value of knowledge does not make it any more plausible that Karen would be better off knowledgeable. There may be situations in which one suffers on account of knowing something, but where knowledgeable suffering is intuitively better than ignorant bliss. Karen's case isn't like that: knowledge of her opponent's superiority doesn't seem valuable enough to trump the value of (potentially) winning. Better: Karen's case doesn't need to be imagined like that. For example, we can imagine that Karen cares deeply about winning and cares not a whit for knowledge of her opponent's relative abilities. That knowledge has constitutive value doesn't yet tell us anything about the strength (as it were) of the eudaimonic value of knowledge compared to other eudaimonic goods. So to recognize the constitutive value of knowledge does not require that we rethink what we said about Karen's case. Our examination of the eudaimonic ideal of knowledge continues.

1.1.6 *Desire-fulfillment vs. desire-independent theories of wellbeing*

How can we adjudicate the dispute between someone who maintains that knowledge has constitutive value and someone who maintains that knowledge doesn't have constitutive value? And how can we adjudicate disputes about the relative strength (as it were) of the constitutive value of knowledge? Imagine a follower of Hippias (Plato, *Hippias Major*, 289e) who maintains that *owning gold* has constitutive value, and