

Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Title Pages

(p.i) Disagreement **(p.ii)**

(p.iii) Disagreement

OXFORD
UNIVERSITY PRESS

(p.iv) **OXFORD**
UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

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 in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States

by Oxford University Press Inc., New York

© the several contributors 2010

The moral rights of the author have been asserted
Database right Oxford University Press (maker)

First published 2010

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, 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 book in any other binding or cover
and you must impose this same condition on any acquirer

British Library Cataloguing in Publication Data

Data available


Library of Congress Cataloging in Publication Data

Library of Congress Control Number: 2010927207

ISBN 978-0-19-922607-8 (hbk)

978-0-19-922608-5 (pbk)

10 9 8 7 6 5 4 3 2 1



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010


Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Dedication

(p.v) *For Ernie Sosa, despite occasional disagreements* **(p.vi)**



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

(p.ix) Contributors

Earl Conee is Professor of Philosophy at the University of Rochester.

Andy Egan is Associate Professor of Philosophy at Rutgers University.

Adam Elga is Associate Professor of Philosophy at Princeton University.

Catherine Z. Elgin is Professor of the Philosophy of Education at Harvard University.

Richard Feldman is Professor of Philosophy at the University of Rochester.

Richard Fumerton is F. Wendell Miller Professor of Philosophy at the University of Iowa.

Alvin I. Goldman is Board of Governors Professor of Philosophy and Cognitive Science at Rutgers University.


Thomas Kelly is Associate Professor of Philosophy at Princeton University.

Hilary Kornblith is Professor of Philosophy at the University of Massachusetts, Amherst.

Peter van Inwagen is John Cardinal O'Hara Professor of Philosophy at the University of Notre Dame.

Ted A. Warfield is Professor of Philosophy at the University of Notre Dame.

Ralph Wedgwood is Professor of Philosophy, Merton College, University of Oxford. (p.x)



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Introduction

Richard Feldman (Contributor Webpage)

Ted A. Warfield (Contributor Webpage)

In this introduction* we do three things. First, we describe the real-world issue that inspires philosophical reflection on disagreements and show how the issue is typically approached by philosophers. Secondly, we identify a few distinct key questions deserving serious philosophical attention and highlight some of the considerations that come into play in discussing these questions. Thirdly, we briefly introduce the chapters in the volume. We hope that this brief introduction will provide the background necessary to appreciate the chapters in the volume and that these essays help shape the ongoing discussion of the epistemology of disagreement.

1. Disagreements and the Epistemology of Disagreement

Disagreement with intellectual peers is common. Two expert weather forecasters disagree about the weekend forecast. Two equally well-informed economists disagree about the most likely movement in interest rates. Two chess players with the same ranking disagree about whether 'white' stands better in a given board position. The available examples are limitless and range widely over nearly all aspects of life.

Parties involved in disagreements such as these will frequently acknowledge that their interlocutors are as capable of making informed judgments **(p.2)** about the subject matters as they are. This may well inspire critical reflection on the epistemology of the situation. One involved in such a disagreement might focus on the apparent intellectual parity between the disagreeing parties as illustrated by the following imaginary monologue:

I think that interest rates will move down. My colleague disagrees. Furthermore, my colleague has examined all the same information as I have, knows as much as I do about the issue, and is as well trained as I am. This worries me. After all, I take his opinion seriously and would welcome learning that he agrees with my view of this matter. Is it reasonable for me to retain my belief in light of this disagreement? Or is some adjustment rationally required?

Philosophers have routinely approached the kinds of questions raised here with the help of an idealization. In the example, the economist regards his colleague as a “peer.” The notion of a “peer” in play is the unanalyzed ordinary language understanding of that term. In an attempt to focus and regiment the philosophical discussion, those addressing this epistemological issue typically assume that the agents in question are peers in a more specific semi-technical sense. In the stipulative sense of “peer” introduced, peers literally share all evidence and are equal with respect to their abilities and dispositions relevant to interpreting that evidence. Of course, in actual cases there will rarely, if ever, be exact equality of evidence and abilities. This leaves open questions about exactly how conclusions drawn about the idealized examples will extend to real-world cases of disagreement.

Nevertheless, it seems plausible that the idealization will be useful for thinking about at least a fairly large number of the cases that motivate reflection on the epistemology of peer disagreement. Notice, for example, that the idealization is consistent with a maximally wide range of information and skill levels: the idealization is to comparative equality of information and skill and does not involve further idealization to maximal (or even a high level of) information and extraordinary skill. In fact, philosophers have typically focused on “expert” disagreement, but the real-world problem applies more broadly. We must, of course, be careful not to draw conclusions about all cases of disagreement from reflection exclusively on cases of expert disagreement.

However, puzzles about disagreement do extend beyond cases in which there is equality of evidence. Instead of being puzzled about disagreement **(p.3)** with a peer in light of shared information, one might be puzzled about how to respond to disagreements in which there is uncertainty about who has better evidence. Consider the following monologue expressing this apparently different kind of worry.

My weather forecasting competitor is generally as good as I am and today she is disagreeing with me about the weekend forecast. Perhaps she has access to some information I don't have relevant to this forecast. It might, of course, be that I am the one with additional information on this occasion, but without knowing what is accounting for the disagreement I am troubled by the fact that my respected competitor disagrees with me about the forecast.

In this case, the parties are unsure whether one of the two has special knowledge regarding this particular judgment about which they disagree. Whether the idealization involving the philosophical notion of an “epistemic peer” is well suited to evaluating real-world examples of this second sort is an open question.

As the disagreement literature continues to mature, we expect that part of the evolution of the literature will include more reflection on the methodology employed in thinking about real-world epistemological disagreements with the help of the various philosophical idealizations that have been put to work for this purpose.

2. Some Questions and Considerations

Consider two epistemic peers in the stipulative sense introduced earlier: they are familiar with all the same evidence and arguments and are equals with respect to the general intellectual virtues. Assume further that the peer disagreement is out in the open: they know that there is disagreement and that the disagreement is with a peer. Call this “revealed peer disagreement.” What are the epistemic consequences of this situation? Are there any general conclusions that can be drawn about revealed peer disagreement? Does it follow, for example, that at least one of the parties must, because of the disagreement, revise her belief on pain of irrationality? Does it follow that at least one of the parties has an unjustified belief? If either or both of these conclusions do follow, do symmetry considerations lead to the conclusion **(p.4)** that both must revise and/or that both have unjustified beliefs? These are some of the questions on center stage in the disagreement literature. We make no attempt to list all of the interesting questions worth exploring about this type of situation. We instead introduce and briefly explore the contours of a couple of exemplar questions as a way of showing some of the issues involved. The chapters in this collection further explore these and other questions in detail.

The general shape of the debate seems to be as follows. Some philosophers think that general negative consequences result from revealed peer disagreement; that some sort of revision or abandonment of belief is epistemically required in light of revealed peer disagreement. There are disagreements within this camp about just what particular revision is called for, some thinking that peer disagreement requires suspension of judgment about the topic of disagreement or “splitting the difference” with one's disagreeing peer. Given the extent of disagreement in the world, these conclusions have a distinctly skeptical flavor.

Other philosophers think that, though in some cases of revealed peer disagreement negative epistemic consequences attach, this is not the case in all such cases and that therefore there are no general epistemic consequences of revealed peer disagreement. These philosophers are more inclined to think that one can justifiably stick to one's beliefs in the face of some disagreements. One will not find any defense of the view that revealed peer disagreement never has epistemological significance and no defense of the view that one can always safely ignore peer disagreement.

An initial answer to the question "are there any general conclusions that can be drawn about revealed peer disagreement?" as that question has so far been unpacked is easy. The answer is no, at least if a general conclusion would say that one should always modify one's belief in the direction of one's peer's view or that one should always stick to one's initial belief. One will need to build in at least one additional assumption in order to have any hope of arguing for any interesting general conclusion. As many have observed, one will have to add to the case at least the assumption that neither party has special reason to believe that her peer, though a peer, was in some way distracted or inadequately focused on the issue or for some other reason committed a simple performance error in the situation. There will be little or no pressure toward deference to a peer who was not, in forming the belief with which one disagrees, acting as a peer. For **(p.5)** the discussion of the epistemic consequences of peer disagreement to be interesting we need to set aside this sort of possibility.

Exactly what the best way is to set it aside, however, is not immediately obvious. We could either add in a stipulation that both parties know that the disagreeing party is making no such performance error. Or we could add the alternative stipulation that neither party has special reason to believe that the disagreeing party is making such an error. The latter stipulation seems more realistic, while the former seems as though it might helpfully simplify the discussion and increase the focus on more theoretically interesting epistemological issues.

Focus, for now, on a revealed peer disagreement in which neither party has reason to believe that the other has made any sort of performance error. Among the questions that might be asked about this situation: Does it follow that at least one of the parties must, because of the disagreement, revise her belief on pain of irrationality? Does it follow that at least one of the parties has an unjustified belief?

One reason for thinking that at least one of the parties to the disagreement should modify her belief is based on the idea that the revealed peer disagreement itself is new evidence or relevant information. Think of this as a diachronic issue. First one reviews the evidence and arguments and arrives at one's view. One then learns that a peer has reviewed the very same evidence and arguments and arrived at an opposing view. Does this revelation (that a peer disagrees) force revision?

What if each peer confidently reasons as follows: "my peer has looked at the same evidence and reached an opposing view; so she must have made a mistake in thinking about the issue"? This seems like overconfidence. Given that one is dealing with a known peer, one should presumably think that, in general, one is no more likely to be right than one's peer. In the absence of special information that one's peer has made a mistake on this occasion, it might seem that the more rational attitude would be to think that "one of us seems to have made a mistake here but I don't know which one, so I should at least give some weight to the view that I'm mistaken and therefore should revise."

The previous point suggests a line of thinking in favor of the view that at least one party to a peer disagreement (and probably, by symmetry, both) must revise their beliefs in light of revealed peer disagreement. But is this point correct? The reasoning relied on the fact that, in the disputes we are **(p.6)** considering, given the disagreements, at most one of the two parties holds a true belief. It does not seem clear, however, that from this fact combined with peer disagreement we see that belief revision is epistemically required. Why cannot it be, for example, that, though the two parties disagree and though at most one of the two parties is correct, both parties are epistemically permitted to maintain their differing beliefs in light of the available evidence and revealed peer disagreement. Cannot the parties, in other words, be having a disagreement in which neither party is required to revise despite the fact that a peer disagrees?

The answer to this last question turns in part on what has come to be known as "the uniqueness thesis." This is the thesis that a given body of evidence justifies exactly one attitude toward any particular proposition. If this is correct, then peers with the same evidence cannot be justified in holding competing beliefs. In actual cases, people are probably apt to reject the uniqueness thesis, holding that people can be justified in drawing different conclusions from the same evidence. But such judgments may be grounded in the assumption that the parties to the disagreement do not really have the same evidence or, perhaps, in the assumption that a justified belief need only be plausible or consistent with one's other beliefs. Whether a body of evidence can actually justify multiple attitudes is an open question, central to the literature on disagreements. A related question, as noted above, has to do with the evidential impact of the kind of "higher-order" evidence one acquires when one learns that one is in a disagreement with a peer.

Of course, we should not assume that all cases of peer disagreement are alike. There are questions about whether the rationally mandated responses to disagreement are the same in cases in which one's peer disagrees about something one takes to be obviously true and cases in which one has less confidence. There are also questions about whether the proper treatment of disagreements in purely empirical domains should be the same as the treatment of disagreements in moral and other more "evaluative" domains.

Much more could be said about these and other epistemological issues arising in the face of disagreements. We hope to have given an initial feel for some of the issues and considerations in play. The chapters that follow will provide much more food for thought. We turn now **(p.7)** to our final section—a brief summary of the chapters collected in this volume.¹

3. The Chapters

Peter van Inwagen introduces many of the philosophical puzzles about peer disagreement. Starting with a discussion of disagreements in religion, and then extending his comments to philosophical, political, and other disagreements, van Inwagen assesses arguments for and against the skeptical view that the symmetry present in cases of peer disagreements makes suspension of judgment the appropriate attitude. He says at the end that he is unable to give up his beliefs in many of these cases and unable to accept the conclusion that his own beliefs are not rational, but also unable to answer satisfactorily the arguments for the skeptical view.

Hilary Kornblith examines differences among disagreements in different domains. He notes that in some domains disagreements can usually be readily resolved, in others consensus often emerges over time. However, he claims, the dynamics of consensus building in philosophy is quite dissimilar to what is found in mathematics and science. These considerations lead him to the conclusion that philosophical beliefs held in the face of peer disagreement are typically not justified.

Catherine Elgin responds to arguments for the conclusion that participants in persistent peer disagreement ought to suspend judgment about the disputed proposition by noting that "ought implies can" and that belief (and suspension of judgment) are typically not under the relevant kind of voluntary control. She suggests that issues about disagreement are better seen as being about acceptance rather than belief. She contends that continuing to accept propositions in the face of disagreement can have sufficient value to make it rational, and thus that peers can rationally accept conflicting propositions.

Earl Conee formulates and assesses several versions of a principle holding that epistemic peers with the same evidence cannot be justified in holding **(p.8)** different beliefs. He rejects these principles, primarily on the grounds that peers may rationally regard themselves as differing in their basis for rational belief, or their evidence, on the topic. The rationality of their differing perspectives can justify different attitudes toward the disputed proposition.

Richard Fumerton acknowledges that learning of peer disagreement provides a reason against a philosophical belief one holds. But he contends that one can discount this sort of reason against one's view on the grounds that this peer (and philosophers generally) is unreliable on the sort of issue in question. He admits, however, that this way out of the skeptical puzzle posed by disagreements leads to another problem, since the reliability of a philosopher making such an argument is also called into question.

Thomas Kelly distinguishes and assesses a set of views about the rational response to peer disagreement. He raises a series of objections to views according to which one should assign "equal weight" to one's own view and the views of those with whom one disagrees. He concludes by defending "The Total Evidence View," according to which what one is justified in believing in cases of peer disagreement depends upon what is supported by one's total evidence.

Adam Elga argues for the incoherence of unrestricted conciliatory views according to which one ought always to modify one's view in the direction of one's peer in cases of peer disagreement. He contends that such views face incoherence when applied to themselves. He goes on to defend a partly conciliatory view, recommending conciliation on topics other than disagreement itself.

Alvin Goldman argues that there is a plausible though non-standard conception of epistemic relativism under which relativism is compatible with objectivism or absolutism. The key element of this conception is that people in different communities can justifiably accept different principles about reasoning. As a result, people with the same first-order ("material") evidence for a proposition can have divergent but reasonable attitudes toward it.


Ralph Wedgwood discusses moral disagreement and moral relativism, and places the long-standing discussion of these issues in the context of recent discussions of the epistemology of disagreement. He contends that it is rational to have a special sort of "fundamental trust" in one's own moral intuitions, but it is not even possible to have the same sort of trust in **(p.9)** the intuitions of others. As a result, it can be rational for both parties to a peer disagreement to have more confidence in their own views than in the incompatible views of their peers.

Andy Egan extends the discussion of disagreement to puzzling disputes about matters of taste. He contends that these disputes are best explained by saying that what is at stake in them is the self-attribution of certain properties and that resolution of the disputes amounts to the parties making the same self-attributions. He contends that this account explains a special sort of defectiveness to which disputes about taste are subject.

Notes:

* We would like to thank Robert Audi, Marian David, and Nathan King for helpful discussions of these topics.

(1) The chapters in this volume were completed in 2007 and 2008. Unforeseen complications delayed bringing the volume to press.



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

We're Right. They're Wrong

Peter Van Inwagen (Contributor Webpage)

DOI:10.1093/acprof:oso/9780199226078.003.0002

Abstract and Keywords

This chapter introduces many of the central philosophical puzzles about peer disagreement. It starts with a discussion of disagreements in religion, and then extends the discussion to philosophical, political, and other disagreements. It assesses arguments for and against the skeptical view that the symmetry present in cases of peer disagreements makes suspension of judgment the appropriate attitude. The author of the chapter is unable to give up his beliefs in many of these cases and unable to accept the conclusion that his own beliefs are not rational, but is also unable to answer satisfactorily the arguments for the skeptical view.

Keywords: disagreement, peer disagreement, exclusivism, rationality

Discussions of “exclusivism” began with discussions of religious exclusivism. It was soon recognized, however, that, if there was such a thing as religious exclusivism, there were also such things as philosophical, political, and scientific exclusivism (and no doubt other forms as well).¹

I shall first discuss religious exclusivism, since religion is the area in which the concept of exclusivism was first applied, and the area in which our intuitions about this concept are likely to be the most clear. When I have done that, I will look at the ways the concept might be extended to or applied in other areas.

I will not attempt to define the term 'religion,' a large project that would take us far afield.² I am going to have to assume that we all have some sort **(p.11)** of grasp of this term and that we all mean more or less the same thing by it. But note that in what follows I use the word 'religion' as a count-noun, not as a mass term. I speak not of some phenomenon called 'religion' but of various institutions called 'religions.'

I will, however, present some definitions, definitions of my terms of art. Let us use 'Ism' as a dummy term that can be replaced by the name of any religion. I will say that the religion Ism is *weakly exclusivist* if it requires its adherents to subscribe to the following two theses (or perhaps it would be more realistic to say: if these two theses are logical consequences of the theses that Ism requires its adherents to subscribe to):

(1) Ism is logically inconsistent with all other religions. That is to say, any system of belief or thought (besides Ism itself) that is logically consistent with Ism is not a religion. If, for example, Christianity is weakly exclusivist, then any Christian who thinks that Berkeley's metaphysic is consistent with Christianity is logically committed to the thesis that Berkeley's metaphysic is not a religion. And if, according to Hinduism, Hinduism and Islam are both religions, and if it is a tenet of Hinduism that one can consistently be both a Hindu and a Muslim, then Hinduism is not a weakly exclusivist religion.

(2) According to Ism, it is rational to accept Ism. More precisely: it is rational for people whose epistemic situation is typical of the epistemic situations of Ismists to accept Ism. More precisely still: people whose epistemic situation is typical of the epistemic situations of Ismists and who accept the teachings of Ism do not thereby violate any of their epistemic duties; it is at least epistemically permissible for such people to accept the teachings of Ism. Two comments: (a) By someone who “accepts the teachings of Ism,” I mean someone who accepts all the theses or propositions that Ism requires its adherents to accept. Since Ism may well require other things of its adherents than that they accept certain propositions, there is obviously a distinction to be made between someone who accepts the teachings of Ism and an Ismist. And, of course, it will be possible to be an Ismist—in several senses of “be an Ismist” that I can think of—and not to accept the teachings of Ism. (b) No doubt it would be hard to find an example of a religion (or a system of belief of any kind) that did *not* have this feature. But one might imagine (one could probably **(p.12)** actually point to) some defiantly anti-rational religion or philosophy or *Weltanschauung* that conceded, and even gloried in the concession, that it was positively irrational of its adherents to accept its teachings. (Kierkegaard has certainly been accused—but only, I think, by those who dislike him very much—of approvingly ascribing this feature to Christianity.) In any case, if a religion does not hold that its adherents are (typically, at least) rational in accepting its teachings, that religion will not be what I am calling a weakly exclusivist religion.

It obviously does not follow from a religion's being weakly exclusivist that it requires its adherents to believe that the adherents of other religions are necessarily irrational. What does follow is that, if Ism is a weakly exclusivist religion, those who accept its teachings will (at least if they are logically consistent and capable of a little elementary logical reasoning) reach the following conclusion: "The teachings of all other religions are at least partly wrong, and it is rational for us to believe this about the teachings of all other religions." Suppose, for example, that Christianity is a weakly exclusivist religion. Suppose that I, a Christian, consider some other religion—Zoroastrianism, let us say. Since my religion is weakly exclusivist, one of the teachings of my religion is that, given that Zoroastrianism is a religion (which we shall stipulate: that Zoroastrianism is a religion and that it is rational to believe that it is are things that can be objectively established to anyone's satisfaction), Zoroastrianism is inconsistent with my religion. And, on this matter, what my religion teaches is demonstrably right: according to the Zoroastrians, for example, evil is an uncreated power (even if, as some have maintained, this is not what Zoroaster himself believed), and, according to Christians, evil is not an uncreated power. I can easily infer from these things that Zoroastrianism teaches something false (even if I do not know that Zoroastrianism teaches that evil is an uncreated power and in fact have not the faintest idea what Zoroastrianism teaches). And, if it is rational for me to accept Christian teaching (which is implied by Christianity and which I therefore believe, if Christianity is a weakly exclusivist religion), it is rational for me to accept what it is rational for me to believe follows logically from it. If, therefore, Christianity is a weakly exclusivist religion, one of its teachings will be (or its teachings will logically imply) that it is rational for me to believe that Zoroastrianism teaches something false. The point is easily generalized: it is a teaching **(p.13)** of any weakly exclusivist religion (or an immediate logical consequence thereof) that it is rational for its adherents to believe that all other religions teach something false.

Now a second definition. Let us say that a religion, Ism, is *strongly exclusivist* if it is weakly exclusivist *and* it teaches (or its teachings entail) that, for any other religion, it is not rational for anyone who is in an epistemic situation of the sort in which Ismists typically find themselves to accept the teachings of that religion.

It is consistent with Ism being a strongly exclusivist religion that it not require its adherents to believe that the adherents of other religions necessarily or even typically violate the norms of rationality. If Islam is a strongly exclusivist religion, its adherents must regard any Muslim who apostatizes and becomes an adherent of some other religion (and who had been in an epistemic position *typical* of Muslims) as irrational. But, consistently with Islam's being a strongly exclusivist religion, a well-informed and orthodox Muslim might regard a pagan or Christian or Jew who had never been properly exposed to the teachings of Islam (or whose knowledge of the world was in some other way incomplete) as at least possibly rational, as not *ipso facto* irrational. It could even be that a well-informed and orthodox adherent of the strongly exclusivist religion Ism thought that all non-Ismists were perfectly rational. (He might ascribe “invincible ignorance” to all non-Ismists, or might think that one is an Ismist if and only if God has bestowed on one an infusion of grace that is entirely independent of one's epistemic condition prior to the moment of its bestowal.)

But a sterner attitude on the part of adherents of a strongly exclusivist religion toward adherents of other religions is also possible. St Paul thought that paganism—at least the paganism with which he was familiar, the paganism of the classical Mediterranean world—was epistemically permissible for no one, and such a belief is certainly consistent with strong exclusivism in religion. If a religion takes this position with respect to *all* other religions, I will call it a *very strongly exclusivist* religion. If, for example, Christianity is a very strongly exclusivist religion, then the teachings of Christianity entail that no religion but Christianity is epistemically permissible for any human being of any culture in any era. It would follow, of course, that, if Christianity is a very strongly exclusivist religion, then Christianity teaches that, before the founding of Christianity, it was not rational for anyone to **(p.14)** accept the teachings of any religion. (This might be called *diachronic* strong exclusivism. Refinements are possible. A Christian might think that it was rational to accept the teachings of Judaism before the Incarnation but not afterwards, and not rational to accept the teachings of any other religion at any time.)

I mention the idea of a very strongly exclusivist religion for the sake of logical completeness. I am fairly sure that there are none. I do not deny that some religions teach that *certain* religions are not epistemically permissible for anyone. (Since St Paul's attitude toward classical paganism was expressed in what Christians regard as an inspired text, it is a defensible position that Christianity teaches the epistemic impermissibility of classical paganism.) And I do not deny that individual adherents of various religions (Christianity and Islam, for example) may have believed that no other religion than their own was epistemically permissible for anyone. I do not even want to deny that for some religions (including my own) it may have been at certain points in history that all adherents of those religions then alive (or all of them who had considered the point) believed both that all other religions were epistemically impermissible for anyone and that this epistemological thesis was a teaching of their religion. What I do deny—at least I strongly doubt whether this is the case—is that this *is* the teaching of any religion. (It not infrequently happens that all the adherents of a religion who are alive at a certain time believe falsely of some proposition, some proposition they all happen to believe, that it is a teaching of their religion.)

The important forms of religious exclusivism are strong and weak exclusivism (or, since strong exclusivism entails weak exclusivism, strong exclusivism and *mere* weak exclusivism). Again, further refinements of these concepts are possible. One might, for example, want to take account of the fact that, according to Islam, the epistemic position of Christians and Jews, while weaker than that of Muslims, is stronger than that of Buddhists and Hindus—and was stronger before the Koran was revealed to the Prophet than it was afterward. But one must at some point leave off making ever finer distinctions, however congenial one may find that occupation to be if one is a philosopher.

This, then, is religious exclusivism (or “alethic” religious exclusivism, religious exclusivism in the matter of the possession of truth). There are at least four “areas” other than religion in which the concept of alethic **(p.15)** exclusivism has obvious application: philosophy, politics, art, and science. (And I might add a fifth: everyday life.) I will say something about all four.

All philosophers would seem to be weak exclusivists as regards their own philosophical positions. Consider, for example, a representative philosopher, Phoebe, who accepts a certain philosophical position that, borrowing a device from our discussion of religious exclusivism, I will call Ism. Phoebe is, I shall say, a *weak exclusivist* as regards Ism if she believes, first, that there are philosophical positions that are inconsistent with Ism, and if she believes, secondly, that her own acceptance of Ism is rational (that her acceptance of Ism is rationally permissible for her). When I say that Phoebe believes that there are philosophical positions that are inconsistent with Ism, I do not mean only that she believes that some of the denizens of the Platonic heaven are propositions inconsistent with Ism and that some of these propositions count as “philosophical positions,” albeit not ones actually *held* by anyone down here in the world of flux and impermanence; I mean that she thinks that there *are* people who hold philosophical positions inconsistent with Ism. Note that what I have said is not in every respect parallel to what I said when I defined weak *religious* exclusivism: I have not said

Ism is logically inconsistent with all other philosophical positions.
That is to say, any thesis (besides Ism itself) that is logically
consistent with Ism is not a philosophical position.

Philosophical positions are not logically related to one another as religions are logically related to one another, or even as the sets of teachings of the various religions are logically related to one another, since one philosophical position can easily be consistent with another philosophical position: intuitionism in ethics is consistent with mathematical intuitionism (at least assuming that neither is a necessary falsehood), and Cartesian dualism is consistent with dualism (with the same qualification). I will not, moreover, speak of philosophical positions *themselves* as being exclusivist; I shall rather speak of *people* adopting an exclusivist *stance* or *attitude* in respect of philosophical positions they hold, for statements of philosophical positions do not generally contain clauses describing the epistemic situations of people who hold those positions and competing positions. But the fact that “philosophical weak exclusivism” and “religious weak exclusivism” (**p.16**) are not parallel in these two respects raises no important barrier to applying the concept of exclusivism to philosophical positions.

But why *do* I call Phoebe's position as regards Ism an "exclusivist" position? Simply because (1) logic requires her to believe that all those who accept positions logically inconsistent with Ism believe something false, and (2) she believes that she is rational in accepting Ism, she (almost certainly) believes that it is rational for her to believe that those who accept positions logically inconsistent with Ism believe something false. "Almost certainly?" Well, if Phoebe has the latter belief, she presumably got to it by an application of the following epistemological principle or one very like it: 'If one accepts p and believes that one's acceptance of p is rational, and if one accepts some immediate and self-evident logical consequence of p , and accepts it *because* one sees that it is an immediate and self-evident logical consequence of p , one should also believe that one's acceptance of that logical consequence is rational.' There may be reasons to reject this principle (some would say it leads to skepticism), but I will not go into the delicate issues involved in the question of whether one should accept this principle. I will assume without further argument the general thesis of which the following thesis is a special case:

If Phoebe is a philosopher who accepts materialism, and if she believes that her acceptance of materialism is rational, then she believes that it is rational for her to believe that anyone who accepts dualism believes something false.

(If I heard some philosopher say, "I accept materialism, and I believe that it is rational for me to accept materialism, but I do not believe that it would be rational for me to believe that all dualists believe something false," I would certainly cast a very puzzled glance in the direction of that philosopher.)

Similar points can be made about weak exclusivism in politics, art, and science. Consider these three theses: 'The so-called right to privacy supposedly implicit in the US Constitution was made up out of whole cloth by twentieth-century jurists as a legal rationalization of rulings they wanted to make on moral grounds;' 'Technical facility is not everything; Tolstoy was a greater novelist than Flaubert despite the clumsiness of his narrative technique;' 'Neanderthal Man was a genetic dead end; no modern human being has Neanderthal ancestors' (this last statement, by the way, is—or at least has been—no less controversial than the **(p.17)** first two). Anyone who accepts any proposition of the sorts these three propositions exemplify—controversial political or aesthetic or scientific propositions—believes (1) that other people accept propositions logically incompatible with that proposition, and (2) that it is rational for him to believe that they are mistaken to accept those propositions. Physical anthropologists who believe that Neanderthal Man was a genetic dead end believe that their colleagues who think that the modern human genome contains Neanderthal genes are in error—and believe that it is rational for them to believe that those people are in error.

I should be surprised if anyone were to deny that those of us who have opinions about philosophical, political, aesthetic, and scientific matters adopt a weakly exclusivist position toward them, at least in those cases in which those “opinions” are not so self-evidently true that there is no disagreement about them. I want now to turn to a more controversial question: what part does *strong* exclusivism play in philosophy, art, politics, and science? We may say that one adopts a strongly exclusivist stance toward one's adherence to a position in any of these areas—and we may as well continue to use the device we have been using and call this position ‘Ism’—if one adopts a weakly exclusivist position toward one's adherence to Ism, and, moreover, thinks that it would *not* be rational for anyone in one's own epistemic situation to adopt or adhere to any position that is (obviously and uncontroversially) inconsistent with Ism.

Suppose, for example, that Dan is an adherent of Darwinism, of the Darwinian theory of evolution. Dan adopts a *strongly exclusivist* stance toward his adherence to Darwinism if (in addition to the requirements of adopting a weakly exclusivist stance toward this position) he believes that anyone whose epistemic situation or condition is the same as his and who accepts any thesis or proposition (obviously) inconsistent with Darwinism is irrational (that is, has adopted a position that is rationally indefensible). Suppose Dan were to read the following words, which I quote from a book by the English biologist Brian Goodwin.

Despite the power of molecular genetics to reveal the hereditary essences of organisms, the large-scale aspects of evolution remain unexplained, including the origin of species. There is “no clear evidence . . . for the gradual emergence of any evolutionary novelty,” says Ernst Mayr, one of the most eminent of contemporary evolutionary biologists. New types of organisms simply appear on the evolutionary scene, persist for various periods of time, and then become extinct. So Darwin's **(p.18)** assumption that the tree of life is a consequence of the gradual accumulation of small hereditary differences seems to be without significant support. Some other process is responsible for the emergent properties of life, those distinctive features that separate one group of organisms from another—fishes and amphibians, worms and insects, horsetails and grasses. Clearly something is missing from biology.³

What should Dan think about the thesis expressed in this passage, given his strongly exclusivist stance toward Darwinism? It seems that he must think that Goodwin is either ignorant of something evidentially relevant to questions about the mechanisms of evolution—something that he, Dan, knows—or else is irrational. And, since Goodwin is a professional biologist, and in fact a respected biologist, it is unlikely that Dan can plausibly ascribe Goodwin's position on Darwinism to factual ignorance. It would be far more plausible for Dan to contend that Goodwin has failed to believe something that any *rational* person in Goodwin's epistemic situation would believe. And I think it is evident that there are Darwinists who adopt a strongly exclusivist stance toward Darwinism. I would in fact go so far as to say that almost every Darwinist is a strongly exclusivist Darwinist. The merely weakly exclusivist Darwinist is very rare indeed. (I do not mean to suggest that very many Darwinists go about *saying* that people like Goodwin are irrational. That, after all, would not be very polite. But that, surely, is what they think, or what they would think if they applied their views with rigorous consistency.)

What happens in science certainly happens in politics. An extreme example is provided by the case of “Holocaust deniers”: no one, I believe—no one *at all*—who believes that the Nazis murdered six million Jews fails also to believe that those trained historians who deny this thesis (who, say, put the figure far lower, or who attempt to qualify the word ‘murder’) subscribe to a thesis that is rationally indefensible. But I need not defend my contention that there are strong exclusivists in political matters by reference to an extreme thesis like Holocaust denial. Current “red-blue” disagreements in American politics have the same feature. Consider, for example, the proposition

The American-led invasion of Iraq in 2003 was morally and politically indefensible.

(p.19) There are many who accept this proposition and many who accept its denial. But almost no one who accepts it believes that an intelligent, rational person who was in possession of a reasonable proportion of the relevant and available facts would accept its denial. And almost no one who denies it believes that a person who was intelligent and in possession of a reasonable proportion of the relevant and available facts would accept it. In short, each “side” in the dispute regards the other side as not only mistaken (adherence to the principle of non-contradiction requires that much) but irrational (more exactly, irrational unless ignorant of relevant and available facts or so intellectually deficient as to be excused from normal epistemic obligations by the “ought implies can” principle).

There are, therefore, strong exclusivists in science, history, and politics. But what about philosophy? What is the place of strong exclusivism in philosophy? I will leave aside questions about the past, questions about what may have been the attitude of philosophers in the past concerning those who took other philosophical positions than theirs, and concentrate on philosophy as it is now. (And by philosophy as it is now, I mean analytical philosophy as it is now. Generalizations about “philosophers” in what follows are restricted to this domain, the only domain about which I have any real information.) Are there philosophers who would endorse instances of the following schema (in this schema, p represents the name of some substantive philosophical thesis — ‘nominalism’, ‘utilitarianism’, ‘the compatibility of free will and determinism,’ and so on):

I accept p , and I regard all trained philosophers who are in my epistemic position (that is, who are aware of the arguments and other philosophical considerations I am aware of) and who accept the denial of p as irrational. Such philosophers are in violation of their epistemic obligations. They are comparable to ordinary, educated people of the present day who believe that cigarette smoking does not cause lung cancer or that the positions of the stars and planets at the moment of one's birth determine one's fate.

I doubt whether many philosophers would say anything along these lines. The more interesting question is whether philosophers generally believe things that commit them to something like this position.

And there are indeed some very plausible propositions that would seem to have this consequence. For example, **(p.20)**

If it is rational for a person to accept a certain proposition, it cannot also be rational for that person (at the same time, in the same circumstances) to accept its denial.

Here is one consequence of this proposition. If Alice says that Ted's belief that the St Joseph River is polluted is rational, and Winifred says that *if* Ted believed (in these very circumstances) that the St Joseph River was not polluted, this belief *would* be rational, Alice and Winifred cannot both be right.

A similar principle:

If, for some proposition, a person accepts neither that proposition nor its denial, and is trying to decide whether to accept that proposition, to accept its denial, or to continue to accept neither, it cannot be true *both* that it would be rational for him to accept that proposition *and* that it would be rational for him to accept its denial.

A past-tense version:

If a person has just accepted a certain proposition (has just reached the conclusion that that proposition is true), and if it was rational for that person to accept that proposition, it cannot be that it would have been rational for him to accept its denial.

These principles, are, as I said, very plausible. Why? What underlies their plausibility, I think, must be some such argument as this.

Consider any person as he is at a certain moment. That person has available to him, at that time, a body of evidence, his total evidence at that moment. Call it *E*. Of metaphysical necessity, *E* has the following property (of itself, regardless of who may have it) with respect to any proposition *p*: either (i) it would be rational for anyone to accept *p* on the basis of *E*, or (ii) it would be rational for anyone to accept the denial of *p* on the basis of *E*, or (iii) it would be not be rational for anyone to accept either *p* or the denial of *p* on the basis of *E*. This 'or' is exclusive. It *cannot* be that it would be rational to accept *p* on the basis of *E* *and* rational to accept the denial of *p* on the basis of *E*. Loosely speaking, a body of evidence cannot have the power to confer rationality on both a proposition and its denial. And, finally, it is rational for one to accept a proposition at a certain moment if and only if one's total evidence at **(p.21)** that moment bears this impersonal "confers rationality on" relation to that proposition.

This is a plausible argument, but one might find difficulties with it. Suppose that *E* is logically inconsistent, but in a very subtle way, and that the person whose evidence *E* is could not be expected to see this inconsistency. Suppose that, as a result of the inconsistency, there exist (platonically speaking) a valid derivation of *p* from *E* and a valid derivation of the denial of *p* from *E*. Might it not be that, if the person were aware of the former derivation (but not the latter), it would be rational for him to accept *p*, and that, if he were aware of the latter derivation (but not the former), it would be rational for him to accept the denial of *p*? This objection may be met as follows: awareness of one of the derivations (if the person has it) is a part of the person's total evidence. That is, if *E* is a person's total evidence, and he then becomes aware that *p* may be validly deduced from *E*, *E* is no longer his total evidence: his total evidence is now (at least) *E* plus the proposition that *p* can be validly deduced from *E*. Therefore, in the two circumstances we have imagined, the person does not have the same total evidence. (And, of course, if he were aware of both derivations, it would certainly not be true that it would be rational for him either to accept *p* or to accept its denial.) If the argument we are considering is correct, it confers validity not only on "single-person" principles like those I have set out, but on the following principle:

If two people have the same evidence, and if one of them accepts a certain proposition and the other accepts its denial, at least one of them is not rational: either it is not rational for the one to accept that proposition, or (inclusive) it is not rational for the other to accept its denial.

It is not hard to see that, if this principle is correct, then weak exclusivism entails strong exclusivism.

Suppose, for example, that I am a platonic realist and that my attitude toward my platonic realism is one of weak exclusivism. And suppose that my colleague Sally is a nominalist and that her epistemic circumstances, as they bear on the nominalism-realism question, are the same as mine. (That is, she and I are aware of all the same relevant extra-philosophical facts and theories, the same facts of everyday life and the same scientific **(p.22)** facts and scientific theories; she and I, moreover, are aware of the same philosophical considerations that are relevant to the nominalism–realism debate: the same distinctions, the same arguments, and so on.) Can it be that Sally's position is rational? Is it rational for her to accept nominalism? If I, as I have said I do, adopt a weakly exclusivist stance toward my own acceptance of realism, and if I accept any of the epistemological principles we have been canvassing, I must say that it is not rational for her to accept nominalism. If I am a weakly exclusivist realist, then I believe that it is rational for me to accept realism. And, since the evidence relevant to the nominalism–realism dispute that is at Sally's disposal is the same as the evidence at my disposal, it cannot be—by the principle we have just set out—rational for me to accept realism and for Sally to accept nominalism. Here is a second argument for this conclusion, an argument that appeals to a “single-person” principle. It is rational for me to be a realist. It would, therefore, be rational for Sally to be realist, since the evidence she and I have that is relevant to the nominalism–realism question is the same. (Here I appeal to the principle: if it is rational for a person whose total evidence is E to accept p , then it would be rational for anyone whose total evidence was E to accept p . But this principle is, I believe, obviously correct.) But if it is rational for Sally to accept nominalism (on the evidence she has), then it would be both rational for her to accept nominalism *and* rational for her to accept realism on the same evidence. And this is ruled out by the “single-person” principles. Therefore, it is not rational for Sally to accept nominalism. The conclusion of this argument is easily generalized: if any of the epistemological principles we have laid out is valid, then weak exclusivism entails strong exclusivism. And, I remind you, almost all philosophers adopt a weakly exclusivist stance toward the philosophical propositions they accept.

Let us look at the consequences of this fact. Suppose you are a philosopher who accepts various philosophical propositions that are rejected by other trained philosophers. (And this is the normal case. Very few of the philosophical propositions that are accepted by some philosopher are accepted by all philosophers.) Let *Ism* be any such proposition and let *Nism* be its denial. If you accept any of the above principles, you must, after due reflection on the fact that they imply that weak exclusivism entails strong exclusivism, reach one of the following conclusions. **(p.23)**

- (1) It is not, after all, rational for me to accept *Ism*.
- (2) It is not rational for any trained philosopher to accept *Nism*.
- (3) Some trained philosophers who accept *Nism* are in epistemic circumstances that are inferior to mine in the matter of deciding what to believe about *Ism* and *Nism*.
- (4) Some of the philosophers who accept *Nism* are less intelligent than me (or labor under some other relevant cognitive disadvantage; lack of philosophical talent, perhaps).

Let me make some remarks about (3) and (4).

If I accept option (3), I must conclude that at least some of the philosophers who accept *Nism* are unaware of or have not fully grasped some relevant argument or analysis or distinction I am aware of and understand. And I must suppose that the fact that they are in some such inferior epistemic circumstances is not something for which they can be blamed, that their condition is not a result of their having failed to fulfill some epistemic obligation. (Here is an analogous situation in ordinary life: you believe that Jane is honest, and I think she is a crook. Up till a moment ago, you and I had the same evidence in this matter, and we both believed, rationally, given this evidence, that she was honest. A moment ago I stumbled, by merest chance, upon a well-hidden document that demonstrates beyond any possibility of doubt that she has been defrauding her employers, a document that no one could blame you for not being aware of.)

As to (4), if I accept this option, I shall not believe that (all) the philosophers who disagree with me are irrational. I shall, however, excuse them from the charge of irrationality, only because of my allegiance to the "ought implies can" principle. Although each of these philosophers is laboring under the burden of some cognitive deficiency, each of them is doing the best he can according to his own dim lights. And this same deficiency, in each case, blinds the philosopher who labors under it to the fact that he is in this respect, cognitively deficient.

In any real situation, all these options can seem extraordinarily unappealing. I will mention a case that I have used for similar purposes on other occasions. I ask you to consider the case of David Lewis and me and the problem of free will. I am an incompatibilist and David was a compatibilist. David and I had many conversations and engaged in a rather lengthy correspondence on the matter of compatibilism and incompatibilism, and, **(p.24)** on the basis of these exchanges—not to mention his wonderful paper “Are We Free to Break the Laws?”—I am convinced beyond all possibility of doubt that David understood perfectly all the arguments for incompatibilism that I am aware of—and all other philosophical considerations relevant to the free-will problem (philosophical distinctions and philosophical analyses, for example). It seems difficult, therefore, to contend that, in this matter, he was in epistemic circumstances inferior to mine. What, after all, could count as the ingredients of a person's epistemic circumstances (insofar as those circumstances are relevant to philosophical questions) but that person's awareness of and understanding of philosophical arguments (and analyses and distinctions and so on)? If philosopher A and philosopher B are both investigating some philosophical problem, and if each is aware of (and understands perfectly) all the arguments and distinctions and analyses—and so on—that the other is aware of, how can the epistemic position of one of these philosophers vis-à-vis this problem be inferior to that of the other? And one could hardly maintain that David was stupid or lacking in philosophical ability or that he labored under any other cognitive deficiency relevant to thinking about the problem of free will. (Not, at any rate, unless all human beings labor under this deficiency.) At the same time, I am unwilling to say that my own allegiance to incompatibilism is irrational. I can only conclude that I am rational in accepting incompatibilism and that David was rational in accepting compatibilism. And, therefore, we have at least one case in which one philosopher accepts a philosophical proposition and another accepts its denial and in which each is perfectly rational. It is, moreover, a case in which the epistemic circumstances of neither philosopher, as they touch on the question whether to accept this proposition or its denial, are inferior to the epistemic circumstances of the other. (And in which neither philosopher labors under the burden of any cognitive deficiency from which the other is free. I know that David labored under no such deficiency. I like to think that I do not.) And, therefore, the epistemological principles I have laid out, the principles that allowed us to deduce strong exclusivism from weak exclusivism, must be wrong.

This conclusion seems to me to be inescapable—if one's epistemic circumstances (those relevant to philosophical inquiry) are indeed defined entirely by the “philosophical considerations” (arguments, distinctions, and so on) one is aware of and understands. A moment ago, I asked, **(p.25)** rhetorically, “If philosopher A and philosopher B are both investigating some philosophical problem, and if each is aware of (and understands perfectly) all the arguments and distinctions and analyses the other is aware of, how can the epistemic position of one of these philosophers vis-à-vis this problem be inferior to that of the other?” If it seems to one that this rhetorical question is unanswerable, this must be because one regards evidence, the stuff of which one's epistemic condition is made, as essentially *public*. It must be because one regards evidence as “evidence” in the courtroom-and-laboratory sense. And, in matters philosophical, public evidence is that which is expressible in language. A piece of evidence for a philosophical proposition, if it is public evidence, is something that could be expressed as a bit of text (a part of an essay or book or letter), and one who had read and understood the bit of text that embodied it would “have” that piece of evidence; it would be a component of his or her epistemic condition.

It is, however, reasonable to suppose that this conception of “evidence” (if evidence is indeed the stuff of which one's epistemic condition is made, if A and B are in the same epistemic condition just in the case that they “have” the same evidence) is overly restrictive. One of the reasons that constitute the reasonableness of supposing this is that there seem to be plausible examples of “having evidence” that do not conform to the courtroom-and-laboratory paradigm of evidence. I sometimes know that my wife is angry when no one else does, for example, and I cannot explain to anyone how I know this—I cannot give what Plato would call an “account” of what underlies my conviction that she is angry. It seems to me to be plausible to say that in such cases my belief that my wife is angry is grounded in some body of evidence, evidence that lies entirely within my mind and that I cannot put into words. A second example is provided by the case of the chicken sexer, beloved of epistemologists in the far-off days of my graduate studies. (Can anyone tell me whether there *are* chicken sexers? Those of my students who were raised on farms have given conflicting testimony on this matter.) Mathematicians are often intuitively certain that some mathematical proposition is true, although they are unable to prove it. (Gödel, I understand, was convinced that the power of the continuum was \aleph_2 .) Since they often later do discover proofs of these propositions, it seems likely that, prior to their discovery of the proofs, they had some sort of evidence for the truth of those propositions. **(p.26)**

There are, therefore, arguments by example for the conclusion that, in everyday life, at least, and perhaps in mathematics, evidence is not always of the public sort, that some evidence is not *exportable*, that some evidence cannot be passed from one person to another. And what is true of “everyday” evidence (and perhaps of “mathematical” evidence) may also be true of the evidence that grounds philosophical convictions. Some “philosophical” evidence, too, may not be exportable. I can give an argument for the thesis that some philosophical evidence has this feature. The argument takes the form of a dialectical challenge to any philosopher who denies it. Consider, for example, the body of public evidence that I can appeal to in support of incompatibilism (arguments and other philosophical considerations that can be expressed in sentences or diagrams on a blackboard or other objects of intersubjective awareness). David Lewis “had” the same evidence (he had seen and he remembered and understood these objects) and was, nevertheless, a compatibilist. If I know, as I do, that David had these features (and this feature, too: he was a brilliant philosopher), that he had these features is itself evidence that is (or so it would seem to me) relevant to the truth of incompatibilism. Should this new evidence not, when I carefully consider it, lead me to *withdraw* my assent to incompatibilism, to retreat into agnosticism on the incompatibilism/compatibilism issue? This is a question I have discussed elsewhere. Here I will offer only the following brief argument. One's evidence is supposed in some way to direct the formation of one's beliefs. If it was rational for David to be a compatibilist, therefore, it must be that his evidence did not direct him *away* from compatibilism. If it did not direct him away from compatibilism, it did not direct him *toward* incompatibilism. But my evidence is his evidence. It must therefore be that *my* evidence does not direct *me* toward incompatibilism. How then can it be rational for me to be an incompatibilist?

The difficulty of finding anything to say in response to this argument, taken together with my unwillingness to concede either that I am irrational in being an incompatibilist or that David was irrational in being a compatibilist, tempts me to suppose that I have some sort of interior, incommunicable evidence (evidence David did not have) that supports incompatibilism.

If I succumb to this temptation, if I allow even that it is possible that I have such evidence, then the above demonstration that weak exclusivism entails **(p.27)** strong exclusivism fails. I can, consistently, believe that it is rational for me to accept Ism and rational for other philosophers to accept Nism. I can, without logical inconsistency, maintain that the Nismists are, through no fault of theirs, in epistemic circumstances that are (vis-à-vis the Ism/Nism question) inferior to mine. Owing to some neural accident (I might say) I have a kind of insight into the, oh, I don't know, entailment relations among various of the propositions that figure in the Ism/Nism debate that is denied to the Nismists. I *see*, perhaps, that p entails q (although I am unable to formulate this insight verbally) and they are unable to see that p entails q . And this insight really is due to a neural quirk (to borrow a phrase Rorty used for a different purpose). It is not that my cognitive faculties function better than theirs. Theirs are as reliable as mine. But theirs are not identical with mine, and, in this case, some accidental feature of my cognitive architecture has enabled me to see the entailment that is hidden from the Nismists.

In the end, though, this idea, tempting as it is to me, is hard to believe. After all, I accept *lots* of philosophical propositions that are denied by many able, well-trained philosophers. Am I to believe that in every case in which I believe something many other philosophers deny (and this comes down to: in every case in which I accept some substantive philosophical thesis), I am right and they are wrong, and that, in every such case, my epistemic circumstances are superior to theirs? Am I to believe that in every such case this is because some neural quirk has provided me with evidence that is inaccessible to them? If I do believe this, I must ask myself, is it the same neural quirk in each case or a different one? If it is the same one, it begins to look more a case of “my superior cognitive architecture” than a case of “accidental feature of my cognitive architecture.” If it is a different one in each case—well, that is quite a coincidence, isn't it? All these evidence-providing quirks come together in just one person, and that person happens to be me.

It seems more plausible to say (to revert to the example of David Lewis and myself) that David and I have the same evidence in the matter of the problem of free will, and to concede that this entails that either we are both rational or neither of us is.

The position that we are both rational, however, is hard to defend. If I suppose that we are both rational, I hear W. K. Clifford's ghost whispering an indignant protest. Something along these lines (Clifford has evidently **(p.28)** acquired, post mortem, a few turns of phrase not current in the nineteenth century):

If you and Lewis are both rational in accepting contradictory propositions on the basis of identical evidence, then *you* accept one of these propositions—incompatibilism—on the basis of evidence that does not direct you toward incompatibilism and away from compatibilism. (For, if it did, it would have directed *him* away from compatibilism, and it would not have been rational for him to be a compatibilist.) But of all the forces in the human psyche that direct us toward and away from assent to propositions, only rational attention to relevant evidence *tracks the truth*. Both experience and reason confirm this. And, if you assent to a proposition on the basis of some inner push, some “will to believe,” if I may coin a phrase, that does not track the truth, then your propositional assent is not being guided by the nature of the things those propositions are *about*. If you could decide what to believe by tossing a coin, if that would actually be effective, then, in the matter of the likelihood of your beliefs being true, you might as well do it that way.


Here I confess my predicament—as a philosopher who holds particular views, as a citizen who casts his vote according to the dictates of certain political beliefs, as an adherent of one religion among many. (For, although I have been talking about philosophy for some time now, what I have said is equally applicable to politics and religion.) I am unwilling to listen to the whispers of Clifford's ghost; that is, I am unwilling to become an agnostic about everything but empirically verifiable matters of fact. (In fact I am unable to do that, and so, I think, is almost everyone else; as Thoreau said, neither men nor mushrooms grow so.) And I am unable to believe that my gnosticism, so to call it, is irrational. I am, I say, unwilling to listen to these whispers. But I am unable to answer them.

Notes:

(1) The “exclusivism” discussed in this chapter has to do with truth. A religion is exclusivist in this sense if it represents itself as the only religion that has the truth. But other forms of religious exclusivism are possible. Some call a religion exclusivist, for example, if it regards itself as the only possible path to salvation. In this sense, only a religion can be exclusivist, at least if ‘salvation’ is understood to pertain to spiritual matters. But there are analogues of this second sort of exclusivism in philosophy and science. According to Marxism, not only does Marxism provide the only correct account of the historical unfolding of societies and cultures, but all philosophies with which it is in competition are mere repressive “ideologies,” systems of ideas whose existence is to be explained not on intellectual grounds but in terms of their economic function (which is to conceal from the economically oppressed the fact of their oppression). Some philosophers and scientists—Richard Dawkins and Daniel Dennett, for example—see any world view that is inconsistent with the Darwinian account of evolution as not only factually wrong but as dangerously delusive.

(2) I have reservations about the concept of a religion. An account of them can be found in my “Non Est Hick,” in Thomas D. Senior (ed.), *The Rationality of Belief and the Plurality of Faith* (Ithaca, NY: Cornell University Press, 1995), 216–41. The paper is reprinted in Peter van Inwagen, *God, Knowledge, and Mystery: Essays in Philosophical Theology* (Ithaca, NY: Cornell University Press, 1995), and in Kelly James Clark (ed.), *Readings in the Philosophy of Religion*, 2nd edn (Peterborough, Ont.: Broadview Press, 2008). I cannot discuss these reservations within the confines of this chapter. I record my conviction that if the argument of this chapter were rewritten so as to accommodate them, that argument would not be weakened—although it would be much longer.

(3) Brian Goodwin, *How the Leopard Changed its Spots: The Evolution of Complexity* (New York: Charles Scribner's Sons, 1994), pp. vii–ix.



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Belief in the Face of Controversy

Hilary Kornblith (Contributor Webpage)

DOI:10.1093/acprof:oso/9780199226078.003.0003

Abstract and Keywords

This chapter examines differences among disagreements in different domains. It notes that in some domains disagreements usually can be readily resolved, in others consensus often emerges over time. However, it is argued that the dynamics of consensus building in philosophy is quite dissimilar to what is found in mathematics and science. These considerations lead to the conclusion that philosophical beliefs held in the face of peer disagreement are typically not justified.

Keywords: disagreement, disagreement in philosophy, consensus, justification

Disagreement* is ubiquitous. One need not be especially opinionated, or especially well informed, to be aware that, whatever one's beliefs, there are a great many people who have beliefs that are contrary to one's own. As I write these words, the United States is about to head to the polls for another election, and many of the issues that voters will face are ones on which the electorate is deeply divided. The electorate is divided about moral and political issues, and it is divided, as well, about the likely effects of various policies; it is also divided about simple matters of historical fact. Anyone who has views on religious questions of any sort, including agnostics and atheists, is well aware that there are many others who see matters differently. Famously, there is deep division of opinion about evolution, especially in the United States, but, to a growing extent, in other industrialized countries as well. At times, it can seem that there is hardly any subject on which one might have opinions on which there are not many others who disagree. What is the epistemological significance of such disagreement? To what extent, if any, need one modify one's views in the face of disagreement?

Of course, not all disagreement presents even a *prima facie* epistemological challenge to one's own beliefs. I believe that the capital of Maine is Augusta. In fact, I know full well that the capital of Maine is Augusta. In order to make certain of this, I just checked it on a map. I also know that many people believe that the capital of Maine is Portland. This does not make me reconsider my opinion at all, nor should it. I have an explanation **(p.30)** of why it is that so many people are mistaken: people often assume that the capital of a state with which they have little experience is its largest city. More than that, having checked my belief against a reliable source, I also know that those who disagree with me would come to recognize their error, and change their belief, when confronted with the evidence I have. There is not some deep disagreement here about what constitutes good evidence for believing that a city is the capital of a given state. So, in this particular case, I have good reason to believe that I am in a superior epistemic position vis-à-vis the question of the capital of Maine, and the disagreement I have with others is epistemically unproblematic. I should go on believing as I do even in the face of this disagreement. In this case, at least, I know that I have better evidence on this question than do others who disagree with me.

There are other cases as well in which I may shrug off the disagreement of others, not because I have better evidence than they, but because I have better judgment. If, in trying to help a young child with his arithmetic homework, I explain that he has made a simple error of addition—five plus seven is twelve, I try to remind him, not thirteen—I should be completely unperturbed if he should insist that, no, he is right; five plus seven, he tells me, really is thirteen. My arithmetical skills are not, of course, infallible, but surely in this case it is perfectly reasonable for me to go on believing as I did before facing this disagreement. I know full well what the sum of five and seven is, and disagreement from an arithmetical neophyte should not cause me to reassess my views about arithmetic or about my mathematical abilities.

It is clear, however, that not all cases of disagreement are like either of these two cases, and so I cannot simply chalk up my disagreements with others to their inferior evidential situation or their inferior judgment. I certainly cannot insist that others' evidence or judgment must be inferior to my own simply in virtue of their disagreement with me. And it is for this reason that the fact of widespread disagreement raises troublesome epistemological questions. There is a very wide range of topics on which we knowingly disagree with others, and yet we seem to lack any reason to think that their evidence or their judgment is inferior to our own. What should we do in such cases? What is it reasonable to believe in the face of such disagreement?

Disagreements within philosophy constitute a particularly interesting case of this kind of disagreement. Consider the debate between internalists (**p.31**) and externalists about epistemic justification. I am a committed externalist. I have argued for this position at length and on numerous occasions. My arguments for the position are not merely a pose. I have not presented arguments for externalism merely to serve as a gadfly, provoking discussion. I sincerely believe that the arguments I have presented are good arguments, and I sincerely believe their conclusions. At the same time, I recognize, of course, that there are many philosophers who are equally committed internalists about justification, and that the arguments they offer are ones not only to which they are committed, but which they believe are good ones, and whose conclusions they sincerely believe.

It would be reassuring to believe that I have better evidence on this question than those who disagree with me, that I have thought about this issue longer than internalists, or that I am simply smarter than they are, my judgment superior to theirs. It would be reassuring to believe these things, but I do not believe them; they are all manifestly untrue. So, on the question of whether externalism or internalism is correct, I find that I have an opinion, but there are others who disagree with me who are, to adopt a useful term,¹ my epistemic peers: they are just as smart as I am, just as well informed, and have thought about the issue just as long as I have and just as carefully. That my epistemic peers disagree with me on this question is surely relevant evidence that I ought to take into account. It is indirect evidence on the question of internalism versus externalism, but it is important evidence nonetheless. And it surely seems that the proper way to respond to evidence of this sort is to suspend judgment, to suspend belief about the proper resolution of the debate between internalists and externalists.

As Richard Feldman has pointed out,² this is precisely what it seems we ought to do, and what we in fact typically do, in perceptual cases. Thus, to take his example, if you and I are each looking out of a particular window, in exactly the same direction, and I see someone standing in the middle of the quad and you see no one there at all, then we have a puzzle. If you and I have equally good vision and we are both in our right minds, then we should each be surprised by the other's judgment. "What do you mean," **(p.32)** I will say, "that you don't see anyone there?" "What do you mean," you will say, "that you really see someone out there in plain view?" Once we become convinced, however, that each of us is playing no joke, that we are looking in the same direction, and that we are, to all appearances, functioning normally, we should suspend belief about what there is in the middle of the quad. Someone here is making a bad mistake, and that person may well have some serious problem, and we will surely both agree that all of this is true. But it would be completely unreasonable for me to conclude, for this reason, that you are the one with the problem. I have no more reason to think that you are making a mistake than that I am, and, for that very reason, I should suspend judgment, as should you. When we give proper weight to the judgment of our epistemic peers in perceptual cases, it seems, suspension of belief is required.

As David Christensen has pointed out,³ the same seems to be true in mathematical cases. To take his example, suppose that you and I go out to a restaurant with a number of friends. After a large meal, the check comes and we agree to split the bill evenly. You and I are each quite good at mental arithmetic. I take a look at the bill and figure out what each person owes, and I put my share in the middle of the table. You look at the bill and figure out what each person owes and put your share in the middle of the table, and then we notice that we have put in different amounts. We are each well aware of the other's mathematical abilities, and we are each convinced of the other's honesty. At least one of us has made a mistake. It would surely be unreasonable for me to conclude that, since one of us has made a mistake, it must be you. The reasonable thing to do in this situation, and surely what most people in fact do in this situation, is suspend belief. We each go back to the bill and try to recompute the proper share.

The dispute about the bill at the restaurant and the dispute about the person in the quad are importantly different from the philosophical dispute about internalism and externalism. In the restaurant case, our disagreement will, in actual practice, quickly be resolved. One of us has made a simple arithmetical error. We are each quite good at mental arithmetic, and we will quickly determine who has erred when we recalculate. We suspend judgment when we first note our disagreement, but our judgment is **(p.33)** suspended only briefly. After quickly recalculating, we will figure out which of us was right about the bill, and we will comfortably agree about its proper division. In the perceptual case as well, we are often able to resolve our dispute in short order. We look more carefully; we call in a third person; we go out into the quad. The dispute is, often enough, only momentary, and our puzzlement is resolved. ("Oh, now I see," I say, "it was only a shadow.") While not all such disputes are resolved, of course, they leave few lingering doubts about perception or perceptual belief. The occasional disagreements that arise on such matters and that resist easy resolution do not leave us suspending belief about perceptual matters generally, nor should they.

But philosophical matters are importantly different. If you and I disagree about internalism and externalism, and we have each, as I have supposed, thought about this issue for many years, then a quick run-through of the arguments will not resolve the issue between us; this is, in no way, like the restaurant case. By the same token, we cannot profitably, as in the perceptual case, do anything analogous to looking more carefully (since we have already scrutinized the arguments with great care), call in a third party (since we know that there are many third parties on each side), or head out into the quad for a closer look. If we set aside our opinions about internalism and externalism as a result of our dispute, then we are likely to be suspending judgment on this issue for quite some time. When we disagree about our share of the bill at the restaurant, or about what is going on in the quad, we recognize that these are isolated disagreements against the background of extraordinarily broad agreement about arithmetic and perceptual matters generally, and, for this very reason, there is no threat of skepticism here about mathematics or the physical world. Even if our disagreement should prove to be unusually difficult to resolve, for whatever reason, it is an isolated disagreement, one that does not threaten to force us to suspend judgment very widely. But, if the right thing to do in the internalism and externalism case is to suspend judgment, then it seems that we will be suspending judgment on philosophical matters generally, and not just for a moment. And, since most of us have very deep philosophical commitments on a great many matters, this would involve substantial revisions in our corpus of beliefs, revisions most of us are not eager to undertake. A broad skepticism about philosophical matters threatens. **(p.34)**

I will assume, for the purposes of this chapter, that the proper resolution of the disagreements in the perceptual and mathematical cases is, in fact, as I have described: in the face of disagreement from epistemic peers, given no special reason to believe that either party is mistaken, one must, if one is to be epistemically rational, suspend belief. I believe the case for this has been well made by Feldman,⁴ Christensen,⁵ and Adam Elga,⁶ and in related work by Roger White.⁷ Supposing that this is correct about such mundane matters, however, how should we respond in the case of philosophical opinion? Are we to suspend judgment there as well? If so, does this force us to a broad skepticism about philosophical matters? And, if we are rationally forced to such a skepticism about philosophical matters, how broadly does this skepticism generalize? These are the questions I wish to address in this chapter.

1.

When I find that others disagree with me on a certain question, this gives me, *ceteris paribus*, reason to be less confident than I was that I am right. In the cases we have been looking at, disagreement from epistemic peers gives one reason to suspend judgment entirely on the disputed question. It is clear, moreover, that the mere possibility that someone might disagree with me does not have the same epistemic significance. Indeed, if it did, then, since there might always be people who disagree with one on any question at all, treating merely possible disagreement as on a par with actual disagreement would result in total skepticism. The worries generated by problems of disagreement, however, broad as they are, are not of this sort. So there seems to be an important asymmetry between actual disagreement and merely possible disagreement.

As Thomas Kelly urges,⁸ however, one should not overstate the difference here. "Suppose," as Kelly suggests, "that there would be considerable disagreement with respect to some issue, but all of the would-be dissenters have been put to death by an evil and intolerant tyrant." In such a case, **(p.35)** disagreements that do not in fact occur, disagreements that are not actual, are no less epistemically significant than actually occurring disagreements. "The significance of actual disagreement," Kelly concludes, "need be no more intellectually threatening than certain kinds of merely possible disagreement." How then should we determine the role that actual and merely possible disagreement should play in determining what to believe? Kelly argues that considerations of disagreement, in the end, drop out of account.

Whether we find the possibility of disagreement intellectually threatening, I suggest, will and should ultimately depend on our considered judgments about *how rational* the merely possible dissenters might be in so dissenting. And our assessment of whether rational dissent is possible with respect to some question (or our assessment of the extent to which such dissent might be rational) will depend in turn on our assessment of the strength of the evidence and arguments that might be put forward on behalf of such dissent. But if this is correct, then the extent to which merely possible dissent should be seen as intellectually threatening effectively reduces to questions about the strength of the reasons that might be put forward on behalf of such dissent. Now there might be cases in which we judge that the arguments and evidence that could be brought forth on behalf of a hypothetical dissent are truly formidable, and this might justifiably make us doubt our own beliefs. But in that case, the reasons that we have for skepticism are provided by the state of the evidence itself, and our own judgments about the probative force of that evidence. The role of disagreement, whether possible or actual, ultimately proves superfluous or inessential with respect to the case for such skepticism.⁹

Now Kelly clearly has in mind cases in which disagreements actually arise, and arguments are offered on behalf of competing views, or, alternatively, although no actual disagreement has arisen, arguments against a particular view are available. And, in these cases, Kelly urges that we simply look at the arguments for or against the views in question, and let the fact of disagreement, when there is one, drop out of consideration. Whether such a disagreement actually arises, or, instead, is merely possible, is of no epistemic weight at all, according to Kelly.

It is important to note, however, that in many cases where disagreements arise, arguments of the sort Kelly has in mind are not in play. Thus, remember Christensen's restaurant case, where you and I reach different conclusions about the fair division of the check. Neither of us offers reasons **(p.36)** for our conclusion about how the bill should be split, and yet we are each faced with a significant challenge to our respective beliefs. When I suggest that a fair division requires one payment and you suggest, instead, that it requires another, the problem is already set. Given that we each are aware of the fact that we are both highly reliable in figuring out such answers, the problem is set precisely by the fact of disagreement, and, at this point in the conversation between us, we cannot explain the epistemic problem we each face without pointing to that disagreement itself.¹⁰ So, at least in cases such as this, the role of the disagreement itself is anything but superfluous, contrary to Kelly. Perceptual cases are often much the same.

Kelly is right, of course, that my belief will not be threatened when you disagree with me if I know that your contrary opinion is unreasonable. This does not mean, however, that I must know the basis for your belief before it can threaten my own. As in the restaurant case, and in perceptual cases as well, I may know that your beliefs on this sort of matter are generally reliable, without knowing anything specific about the basis of your belief; I need not, in particular, know any of your reasons at all. In cases of this sort, the fact of disagreement, in light of my background knowledge of your reliability, is an ineliminable part of my reason for suspending belief.

What should we say, however, about the kinds of cases Kelly clearly has in mind, cases where each party to the disagreement has actually offered arguments for his or her view—at least when there is an actual disagreement—or, alternatively, when there is no actual disagreement, and yet arguments are available for each of two conflicting views? This is precisely the kind of case which arises in philosophy. In the internalism and externalism case, there are arguments that have been widely noted on both sides of the issue, and the participants to the debate, whatever position they may hold, are aware of these arguments. Similarly, as Kelly points out,¹¹ there are issues in philosophy where arguments are available on either side of an issue, and yet one side of the argument has no actually existing adherents. Kelly's view, in both sorts of case, is that the fact of disagreement, or the lack of disagreement, should drop out of rational consideration. We should consider the arguments, and let it go at that. **(p.37)**

Now there is no question that this accords well with standard philosophical practice. We simply do not tend to see arguments of the following sort in the philosophical journals.

Smith has offered an argument that p , and I must admit that the argument appears to be sound. But Smith has failed to take account of important evidence that p is false: I am one of Smith's epistemic peers, and although I grant that Smith's argument appears sound, I find that I am not convinced by it. As a matter of fact, I believe that not p . So there you have it: Smith believes that p ; I believe that not p . It is a tie. Given this evidence, we should just both suspend belief.

Nor does one see probabilistic variants of this kind of reasoning, in which the distribution of opinion within the field is offered as evidence for the degree of confidence that should be assigned to a disputed claim. Standard practice is to do as Kelly recommends: ignore the distribution of opinion and focus on the arguments. Kelly's suggestion that we focus on arguments in such cases is not meant as merely pragmatic advice. He is not suggesting that focusing on arguments in these cases will advance the cause of inquiry, or that it is more likely to help get one's work published in respectable journals, or that it will be likely to advance one's professional standing. His point, of course, is that the arguments are where the relevant evidence is to be found. That there is an actual disagreement, or that the preponderance of opinion falls on one side or other of the disagreement, is, on Kelly's view, simply irrelevant to the question at issue.

So let us imagine following Kelly's prescription. I have offered arguments for externalism, and, when I examine them, I find them thoroughly convincing. You have offered arguments for internalism, and, when I examine them, I do not find them convincing. Now I know that the situation here is perfectly symmetrical: you are convinced by your arguments, and you are not convinced by mine, but these facts, Kelly tells us, are ones that I should not take into consideration in evaluating what to believe; they are simply irrelevant here. I should focus on the arguments and forget about who believes what.

But why should I do this? I know that you have thought about this issue as long as I have; that you are as familiar with the arguments on both sides of this issue as I am; and I know that you are just as smart as I am. We **(p.38)** are epistemic peers. So, just as in the restaurant case, I have good reason to regard you as reliable, or, at a minimum, as reliable as I am. And, if this can serve as a reason for taking the fact that your opinion is different from mine as reason for doubt when I know nothing of your reasons, it is hard to see why it should be that it cannot equally serve as reason for doubt when I do know your reasons.

After all, once our reasons are on the table, the dispute between us comes down to our abilities to assess the cogency of complicated arguments.¹² And we agree that we are each quite good at that; we are epistemic peers. But we disagree about the cogency of the arguments at issue. So the disagreement about internalism and externalism has now been replaced by a different disagreement, a disagreement about the cogency of various arguments. How are we to resolve this dispute? On Kelly's view, we would have to give reasons here as well, and it is these reasons, presented in the form of arguments, on which the rational resolution of the dispute must turn.

Now no one, I think, will disagree with the suggestion that formulating such explicit arguments may be a useful thing, and it may serve to advance our inquiry. We should each do our best to formulate the arguments that we find most revealing, as well as the counterarguments that, we each believe, show the errors in our opponent's view. But these points are pragmatic ones, points about what we should do in order to advance inquiry. And Kelly's claim is not one about such pragmatic concerns, but one about what it is reasonable to believe right now in the face of disagreement.

So let us return to the disagreement you and I have about, as I have been supposing, internalism and externalism. You and I have opposing views, and we are each aware of the other's view; we are also each aware of the other's arguments. And the question at issue is not what we should do now to advance our inquiry; looking at arguments is, no doubt, a good **(p.39)** strategy to pursue for that purpose. The question at issue is what, at this very moment, each of us is justified in believing in the face of all of the evidence we have. Kelly's view is that the fact of our disagreement is not even relevant evidence here.¹³

In order to see how best to handle this sort of case, it will be useful to examine a different case in some detail. So let us look at what Kelly has to say about the Newcomb Problem.

2.

Robert Nozick presented the Newcomb Problem, a problem in decision theory, in 1969.¹⁴ There are two possible solutions to the Problem: the one-box solution and the two-box solution. The Problem arises because there is a highly intuitive argument for each of these incompatible solutions. As Kelly points out,¹⁵ when Nozick first presented the Problem, he remarked¹⁶ that opinion within the community of decision theorists seemed to be roughly evenly divided. Currently, however, Kelly notes that “the by-now over three decades of sustained debate . . . has resulted in a significant shift in the original distribution of opinion in favor of Two-Boxing.”¹⁷ How should the distribution of opinion on this question affect our judgment about its proper resolution? Kelly's answer, as we have seen, is that we must look to the arguments on either side, and ignore the distribution of opinion.

Why should we ignore the distribution of opinion on this case? Kelly asks us to imagine a student studying the Problem at the time it was introduced; in talking to others, she discovers that opinion is evenly divided on the question. There are other possible worlds, however, in which everyone **(p.40)** she meets is a One-Boxer. Is it rational to believe different things in these two worlds? Not according to Kelly.

Should she take a different view about Newcomb's Problem in the other, unanimous world than she does in our fragmented and divided world? *Despite* the fact that she has access to exactly the same arguments in both worlds? This seems extremely dubious—after all, can't the student in the unanimous possible world simply look over at our own fragmented world, and realize that *here* she has epistemic peers who extol Two-Boxing?¹⁸

As Kelly notes, “whether there is any actual disagreement with respect to some question as opposed to merely possible disagreement might, in a particular case, be an extremely contingent and fragile matter.”¹⁹ But surely this leads, he argues, to the conclusion that the distribution of opinion should carry no weight for us.

Now I think it is significant that Kelly has chosen an example from decision theory. Areas of formal investigation within philosophy, such as logic, probability, and decision theory, are ones in which an extraordinary amount of progress has been made. There can be little doubt that there is as much reason for confidence in some of the results in these areas as there is in the empirical sciences. A look at the history of these fields, moreover, shows some interesting trends.²⁰ As in mathematics, intuitions about results *and about arguments* in these areas are extremely unreliable at the early stages of investigation. It is only as theory advances that we come to be able to understand the fundamental concepts at issue in these fields, the kinds of intuitions that are genuinely insightful, and the kinds of arguments that are genuinely sound. To take a single example, Stephen Stich notes²¹ that, as late as the end of the nineteenth century, one could find logic textbooks that defended the gambler's fallacy as a legitimate statistical inference. Diversity of opinion is a commonplace in these areas early on in the development of theory; as theory advances, the opinions of experts in the field tend to converge. Even among experts, of course, convergence of opinion is no guarantee of truth, but one would have to be a radical skeptic about mathematics, logic, probability, and decision theory **(p.41)** to think that convergence of opinion is not, at this point in the history of these fields, evidence of truth. And at this point in the history of these fields, I think it is fair to say, radical skepticism is no longer a rational option.

Thus, as I see it, when we see that, after three decades of work on the Newcomb Problem, there is an emerging consensus on a solution, this is itself, given the history of the field, strong evidence that the consensus is correct. What should we make of Kelly's suggestion that we may always look over at other worlds in which the consensus is different, or in which there is no consensus? And what should we make of his suggestion that consensus may be "an extremely contingent and fragile matter"? I believe that these suggestions give a misleading view of the dynamics of consensus within formal fields in philosophy.

Let us consider, by way of contrast, views about the aesthetics of clothing. If I were to head off to buy a suit, I would find that there is some consensus in the fashion world about the proper width and shape of lapels. The consensus today is quite different from what it was five or ten years ago, and different still from what it was in the 1930s. Someone might argue that consensus among fashion experts is evidence of truth: the lapels currently approved of are ones, we should believe, that are aesthetically correct; the fashion world is converging on the timelessly right answer here. Whatever one thinks of the kind of realism about aesthetic value that such a view presupposes, it is surely unreasonable to think that consensus here is any evidence of truth at all. The history and dynamics of consensus in the fashion world are strikingly different from what one sees in mathematics. In mathematics we see periods of disagreement followed, after intensive study, by a growing consensus; in the fashion world, we see a large consensus that simply changes from year to year. The second of these two patterns does not support any kind of confidence that this year's consensus is more likely to get it right than last year's consensus, or the consensus of the year before, or the year before that.

The consensus we find in the fashion world is one that we might accurately describe, to use Kelly's words, as "an extremely contingent and fragile matter." Although lapels of a certain width are currently in fashion, they might just as easily have been seen as evidence of bad taste. We need change little about the underlying dynamics of opinion about fashion if we are to imagine a world in which much broader or narrower lapels are approved of. **(p.42)**

But now contrast this kind of case with cases of emerging consensus in formal areas of philosophy. Thus, consider the way in which a consensus was achieved following the publication of Gödel's result on the incompleteness of arithmetic. We can easily imagine worlds in which Gödel did not discover his proof, and in which no proof of incompleteness was discovered until much later, or at all. We can certainly imagine worlds in which someone else discovered the famous proof, or in which it was discovered earlier. Some of these worlds are not very different from our own. What would, however, be very much different from the world we inhabit is one in which Gödel discovers his proof, the very proof that he presented in the actual world; he presents it to the mathematical community, and a consensus emerges that it is mistaken. Such a world would be one that is very different from our own. It would be a world in which the abilities of mathematicians are utterly different from those in the actual world, and, as a result, one whose history would look entirely different from the history of mathematics as it has actually unfolded. The fact that we can imagine a world in which the people who are called mathematicians are all incompetent, and that they reach consensus in much the way that consensus about fashion is reached in the actual world, says nothing whatsoever about how we should allow consensus among actual mathematicians to affect our opinions.

Consensus in the mathematical community, except in rare cases, is not a fragile thing at all. When there is consensus in this community, the consensus is typically very robust: one could not easily change the consensus, or undermine it, without making very substantial changes in the underlying facts (for example, by changing the facts about what arguments are available on either side of an issue, or by changing the facts about the basic competence of the members of the community). The same is true within the scientific community. The fashion community, of course, is a different matter entirely.

So what should our student do when she examines the arguments for each side in the case of the Newcomb Problem? If she wants to understand anything about the Problem at all, she needs to examine the arguments on either side with tremendous care and try to figure out which of the two sides seems to have the better argument. And, if she wants to have any chance at all of advancing the community's understanding of the issues, she will, again, need to examine the arguments. But, if the question is, instead, **(p.43)** what she is justified in believing when she knows that the community of experts is deeply divided on the issue, or, alternatively, what she is justified in believing when she knows that the community is unanimous in favor of one side rather than another, then the answer is that she should go with the community opinion.²² And this is not simply true of students, who are not yet peers with the experts. It is true of the experts themselves as well.

Consider, once again, Christensen's restaurant case. In the simple version of the case, already discussed, you and I come to different opinions about how the bill is to be fairly divided. But Christensen also presents another restaurant case, in which I am dining with seventeen other people, each of whom is known to be as reliable at arithmetic as I am, and each of whom has come to the same conclusion about the fair division of the bill, a conclusion different from my own. Clearly in this case, as Christensen argues, I should believe that it is I who have made the mistake. Things are no different if we move from dividing the bill at a restaurant to solving a problem in decision theory. If my views on a problem in decision theory are entirely at odds with the experts in the field, then, even if I am myself such an expert—well known to hold a minority position on this issue—I would not be justified in continuing to hold the belief in the face of such opposition.²³

This is not, in many cases, how experts behave. Experts will often go on holding a view even when they know that they are in the minority. If we restrict ourselves to formal areas like mathematics, logic, probability, and decision theory, however, or to the empirical sciences, then the long-run prospects for such dissenters is not very good. These are all areas in which there is a well-established track record of reliable results issuing from the community, and, although dissenters are sometimes proven right, these communities are sufficiently reliable, and have been so for long enough, that one would always be ill advised to bet on the dissenter in the face of an overwhelming majority opinion. My point is simply that this is something that dissenters should realize as well. If the question at issue is thus whether one is justified in siding with the dissenter in the face of an overwhelming majority, the answer is that one is not, and this is true not only of bystanders **(p.44)** who seek to inform their opinions by looking to the experts; it is true of the dissenter him- or herself.

Kelly's suggestion then that we should look only to the arguments directly bearing on a disputed question, and simply ignore the distribution of opinion, cannot be sustained.

3.

So let us return to the internalism/externalism dispute and the question, more generally, of what we should make of disagreement within philosophy outside of the most formal areas. I have been arguing that the distribution of opinion is, contrary to Kelly, importantly relevant to what we should think about disputed matters. In the case of the internalism/externalism dispute, I will suppose, what I believe to be in fact true, that there is not an overwhelming majority of opinion among the experts within epistemology as to which of these approaches is correct. If opinion is not evenly divided on this issue, there is certainly not an overwhelming majority of expert opinion on one side rather than the other. What is it reasonable then to believe about this issue, given the current state of play? I believe that we are justified in withholding opinion on this issue, and that one would not be justified in believing either that internalism is correct or that externalism is correct.

Suppose, first, that one thinks that the philosophical community is, in relevant respects, like the mathematical or scientific communities. Suppose, that is, that one believes that individual philosophers are, by and large, quite reliable in their opinions in matters that touch on their areas of expertise, and that the community, overall, has shown a long history of steady progress on the issues it addresses. If one were to believe this (and be justified in believing it), then, as I have argued, the distribution of opinion is straightforwardly relevant to what one should believe, and, when opinion is closely divided, one ought to suspend judgment. So, in this case, suspension of judgment would be mandatory.

But surely it is not reasonable to believe that the philosophical community is like the mathematical or scientific communities in relevant respects. We do not have a long history of steady progress on issues, and, as a result, the case for deferring to community opinion is thereby weakened. But **(p.45)** this hardly strengthens the hand of those who would form an opinion one way or the other on matters such as the debate between internalists and externalists. When the community is composed of individuals each of whom is reasonably believed to be reliable, we must bow to majority opinion. But, if we do not have reason to believe that the community is composed of individuals who are reliable, if the history of the field gives us no reason for confidence in the judgment of individual practitioners, then this, by itself, gives us reason to suspend judgment on questions that confront the field. If the history of the field shows no track record of success in addressing the issues it confronts, the only conclusion we can reasonably reach is that there is no basis for opinion here on anyone's part at all. It certainly does not give one free rein to believe whatever one pleases.

The sad truth, it seems, is that the history of philosophy does not look remotely like the history of science or mathematics when it comes to the dynamics of consensus among its most esteemed practitioners, and this has a striking bearing on the question of its epistemic credentials. One might try to carve out a recent piece of this history, and some particular subject matter, where one believes that real progress is being made and that we are finally getting to the truth on some important issue. I must confess that, in my more optimistic moments, I find such a view tempting. But, if we are to take any such view seriously, and subject it to real scrutiny, we would surely find that this view of the particular question at issue is itself a subject of real controversy among acknowledged experts in the field, and so it too must be seen, on careful consideration, as an issue on which we ought to suspend judgment. The field of philosophy in general, outside of the more formal areas that are most akin to mathematics, simply does not have anything like the epistemic standing of the empirical sciences. So, much as we all find ourselves forming beliefs about disputed philosophical questions when we immerse ourselves in the arguments, we should acknowledge in quiet moments of reflection that these views we form are ones that are not epistemically justified. It would be as presumptuous to claim that we are justified here as it would in Christensen's restaurant case when we find that our mathematically reliable dinner companion has reached a different conclusion about the division of the check.

It is not my view that the distribution of opinion is all the evidence one needs to determine what one should believe on any given question; such a view is completely untenable. It must be acknowledged, however, **(p.46)** that the distribution of opinion among acknowledged experts must carry a tremendous amount of epistemic weight, and anyone who would reject a consensus among experts, or adopt a specific view in the face of deep division among experts, faces a very large hurdle. One must show that, in such a case, there is independent reason to believe that a substantial group of experts have gone wrong on the disputed matter.²⁴ To take a purely imaginary case, one might be able to show that a large portion of some field of experts had been kidnapped and forced to express certain views at gunpoint, and, in such a case, one would be fully justified in ignoring the expressed consensus in the field. Here one's evidence of the kidnapping would be entirely independent of the disputed question within the field, and, once the expressed opinions of the kidnapped experts were thereby discounted, one might well be in a position to have confidence in one's own judgment. But here, too, one would need to attend to the opinions of the remaining experts, if any. And I think it is perfectly clear that the high hurdle that bucking the consensus of expert opinion presents us with, even when we are among the experts ourselves, is one we are rarely, if ever, in a position to meet. I have no reason to believe, for example, that internalist epistemologists have been defending their views under threat from subversive kidnappers.²⁵

I am thus forced to conclude, very reluctantly, that the opinions I hold on most philosophical matters—and I have a great many of them—are not epistemically justified. Given the current state of the field, no one's opinions on these matters, it now seems to me, are epistemically justified.

More than this, this conclusion seems to generalize quite broadly. There are, for example, a great many moral and political issues, issues about which I have, in some cases, rather strong opinions, that are subjects of **(p.47)** dispute among very intelligent, thoughtful, and well-informed individuals. In some of these cases, we can explain away the disagreement of otherwise reasonable people by way of moves less desperate than the suggestion that our opponents have expressed their views only under threat by kidnappers. We do, after all, sometimes have good, independent reason to believe that someone who disagrees with us on a particular question is biased on that very question. We should, however, be wary of making this move too easily. The requirement of independence means that we cannot reject the opinions of our opponents as the product of bias merely on the basis of the view that their opinions are false—since they disagree with our own—and therefore must inevitably be due to some sort of biasing factor. What this means, in the end, is that this sort of maneuver will help us only rarely in dispensing with the challenge of disagreement from peers. And what follows from this, of course, is that a broad skepticism threatens.

This is not the sort of total skepticism of the Cartesian demon, but it would certainly force a radical revision in our body of beliefs.

4.

One might try to resist this conclusion, not, as Kelly does, by denying the relevance of the opinions of others, but, instead, by denying that giving others' opinions their due weight in the interesting cases of disagreement—in philosophy, for example, and on moral and political questions—forces such a broad withholding of opinion. On such a view, idealized cases like Christensen's restaurant example do genuinely show that one's own opinions should be given no more deference than the opinions of one's epistemic peers, but features of the more complex cases—somehow—allow one to avoid the unwanted and widespread change in our bodies of belief to which the general principle seems to lead. This, to be sure, is a consummation devoutly to be wished. It is, moreover, precisely the position defended by Adam Elga.²⁶

As Elga points out, when you and I disagree about the fair division of the check at the restaurant, our disagreement is highly isolated. We do not disagree about arithmetical questions in general. Indeed, I count you as **(p.48)** my arithmetical peer precisely because we agree so broadly on arithmetical questions. More than that, it is not just that we agree. You and I might agree on some questions, knowing full well that we are rank amateurs in an area where others are far better informed. But this is not the case in the restaurant example. We not only agree. We have good reason to believe that we are as well informed as just about anyone when it comes to these sorts of questions. We are fully justified in believing that we are both highly reliable in forming judgments about simple arithmetical issues. We simply disagree about an isolated question within arithmetic when we are each able to assess the other's track record on arithmetical questions in general in a perfectly straightforward way.

This is not so, as Elga points out, in many other areas. Thus, Elga imagines two friends, Ann and Beth, who disagree about abortion. Their disagreement about abortion, if we are to make the case realistic, as Elga does, is not an isolated disagreement. Instead, there is a very wide range of related moral questions about which Ann and Beth disagree. But this is no minor complication, as Elga argues, for now the very basis each has for regarding the other as generally reliable on the kinds of questions at issue has been undermined. I may regard you as reliable about arithmetic even when we disagree about how to divide the check at the restaurant because we agree, it may be supposed, about every other arithmetical question we have ever jointly considered. But Ann and Beth cannot regard each other as reliable about moral matters generally precisely because their disagreement is so broad. We will regard others as suitably reliable only when they agree with us very broadly. It is for this very reason, Elga argues, that we may eat our cake and have it too: we may acknowledge that the opinions of our epistemic peers count just as heavily as our own antecedent opinions, and yet deny that in areas of moral and political dispute, for example, we must simply withhold opinion. Thus, Elga concludes, “with respect to many controversial issues, the associates who one counts as peers tend to have views that are similar to one's own. That is why—contrary to initial impressions—the equal weight view does not require one to suspend judgment on everything controversial.”²⁷

One might try, in response to this move, to think of the dispute between Ann and Beth not as one about abortion alone, but as one about a cluster **(p.49)** of related issues that includes the abortion question. Ann and Beth, one might argue, regard themselves as reliable on issues outside this cluster, and it is for this reason that they are troubled by their disagreement about the cluster itself. They should regard each other as reliable on the basis of their agreement on issues outside the cluster surrounding abortion, and this then forms the basis for taking the disagreement over the cluster of issues surrounding abortion as seriously as they do. Once one views the disagreement in this way, it seems, the parallel with the restaurant case is restored.

But Elga denies this.

Contrary to what the objection supposes, Ann does *not* consider Beth a peer about that cluster [of issues related to the abortion question]. In other words, setting aside her reasoning about the issues in the cluster, and setting aside Beth's opinions about those issues, Ann does not think Beth would be as likely as her to get things right. That is because there is no fact of the matter about Ann's opinion of Beth, once so many of Ann's considerations have been set aside . . . Once so much has been set aside, there is no determinate fact about what opinion of Beth remains.

Thus, according to Elga, once again, we see that the basis for seeing our opponents on moral and political matters as generally reliable, when our disagreements are very broad,²⁸ is undermined. And, once our judgment of the reliability of our opponents has been thus undermined, he argues, we are entitled to go on believing as we did before the disagreement was encountered.

I do not believe, however, that Elga is right about this. First, it seems to me that Elga's attempt to undermine Beth's and Ann's judgments about each other's reliability here is unsuccessful. It is true that, in the imagined case, there is a large cluster of moral issues about which Ann and Beth disagree. But we need not exaggerate their disagreement. Thus, for example, Elga describes Ann and Beth as being "at opposite ends of the political spectrum." So it would not be unfair to imagine Ann as, say, someone who characterizes herself as a typical **(p.50)** American pro-choice Democrat, and Beth as someone who characterizes herself as a typical American pro-life Republican. They disagree, to be sure, on a wide range of moral and political issues. But, although Beth and Ann disagree about a great deal, their disagreement is not at all like their disagreement with, for example, Zena, a homicidal sociopath. Zena does not just disagree with Ann or Beth about the cluster of moral issues surrounding the abortion question. She disagrees with them about virtually every moral question one might care to raise.

Now I think it is safe to say that neither Ann nor Beth will be much troubled by Zena's disagreement with them, nor should they. And the reason why they should not is precisely that, when we subtract the moral issues on which Ann or Beth disagree with Zena from the totality of moral issues, there is virtually nothing left at all on which they might base a judgment that Zena is, but for their little disagreement, generally reliable about moral issues. Here we may reasonably say precisely what Elga says about the disagreement between Beth and Ann: once we set aside the issues on which Beth and Ann disagree with Zena, there simply is no basis for forming an opinion about Zena's reliability on moral issues. When we set aside the areas of disagreement, after all, there are no moral issues that are left.

But this, it needs to be emphasized, is not at all the case with the disagreement between Beth and Ann. They disagree profoundly about an important range of moral issues, but neither regards the other as the moral equivalent of Zena. Indeed, it is because they do not regard each other as Zena's moral equivalent that they are so engaged, and so disturbed, by each other's opinions. Beth and Ann regard each other as basically decent, caring, thoughtful individuals whose opinions on a very wide range of moral matters, outside the sphere of issues most closely related to abortion, are trustworthy and insightful. While they rightfully dismiss Zena's moral views out of hand, they are in respectful agreement about a very wide range of moral issues. And it is on this basis that they regard each other as moral epistemic peers, something they simply cannot do with Zena. Given that they do, justifiably, regard each other as moral epistemic peers, their grounds for withholding belief on the cluster of issues on which they disagree is thereby restored.

Note that the same is true when we consider disagreement about philosophical questions. I disagree with others about the proper resolution of the internalism/externalism debate in epistemology. But this is not **(p.51)** like the restaurant case, where there is a disagreement about a single claim against the background of complete agreement about all other issues on the same general subject. Rather, once again, we have a case like the abortion issue. Those whom I disagree with about internalism and externalism are philosophers with whom I disagree about a wide range of related issues. Does this then mean that I am no longer in a position to see these philosophers as my epistemic peers, as Elga suggests is the case with Beth and Ann on the abortion issue? Not at all. Even these disagreements, broad as they are, take place against a background of very broad agreement, agreement about the important issues in epistemology, about which positions are worth taking seriously, about what counts for and against various views, and so on. In short, I view internalists in epistemology in much the same way as Ann views Beth; I do not view internalists in the way that Ann views Zena. What this means, of course, is that there is ample room to view such philosophers as my epistemic peers, which is, in fact, precisely how I view them.

Disagreement on a wide range of related issues, as in the abortion debate or the internalism/externalism controversy, is not automatically a bar to reasonably viewing one's opponents as one's epistemic peers. And it is for this very reason that, on many matters of great controversy, the only rational thing to do is to suspend belief.

5.

It is worth thinking about these controversial matters from a somewhat broader perspective. Consensus or near-consensus on formal issues within philosophy is so weighty epistemically, as I have suggested, because there is a history within these formal areas of undeniable progress. Against that background, consensus among the experts is a formidable matter. By the same token, disagreement among experts in those fields must also carry great weight. The same is true, of course, in the empirical sciences. When we look at the track record of less formal matters in philosophy, however, or on matters of public controversy having to do with moral and political issues, the case for a well-established track record of progress is, to be sure, harder to make out. Without that background of long-standing progress, we must look at individual investigators quite differently. In the sciences, **(p.52)** we have good reason to believe that individual experts are each highly reliable overall, and they should surely regard one another in this way. Their considered opinions should thereby be tempered when there is disagreement in the field. In philosophy, however, and the other areas of controversy we have been considering here, there is no such history of long-standing progress, and for that very reason we should not consider the experts in the field—including, of course, ourselves—to be highly reliable. The history of the field simply does not give any ground for credence in such a view. But it is then, for that very reason, that we must, in the end, withhold opinion on the issues under consideration.

I do not mean to suggest that we should stop thinking about these issues, or that thinking about them, and trying to work out tenable views, is not intellectually respectable. This is not my view at all. And, given the nature of human belief, I very much doubt that philosophers will stop forming views about the subjects they think about for so long and with such care. When we stand back, however, and reflect on our practice and on the beliefs that that practice generates, it seems to me that the history of our field makes epistemic modesty the only rational position available. We may hope that in trying to work out the views we are most sympathetic to that, in the long run, we may somehow contribute to an approach to philosophy that will look more progressive than any we have thus far seen. But, at the present time, we should all recognize that this is merely a hope, and that rational belief must be tempered by the facts about our current situation.

Notes:

* I have received helpful comments from the audience at a presentation of this chapter at the Free University, Amsterdam, and at MIT. I also wish to thank David Christensen, Adam Elga, and Tom Kelly for comments on a draft of the chapter.

(1) This term was introduced by Gary Gutting in his *Religious Belief and Religious Skepticism* (Notre Dame, IN: University of Notre Dame Press, 1982).

(2) Richard Feldman, "Epistemological Puzzles about Disagreement," in Stephen Hetherington (ed.), *Epistemology Futures* (Oxford: Oxford University Press, 2006), 223.

(3) David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review*, 116/2 (2007), 187-217.

(4) Feldman, "Epistemological Puzzles about Disagreement."

(5) Christensen, "Epistemology of Disagreement."

(6) Adam Elga, "Reflection and Disagreement," *Noûs*, 41/3 (2007), 478-502.

(7) Roger White, "Epistemic Permissiveness," *Philosophical Perspectives*, 19 (2005), 445-9.

(8) Thomas Kelly, "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology*, 1 (2005), 181.

(9) Kelly, "The Epistemic Significance of Disagreement," 181-2.

(10) There is, of course, an argument that each of us can give that explains our epistemic predicament, but any such argument makes essential reference to the fact of disagreement. These are not, of course, the kinds of arguments Kelly has in mind.

(11) Kelly, "The Epistemic Significance of Disagreement," 184-5.

(12) Actually, I do not think this is entirely correct, but I will assume it for the sake of argument in the text, since it is most favorable to Kelly's position. In fact, I believe, the reasons we are able to articulate often fail fully to do justice to the reasons for which we believe. It is these latter reasons, rather than the ones we can verbalize, that determine the epistemic status of our beliefs. I have discussed this further in "Distrusting Reason," *Midwest Studies in Philosophy*, 23 (1999), 181-96, and *Knowledge and its Place in Nature* (Oxford: Oxford University Press, 2002), ch. 4.


(13) Kelly does, in addition, present an argument that, even if we were to count the fact of disagreement as relevant evidence, this would not rationally force us to suspend belief in the cases described ("The Epistemic Significance of Disagreement," 185-90). Nevertheless, his view is, as the quoted passage shows, that facts about disagreement are epistemically irrelevant to the issue under dispute.

(14) Robert Nozick, "Newcomb's Problem and Two Principles of Rational Choice," reprinted in Nozick, *Socratic Puzzles* (Cambridge, MA: Harvard University Press, 1997), 44-73.

- (15) Kelly, "The Epistemic Significance of Disagreement," 182 n. 16.
- (16) Nozick, "Newcomb's Problem and Two Principles of Rational Choice," 48.
- (17) Kelly, "The Epistemic Significance of Disagreement," 182 n. 16.
- (18) Kelly, "The Epistemic Significance of Disagreement," 183.
- (19) Kelly, "The Epistemic Significance of Disagreement," 181.
- (20) See, e.g., Ian Hacking, *The Emergence of Probability* (Cambridge: Cambridge University Press, 1975).
- (21) Stephen Stich, *The Fragmentation of Reason* (Cambridge, MA:MIT Press, 1990), 83.
- (22) I am assuming here that our student has not come up with a new argument, for in such a case our student is familiar with an argument that the other side has not yet considered. There are interesting questions about the extent to which one should be confident that one is right even in such cases, but I will not consider them here.
- (23) Again, as in the previous note, I am assuming that I have not discovered some new argument in favor of my position that is as yet unknown to the other members of the community.
- (24) Christensen's discussion of this point about independence is particularly illuminating.
- (25) This bears on a question addressed by David Lewis in "Academic Appointments: Why Ignore the Advantage of being Right?," reprinted in his *Papers in Ethics and Social Philosophy* (Cambridge: Cambridge University Press, 2000), 187-200. Lewis notes the common practice, not only in philosophy departments, of bracketing the truth of a candidate's views in making judgments about whether the candidate should be offered a job. Why should we do this? Lewis's answer is as follows: "We ignore the advantage of being right because we comply with a tacit treaty to do so. It is reasonable for all of us to think that this treaty, and therefore our present compliance that sustains it, serves the advancement of knowledge" (p. 197). While I do not doubt that considerations of this sort may play a role, I suspect that this is not the major factor in explaining the practice, and that it is not the most important factor in explaining the legitimacy of the practice. Rather, in my view, it is a somewhat inchoate awareness of the weak epistemic standing of our opinions on disputed questions that both motivates and justifies the practice.
- (26) Elga, "Reflection and Disagreement."

(27) Elga, "Reflection and Disagreement," 494.

(28) This is an important qualification, as Elga notes. On Elga's view, disagreement on these matters with those who see things largely in the way we ourselves do can serve as an important check on our own opinions. As Elga notes, this, by itself, if taken to heart, would force important revisions in our bodies of belief. It would thus be wrong to paint Elga as some sort of quietist. He is merely trying to resist the very sweeping withholding of belief on controversial matters that I am defending here.



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Persistent Disagreement

Catherine Z. Elgin

DOI:10.1093/acprof:oso/9780199226078.003.0004

Abstract and Keywords

This chapter responds to arguments for the conclusion that participants in persistent peer disagreement ought to suspend judgment about the disputed proposition by noting that 'ought implies can' and that belief (and suspension of judgment) are typically not under the relevant kind of voluntary control. It is argued that issues about disagreement are better seen as being about acceptance rather than belief, and that continuing to accept propositions in the face of disagreement can have sufficient value to make it rational, and thus that peers can rationally accept conflicting propositions.

Keywords: disagreement, belief, acceptance, rationality, ought implies can, voluntary control

1. The Problem

Some disagreements are epistemologically benign. One party is wrong and it is easy to determine which one. Sue, dimly recalling Longfellow's poem, believes that the American Revolution began on April 18, 1775; Sam, relying on his study of history rather than poetry, believes that it began on April 19, 1775. Dan believes that his flute is in tune; Dora, who has perfect pitch, believes that it is not. Maeve believes that, if smoking causes cancer, then she will not get cancer if she does not smoke; Mark, realizing that 'if p then q ' does not entail 'if not- p then not- q ', disagrees. Such differences of opinion are unthreatening. Sam, Dora, and Mark have excellent grounds for their beliefs, and excellent grounds for thinking that their opponents are mistaken. They should retain their beliefs, and be unmoved by their opponent's opinions. Disagreement *per se* does not jeopardize epistemic standing.

More problematic are cases in which opponents are, and consider themselves to be, epistemic peers. Then they have the same evidence, reasoning abilities, training, and background assumptions. If Fred and George are, and realize that they are, equally good at spelling, then Fred's firm belief that the proper spelling is 'ignomineous' when George believes it is spelled 'ignominious' should give them both pause. They should probably suspend judgment and consult a dictionary.

Although inconvenient, such short-term suspensions of judgment are relatively easily handled. The serious difficulty comes with persistent disagreement, where no easy or obvious resolution is available. The **(p.54)** evidence is equivocal. The evidence class contains misleaders, but there is no consensus about which the misleaders are. Should opponents suspend judgment in these cases too? Suppose two paleontologists, Jack and Jill, are epistemic peers who disagree about the fate of the Neanderthals. Jack believes that Neanderthals were an evolutionary dead end. Unable to compete, they simply died out. Jill believes that Neanderthals evolved into later hominids whose descendants are alive today. Because the issue is complex and the evidence is equivocal, they come to different conclusions about it. What should they (and we) make of their disagreement? In particular, should the fact that an epistemic peer disagrees with Jack have any effect on the epistemological status of his belief? Should Jack's knowledge of that fact have any effect?

Opinions diverge. Some philosophers, such as Richard Feldman, Hilary Kornblith, and David Christensen, contend that the existence of peer disagreement undermines one's grounds for belief.¹ If someone with the same evidence, training, background knowledge and reasoning abilities came to the opposite conclusion from Jack's, that is evidence that Jack's grounds are inadequate. Such philosophers think that epistemic agents should moderate their views in light of the disagreement. Others, such as Thomas Kelly and Richard Foley, maintain that it is reasonable for a thinker to retain his opinion in the face of disagreement with an epistemic peer.² They think that epistemic agents should be resolute. Both positions have unwelcome consequences.

2. Resoluteness

Advocates of resoluteness maintain that Jack should hold fast to his belief. To do otherwise would be spineless. Since Jack believes that his reasons are sufficient to support his conclusion, he thinks that Jill is wrong about **(p.55)** the Neanderthals. This is compatible with her being, and his recognizing her as, his epistemic peer. Everyone makes mistakes. So, although she is generally as good a paleontologist as he is, their disagreement is reason enough for him to conclude that in this case she is in error. Even though he cannot point to any flaw in her reasoning, Kelly maintains, Jack should take the disagreement to show that there must be a flaw. The mere fact that they disagree convicts Jill in Jack's eyes. Since they have the same evidence, her error must lie in how she reasons about the evidence. From Jack's point of view, the disagreement demonstrates that Jill is, in this case, not just wrong but irrational.

The situation is symmetrical. Jill, on Kelly's view, should be equally resolute. From Jill's perspective, Jack is being irrational. There currently is no tie-breaker. Such symmetry is disconcerting. It is evidently impossible to determine which of them is rational without determining which of them is right. So the fact that a belief is rational cannot function as a reason to think it is right. The stance of each party to such a disagreement is precarious. It is hard to be confident that one's belief is well founded while conceding that one's epistemic peer reasonably considers it irrational. Moreover, resoluteness seems to deprive epistemic agents of resources for correcting their mistakes. If Jack can so easily dismiss Jill's opinion, her disagreement gives him no reason to re-examine his own position, to seek further evidence, or to develop better methods of assessing the evidence.

Jack and Jill are experts in paleontology. So their disagreement occurs in a context in which they have what is, and what they recognize as, good evidence, and in which they have and recognize that they have, the capacity to reason responsibly about that evidence. But peer disagreements can occur at any level of expertise. Bill and Beth are epistemic peers who are woefully naive about economics. Bill believes that Liberia's dependence on US currency weakens the dollar. Beth believes that it does not. They appeal to the same sparse and dubiously relevant considerations to justify their conclusions. Once their disagreement focuses attention on the inadequacy of the reasons they can offer to support their positions, it becomes clear that neither has a clue whether Liberia has any effect on the value of the US dollar. Plainly, they should suspend belief. Evidently, a threshold of competence has to be reached before resoluteness is even remotely reasonable. Only if epistemic peers are cognitively competent with regard **(p.56)** to the topic under dispute is it plausible that they should retain their belief in the face of disagreement.

3. Moderation

Feldman maintains that epistemic peers who disagree should suspend belief. Christensen maintains that each of them should moderate his degree of belief, although not perhaps to the point of suspending belief entirely. For reasons of simplicity, I will discuss only Feldman's position, but my points extend in obvious ways to Christensen's.

In cognitive contexts, Feldman notes, it is always open to an epistemic agent to suspend belief.³ She should do so whenever she recognizes that her grounds are inadequate. Jill's epistemic peer disagrees with her about the fate of the Neanderthals. She can find no fault with his reasoning. This, Feldman maintains, provides Jill with evidence that her grounds are inadequate. So she should suspend belief. The epistemic situation is symmetrical. Jack should suspend belief too. Symmetry, an unattractive feature of resolute theories, is unproblematic for advocates of moderation, for it leads to a convergence of opinion.

Although suspending belief in such cases may seem reasonable, it pushes in the direction of skepticism. Wherever issues are complicated, if there are epistemic peers, they are apt to disagree. Thus there are apt to be vast areas of inquiry where belief is to be suspended. Moreover, whether one happens to have an epistemic peer seems to be utterly contingent. This suggests that we should consider possible peer disagreement as well. Should Harry suspend belief because he recognizes that if he had any epistemic peers, some of them would disagree?

One might think not. We may imagine that thinkers with no epistemic peers are geniuses—people like Einstein or Darwin. The fact that they have no peers is reason to believe that they at least need not suspend belief. Their reasons and reasoning powers are so strong that no one with those reasons and reasoning powers could disagree. But in the recent debates about disagreement ‘epistemic peer’ is defined quite narrowly. It requires having *the same* evidence and reasoning abilities. So it is not surprising if an **(p.57)** ordinary person lacks epistemic peers with respect to a particular, mundane issue. If Jen and Jon have even slightly different relevant reasoning abilities or evidence pertaining to the causes of the Civil War, they are not epistemic peers with respect to the subject. Given the vicissitudes of education and abilities, and the idiosyncrasies of evidence-gathering, ordinary epistemic agents are apt to have few epistemic peers. But, if the only reason that Harry does not face peer disagreement about a particular issue is that he happens to have no epistemic peers with respect to that issue, then the absence of disagreement is fortuitous. It hardly puts his belief about the fate of the Neanderthals on a stronger epistemic footing than Jack's. In these discussions ‘epistemic peer’ is an idealization. To decide when and how real disagreement should affect real epistemic agents we need either to construe ‘epistemic peer’ more generously or to take seriously possible as well as actual epistemic peers. Rather than introduce possible peers, I will construe peer disagreement more generously, so that epistemic peers are those who have pretty much the same relevant evidence, reasoning powers, training, and background information. This choice is a matter of expository convenience. My discussion could equally be cast in terms of actual and possible peers, where the standard conception of an epistemic peer is used.

Either way, the unwelcome consequence of Feldman's view is that it recommends suspending judgment in a vast number of cases. We would be forced to skepticism about such things as the fate of the Neanderthals, the causes of the Civil War, the problem of free will, the Kennedy assassination, and so on. On Christensen's view we would not necessarily be forced to skepticism, but disagreements about such matters would require us to moderate the strength of our beliefs. The fact that some of Joe's epistemic peers are incompatibilists means that he should not strongly believe that free will and determinism are compatible.

4. Hyperresoluteness

So far, we have considered the cases where peers disagree, and the choices we have entertained are between moderating one's views and standing firm—roughly, between being spineless and being stubborn. Epistemic agents who moderate their beliefs in the face of disagreement seem spineless, **(p.58)** abandoning their convictions as soon as a serious challenge appears on the scene. Resolute epistemic agents seem stubborn, simply insisting that there must be something wrong with their opponent's reasoning, since there is plainly nothing wrong with their own. To see the way out of this predicament, let us look at a more extreme case raised by Peter van Inwagen.⁴ Call this the David Lewis Problem.

David Lewis believed that infinitely many possible worlds exist, each of them just as real as the actual world.⁵ There is no denying that he believed this. Moreover, there is no denying that he was incredibly smart, philosophically gifted, and intellectually responsible. He examined the arguments for and against his position with enormous care. It is no false modesty for me to say that David Lewis was a far better philosopher than I am. Nevertheless, I think he was wrong. I cannot refute his position; it is admirably well defended. But, despite Lewis's intelligence and arguments, I do not believe that there exist real possible worlds, consisting of material objects and inaccessible from the actual world.

I believe that the only world is the actual world. I think that my belief is reasonable. But David Lewis thought otherwise. He was not my epistemological peer; he was my epistemological superior. So should I not revise my opinion to agree with him? After all, if a knowledgeable physicist tells me that, despite what I think, electrons are not material particles, but clouds of energy, I revise my belief to accord with hers. So in some cases, at least, it seems epistemically reasonable to defer to my epistemic superiors. Is my disagreement with Lewis different?

Van Inwagen's answer is similar to Kelly's. In explaining why he thinks it is reasonable to retain philosophical convictions with which Lewis disagrees, he says:

I suppose my best guess is that I enjoy some sort of philosophical insight [with respect to these issues] that, for all his merits, is somehow denied to Lewis. And this would have to be an insight that is incommunicable—at least *I* don't know how to communicate it—for I have done all that I can to communicate it to Lewis, **(p.59)** and he has understood perfectly everything I have said, and he has not come to share my conclusions.⁶

Van Inwagen thinks that the disagreement shows that Lewis must be mistaken, even though he cannot say what the mistake is. Kelly would add that Lewis, being wrong, must be irrational.

I cannot speak for van Inwagen. But, speaking for myself, I think it is exceedingly unlikely that I enjoy any sort of philosophical insight that Lewis lacked (except, perhaps for the utterly question-begging insight that I am right and he is wrong, which even if true is utterly question-begging). Nor can I conclude, as advocates of resoluteness think I should, that in this case Lewis's reasoning is flawed. The position is amply, publicly, brilliantly defended. The number of able philosophers who cannot find a defect in the argument is legion. And, in response to an endless barrage of criticisms and incredulous stares, Lewis re-examined his position often. Granted, there may nevertheless be an extremely subtle flaw in Lewis's reasoning. But, on the available evidence, it is sheer hubris to insist that there must be. I do not believe that Lewis was being irrational. Should I, conceding Lewis's epistemic superiority in metaphysics, endorse realism about possible worlds? If not, should I at least follow Feldman's advice and suspend belief? Although I am not Lewis's *peer*, I might be close enough to a peer for it to be epistemically permissible for me to suspend belief rather than going over to his side. But even this seems excessively open-minded. I do not even think that Lewis might be right on this matter. Am I being hyperresolute? Should I be?

5. The Solution

Luckily, there is an easy solution to the David Lewis Problem. Unfortunately, it simply unmask a further problem. For, if we accept this solution, as I think we should, we must conclude that recent debates about the epistemic consequences of disagreement rest on a mistake.

The solution is this: despite the fact that Lewis's position is brilliantly constructed, admirably defended, and beautifully argued, I find it incredible. **(p. 60)** I simply cannot believe it. Since 'ought' implies 'can,' that I cannot believe it entails that it is not the case that I ought to believe it. And that I cannot believe that it might be true entails that it is not the case that I should suspend belief or lower my degree of belief that the only real world is the actual world. It is philosophically interesting and perhaps troubling that a position I find utterly incredible admits of such a strong defense. But my belief in a unique world is not in jeopardy.

One might think that this solution is available in extreme cases like the David Lewis Problem but not in ordinary cases like the ones we started with. That there are infinitely many real possible worlds is incredible; that the Neanderthals were an evolutionary dead end seems not to be. Jonathan Adler argues otherwise.⁷ Belief is responsive to evidence. Given a body of evidence, there is no choice about what to believe. So, even if it is not a priori incredible that the Neanderthals were an evolutionary dead end, when Jill surveys the evidence she finds it incredible, given that evidence, that the Neanderthals were an evolutionary dead end. In light of the evidence, she cannot believe it. Different epistemic agents might assess the evidence differently and so come to different beliefs. But this is not a matter of choice. They come to different beliefs because the evidence affects them differently.

Belief is not voluntary. Belief aims at truth in the sense that a belief is defective if its content is not true. If believing were something we could do or refrain from doing at will, the connection to truth would be severed. If Jack could believe that Neanderthals were an evolutionary dead end just because he wanted to, then his believing that Neanderthals were an evolutionary dead end would not amount to his thinking that 'Neanderthals were an evolutionary dead end' is true. For nothing about the fate of the Neanderthals is affected by what he wants. This is Bernard Williams's point.

If I could acquire a belief at will, I could acquire it whether it was true or not; moreover, I could know that I could acquire it whether it was true or not. If in full consciousness, I could will to acquire a 'belief' irrespective of its truth, it is unclear that . . . I could seriously think of it as a belief, i.e. as something purporting to represent reality.⁸

(p.61)

Although evidence or other epistemic considerations can move me, since belief is not voluntary, my reaction is not anything I *do*. Being responsive or unresponsive to evidence, argument, or peer pressure is something that happens to me. This means that I cannot follow Kelly's recommendation that I hold fast to my belief in the face of peer disagreement or Feldman's recommendation that I suspend belief. How peer disagreement affects my belief is not up to me. I may find myself with a belief suspended as a result of evidence, argument, testimony, or disagreement; or I may find my belief unmoved by evidence, argument, testimony, or disagreement. But my response is not under my control. Debates about whether I should suspend belief in the face of peer disagreement are wrong-headed. They are like debates about whether I should be less than 6 feet tall. I do not have any choice.

6. Assessment

Although involuntariness does not exempt responses to disagreement from assessment, it affects the sort of assessment they are subject to. Plenty of things that are beyond our control are subject to assessment—the weather, for instance. It is a miserable day today; it would have been preferable if the rain had held off until after the parade. This is a perfectly respectable assessment that does not impute fault. If belief is not subject to direct voluntary control, then assessing someone's reaction to disagreement is similar to assessing the weather. 'It is (or is not) regrettable that Jack suspended belief' is like 'It is (or is not) regrettable that it rained on the parade.' The epistemological issue under dispute, then, is whether a better constellation of beliefs results when one suspends belief in the face of peer (or superior) disagreement or when one resolutely retains one's belief. The locus of assessment is the constellation of beliefs, not the actions of the believer.

If such assessments are assessments of doxastic rationality, then assessments of doxastic rationality are assessments of what happens to us cognitively. In that case, they are like 'Jim is smart' and 'Joan is creative.' If suspending belief (or retaining belief) is cognitively good in cases of disagreement with epistemic peers or epistemic superiors, and 'rationality' is the predicate that characterizes that sort of goodness, then 'He is rational' is the same sort **(p.62)** of praise as 'She is smart.' Both characterize cognitively valuable attributes that their subjects happen to possess. Neither characterizes an attribute that is under their subjects' control.

This construal of rationality would allow us to characterize those who respond correctly to peer disagreement as rational. But it would have the consequence that doxastic rationality and practical rationality diverge. Practical rationality presupposes control. Actions are voluntary, and the rationality of an action depends on what an agent voluntarily does, given her beliefs, desires, preferences, and so on. Behaviors such as sneezes, spasms, and snores are exempt from assessment as rational or irrational, because they are involuntary. If practical rationality depends on what we do and doxastic rationality depends on what happens to us, then the concept of rationality bifurcates. The term 'rational' indicates something quite different when applied to beliefs and when applied to action. This could be so, but it is an awkward result. It raises the question why we use the same term for two such different phenomena.

7. Indirect Control

Since beliefs are not voluntary, an epistemic agent cannot, even through judicious assessment, bring it about that she retains, or lowers her degree of belief, or suspends belief in the face of a disagreement. She may, however, be able to affect her responses indirectly. Pascal recognizes this in his discussion of the wager.⁹ He does not think that one could come to believe that God exists simply by appreciating that it would be prudent to believe that God exists. But he thinks that appreciating that it would be prudent to believe that God exists gives a person reason to put himself in a position to improve his prospects of acquiring the belief that God exists. By engaging in religious practices, interacting with religious people, and avoiding irreligious people and situations, Pascal maintains, a person maximizes his prospects of being moved by factors that foster the belief that God exists. Education has a similar effect. By learning about the cognitive force of evidence, argument, and expertise, students can be put in a position to be moved by considerations of one sort or another. And, as both Pascal (**p.63**) and the educators recognize, epistemic agents can learn to appreciate why it might be worthwhile to maximize their prospects of forming, retaining, revising, and rejecting beliefs of different kinds. Arguably then, we are rational vis-à-vis our belief that *p*, not directly because we are moved by the evidence for *p*, but because we properly put ourselves in a position to be so moved. If so, doxastic rationality is a sort of practical rationality. It applies to strategies for acquiring beliefs, not to beliefs themselves.

If this is so, then the real issue about the epistemological implications of disagreement is not whether an epistemic agent should retain or revise a belief in the face of peer disagreement. It is whether she should put herself in a position to be moved by such disagreement or put herself in a position to stand fast in the face of it. Either option would be, to a significant extent, a consequence of education. If an epistemic agent learns to appreciate the merit of her opponent's position or the value of his insights, she might find herself moderating her views when they conflict with his. If she learns the perils of skepticism and spinelessness, she might find her resistance to epistemic peer pressure strengthened. Such responses are effects of cognitive character formation. The question then is what sort of character we ought to form. There is, of course, no guarantee that our beliefs will respond as we hope they will. But by subjecting ourselves to the right influences, we maximize our prospects.

Such cognitive character formation is epistemically valuable. It is cognitively worthwhile to be able to appreciate how evidence, argument, and expertise bear on the tenability of a thesis, even where we cannot believe its conclusion. Although Hume recognizes that we cannot long sustain a skeptical attitude, he thinks that the arguments that lead to skepticism are important, for they reveal the flimsiness of our grounds for belief.¹⁰ Even if we cannot help but believe, we are better off knowing that our cognitive house is built on sand.

8. From Belief to Acceptance

If we retain the focus on belief, the relation to voluntariness and epistemic responsibility is distant at best. Let us consider an alternative. L. Jonathan **(p. 64)** Cohen distinguishes between belief and acceptance.¹¹ To believe that p is to feel that p is so. To accept that p is to adopt a policy of being willing to treat p as a premise in inferences or as a basis for action. Let us modify this slightly and say that to accept that p is to adopt a policy of being willing to treat p as a premise in assertoric inference or as a basis for action where our interests are cognitive. The reason for the restriction to assertoric inference is to screen off premises used in reductios. The reason for the restriction to cognitive interests is that a premise accepted because it fosters non-cognitive ends—because it is consoling or amusing, for example—is epistemologically irrelevant.

Acceptance and belief are distinct. I can believe that p , and yet refuse to use it as a premise in inference or a basis for action. I might, for example, consider my evidence inadequate. Suppose I have what I consider an unfounded fear of frogs. (I was frightened by a frog at an early age and never quite got over it.) In that case, even though I cannot help but feel that frogs are dangerous, I refrain from using 'frogs are dangerous' in any cognitively serious inference. Nor do I act on my fear of frogs. I do not, for example, arm myself against them. Analogously, I can accept that p , even though I do not believe that p . I accept that frogs are not dangerous, when I include 'frogs are not dangerous' among the premises I am prepared to use when deliberating about the perils of summer camp, and when I allow my toddler to wade in a frog pond.

I suggest that the epistemic issue raised by peer disagreement is best seen as an issue pertaining to acceptance rather than belief. Accepting is something we do; it is an action. Hence it is voluntary. So we can, at will, continue to accept that p in the face of peer disagreement, or suspend acceptance that p . On this reading, if moderationists are right, peer disagreement should affect what we accept. It should influence the inferences we are prepared to make and actions we are prepared to perform. If advocates of resoluteness are right, it should not.

It might seem that this brings us back to where we started, or close enough to make no difference. The dispute, as I reconstrue it, concerns accepting for cognitive purposes rather than believing. But the issues are the same; the considerations favoring each side are the same; and we are no closer to finding a decisive reason to favor either holding fast to one's **(p.65)** cognitive commitment or suspending that commitment in the face of peer disagreement.

I do not think this is so. The shift from belief to acceptance reconfigures the epistemic terrain. As Feldman points out, when our goals are purely cognitive, suspending judgment is always an option. He argues that it is the option we should exercise when faced with a disagreement with our epistemic peers. He recognizes that this pushes in the direction of skepticism, but considers such skepticism plausible. If we focus exclusively on the question whether to affirm p or to affirm that not- p or to affirm neither, suspending belief in the circumstances seems the safest thing to do. But the switch to acceptance highlights the fact that all three options have costs. To suspend acceptance vis-à-vis p is to adopt a policy of refraining from using either p or not- p as a premise in assertoric inferences or as a basis for action. This is a cognitively impoverishing stance. We have fewer premises available to reason with. If, for example, Jack suspends acceptance of 'Neanderthals were an evolutionary dead end', he cannot use it as a premise. If 'Neanderthals were an evolutionary dead end' were the only problematic claim in the neighborhood, suspending acceptance might be reasonable. But in cases like this, where evidence is sparse and equivocal, peers are apt to disagree about a host of related issues.¹² The extent of Neanderthal tool-making, the structures of their communities, their level of cognitive development, their resistance to disease are all controversial matters. What a paleontologist thinks about one of these issues is apt to be enmeshed with his position on others. There are apt to be complicated patterns of agreement and disagreement across the community. Jack may disagree with Jill about some matters and with Jen about others and with Joe about yet others. To suspend acceptance of all of them leaves the paleontological community with few premises about their subject matter, yielding a flimsy and moth-eaten fabric of cognitive commitments. It is not clear how they should reason about the paleolithic period, if they can deploy only premises about which no peer disagrees. Suspending acceptance in the face of peer disagreement in cases like this would probably be excessively cognitively costly. **(p.66)**

Whether, in the face of peer disagreement, to continue to accept that p , to come to accept that $\text{not-}p$, or to suspend acceptance vis-à-vis p is a practical question. It depends on what we gain and what we lose under each alternative, and what costs are worth paying. Which premises must we abandon? How central are they? Again, this might seem to collapse into the original debate, or something close to it. Standardly, epistemologists take it for granted that, if the evidence favors p strongly enough, it is rational to accept or believe that p ; if it favors $\text{not-}p$ strongly enough, it is rational to accept or believe that $\text{not-}p$. If the evidence is about equally balanced, it is rational to suspend judgment. The standard of rationality is utterly general; everyone with the same evidence should respond the same way. The issue raised by peer disagreement is this: if Jack takes the evidence to support p strongly enough, and Jill, his epistemic peer, takes it to support $\text{not-}p$ strongly enough, then Jack either needs somehow to discredit his epistemic peer, or take the disagreement to show that the evidence does not support either p or $\text{not-}p$ strongly enough to justify belief or acceptance. Insisting that we count the costs does not seem to settle the issue.

9. The Value of Disagreement

The objectionable feature of resoluteness is the requirement that, to hold fast to their opinions, peers must construe each other as irrational. This is awkward, since they cannot point to any flaws in their opponents' reasoning and since they recognize that their opponents construe them as irrational on the same grounds. Moreover, if we think about actual cases of sustained peer disagreement, the charge of irrationality seems unfounded.

Jury deliberations are a familiar example. Some jurors think the defendant is guilty; others think she is not guilty. They all have the same evidence and, let us assume, the same reasoning powers. They disagree because they assess the evidence differently. It is clear to everyone that some of the evidence offered at the trial is misleading. A member of an opposing gang placed the defendant near the scene of the crime. A member of her own gang said that she was across town. The jurors disagree about which witness is reliable. Some doubt the first, since she bears the defendant a grudge; some doubt the second witness, since she seems like the sort who would lie to help her friend. Neither of the witnesses comes off as a stellar character. Jury **(p.67)** members might also disagree about the weight that attaches to various bits of evidence. How significant is it that the weapon was never found? How directly does the statistical evidence bear on a case like this? What should they make of the absence of fingerprints? It is not obvious that any of the jurors is irrational. If they cannot resolve their disagreement, the result will be a hung jury, the judicial equivalent of suspension of judgment. But, it seems, each juror could rationally retain his belief, while recognizing the rationality of his opponent's.

The drive to consensus in jury deliberations derives from their role in criminal trials. A hung jury is a disappointing outcome, for obvious practical reasons. But it is not clear that this point generalizes. Persistent disagreement in science or philosophy is not obviously a bad thing.

Consider the disagreement between materialists and dualists in the philosophy of mind. Materialists accept that whatever is is material; dualists accept that there are irreducibly mental entities or processes as well as irreducibly physical entities and processes. Each side can point to some conspicuous explanatory successes. But each side faces serious difficulties. Either there are outstanding problems that it cannot solve, or the solutions it offers seem inelegant, strained, and ad hoc. What is worse, the serious problems that each faces seem straightforwardly handled by the other. The dualist has a problem explaining the causal link between the mental and the material; the materialist can simply maintain that the connection is straightforward physical causality. The materialist has a problem accounting for qualia and what-is-it-like-ishness; the dualist takes these features to be distinctive marks of the mental. This is all familiar.

The standard epistemological view would maintain that in such a case everyone should go with the balance of evidence. There is a tipping point. Until the evidence reaches the tipping point, everyone should suspend acceptance. Once it is reached, everyone should accept whichever side the evidence favors. But, as Philip Kitcher argues, it is not obvious that our cognitive objectives are best achieved by everyone's marching in lock step to the same conclusion.¹³ When the reasons favoring each side of a dispute are sparse or exceedingly delicate, or the evidence is equivocal, or each side can solve important common problems that the other cannot, it may be better for the epistemic community as a whole that some of its members **(p.68)** continue to accept each position. In that case, materialists can in good conscience continue to accept materialism. Dualists can in good conscience continue to accept dualism. Agnostics can suspend judgment. Each group then can draw on a different range of commitments for premises in their reasoning and as a basis for their actions. By developing their positions, they put them to the test. Arguably, the only way we will ever find out whether materialism can solve the hard problem of consciousness is for materialists wholeheartedly to accept materialism and push it to its limits.

This position does not require denying that the overarching epistemic goal is accepting or believing only what is true. It simply notes that, where there is a significant chance that my opponent's view is true, if I want to believe only what is true, I would be well served by not foreclosing inquiry prematurely. If I can recognize that my opponent is rational and might (although I strongly doubt it) be right, then I have reason to hope that she retains her position, develops it, and either comes (as I believe she will) to see the error of her ways or (however unlikely) develops an argument that will demonstrate to me the error of mine. A convinced materialist then has sound epistemic reasons to tolerate dualism.

Such tolerance has limits, though. As we have seen, the mere fact of disagreement is not enough to make each party's position creditable. Nor is the mere fact of peer disagreement. Tolerance of disagreement is epistemically valuable only when the disagreement is among parties who have sufficient expertise in an area that their opinions are individually worth accepting, and where the evidence at hand is equivocal. When these conditions are met, a community of inquiry may best be served if epistemic peers resolutely reason and act on opinions about which other equally competent inquirers disagree. In such cases, peers who disagree have reason to consider each other wrong but not irrational. Perhaps in the fullness of time the disagreement will be resolved. That remains to be seen.

Notes:

(1) Richard Feldman, "Epistemological Puzzles about Disagreement," in Stephen Hetherington (ed.), *Epistemology Futures* (Oxford: Oxford University Press, 2006), 216–36; Hilary Kornblith, "Belief in the Face of Controversy," Ch. 2, this volume; David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review*, 116/2 (2007), 187–217.

(2) Thomas Kelly, "The Epistemic Significance of Disagreement," in Tamar Szabo Gendler and John Hawthorne, eds., *Oxford Studies in Epistemology*, i (Oxford: Oxford University Press, 2005), 167–96; Richard Foley, *Intellectual Trust in Oneself and Others* (Cambridge: Cambridge University Press, 2001), 83–131.

(3) Feldman, "Epistemological Puzzles about Disagreement," 229.

(4) Peter van Inwagen, "It is Wrong Everywhere, Always, and for Anyone to Believe Anything upon Insufficient Evidence," in Eleanor Stump and Michael J. Murray (eds.), *Philosophy of Religion: The Big Questions* (Oxford: Blackwell, 1999), 274. I consider only one of the disagreements with Lewis that van Inwagen mentions.

(5) David Lewis, *Counterfactuals* (Cambridge, MA: Harvard University Press, 1973); *On the Plurality of Worlds* (Oxford: Blackwell, 1986).

(6) Van Inwagen, "It is Wrong Everywhere, Always, and for Anyone," 274.

(7) Jonathan Adler, *Beliefs Own Ethics* (Cambridge, MA: MIT Press, 2002), 55-64.

(8) Bernard Williams, "Deciding to Believe," in *Problems of the Self* (Cambridge: Cambridge University Press, 1973), 148.


(9) Blaise Pascal, "The Wager," in *Pensées*, ed. Roger Ariew (Indianapolis: Hackett, 2005), 212-14.

(10) David Hume, *An Enquiry Concerning Human Understanding* (Indianapolis: Hackett, 1977), 25-37.

(11) L. Jonathan Cohen, *An Essay on Belief and Acceptance* (Oxford: Oxford University Press, 1992).

(12) Adam Elga maintains that, if they disagree on a host of related issues, they do not share enough to be epistemic peers. This strikes me as implausible. See Adam Elga, "Reflection and Disagreement," *Noûs*, 41/3 (2007), 478-502, and Hilary Kornblith, Ch. 2, this volume, for an argument to the contrary.

(13) Philip Kitcher, "The Division of Cognitive Labor," *Journal of Philosophy*, 87 (1990), 5-22.



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Rational Disagreement Defended

Earl Conee (Contributor Webpage)

DOI:10.1093/acprof:oso/9780199226078.003.0005

Abstract and Keywords

This chapter formulates a rational uniqueness principle holding that those who are epistemic peers on a proposition, in that they know that they share all rational considerations concerning the truth of the proposition, cannot be justified in having different attitudes toward it. It then argues against the principle, primarily on the grounds that such peers may rationally regard themselves as differing in their basis for rational belief, or their evidence, on the topic. The rationality of their differing perspectives can justify having different attitudes toward the disputed proposition.

Keywords: disagreement, epistemic peers, rational uniqueness, evidence, rationality

1. Introduction

We are rational only if we heed the dictates of reason. We heed our epistemic reasons by taking the doxastic attitudes that align with their support. If the bearing of our reasons on the truth-value of a proposition is unequivocal, then just one doxastic attitude is epistemically rational for us to take toward the proposition. Rational uniqueness principles imply that our epistemic reasons determine the doxastic attitudes that are uniquely epistemically rational for us to take.¹ The most defensible of these principles are restricted to cases of extreme duplication among those whom the principles imply to have the same justified attitudes. These epistemic peers must have identical total evidence, or they must have identical reasons, reasoning abilities, and inclinations to reason. Such restrictions are apparently excessive. It seems equally credible to require only a sharing of all that is rationally relevant to a proposition's truth-value.

Are some such relaxed uniqueness principles correct, or can thoroughly reasonable people share everything relevant to a disputed proposition and still rationally disagree? It will be argued here that the relaxed principles are not correct. Sharing the relevant reasons allows that the parties involved **(p.70)** rationally hold certain views about rational differences. Some such views can yield actual rational differences among them.

A few terminological stipulations will refine the issue. Someone has a *rational doxastic attitude* toward a proposition when the person's belief, disbelief, or suspension of judgment on the proposition is fully epistemically reasonable under the circumstances. For instance, in an evidentialist view, an epistemically rational doxastic attitude is one that fits with the person's evidence concerning the proposition. Other views of epistemic rationality have it depend on coherence or probability, given the person's beliefs or knowledge. No theory of rationality need be employed, though. The notion of being fully epistemic reasonable is well enough understood pre-theoretically for it to be applied directly.

We seek a notion of *epistemic peers* that is less restrictive than that of having exactly the same total rational bases. We do want to require of the peers a mutual possession of all that seems rationally relevant, in the interest of preserving maximal plausibility for the resulting uniqueness principle. To this end, a first requirement is that the peers have a thoroughly "shared" basis for reasonable attitude formation. People have thoroughly shared rational bases when either their bases do not differ, or, if they do, the differences are mutually known.²

Requiring epistemic peers to have mutual knowledge of any differences in their bases for rational attitudes gives us a demanding notion of a shared basis. A demanding notion makes plausible the idea that the same attitudes would be rational for epistemic peers. On the other hand, requiring mutual knowledge of any differences in bases for attitudes, rather than identical bases, allows for the virtually inevitable differences in details of some relevant experiences.

This sort of sharing also avoids the issue of whether all bases for rational belief can be identical across individuals. It may be, for instance, that we can have what Peter van Inwagen calls “incommunicable insights” in support of propositions.³ It may be that some can receive some such special reason **(p.71)** in support for a given proposition by having an incommunicable insight while others cannot. Even though the specific reason acquired by having any such insight cannot be communicated, the fact that one is having an incommunicable insight favoring some proposition can be reported. This fact about someone can be known by another on the basis of sufficiently trustworthy testimony. Having such knowledge is “sharing” the reason in the intended sense. The rest of what is required by the current notion of epistemic peers is that they have the same abilities to form reasonable doxastic attitudes.⁴

We all differ in the details of our bases for rational attitudes, and none of us knows all about the differences in anyone else's rational bases. So genuine epistemic peers are ideal cases. But there may be experts who are epistemic peers on topics of their expertise. And it is likely that many people share their limited information on a given topic. People are *epistemic peers on the topic of a proposition* when they have a thoroughly shared basis and capacity for reasonable doxastic attitudes concerning the proposition. An evidentialist thinks of topical epistemic peers as having thoroughly shared the evidence that they have that bears on the truth-value of the proposition and having equal capacities to respond to that evidence with fitting doxastic attitudes. Other theories have other things shared by topical peers. Again, no theoretical conception of the basis for peerage is needed.

Finally, people and their doxastic attitudes *disagree* about a proposition when their doxastic attitudes toward the proposition differ. Disagreements, in this sense, go beyond belief versus disbelief. They include a suspension of judgment on a proposition by one and belief or disbelief in the proposition by another. They also include differing strengths, or degrees, of belief in the same proposition, if belief varies in that way.

Here is a credible thesis that requires agreement among epistemic peers.

(RU) Rational Uniqueness. If A and B are epistemic peers on the topic of proposition X, then the same doxastic attitude toward X is epistemically rational for A and B.

(p.72)

(RU) is plausible. Epistemic peers on the topic of a proposition share epistemic reasons on the topic. It is reasonable to think that, on balance, any given batch of reasons either supports a proposition, or supports its negation, or does neither. It is reasonable to think that having on-balance support of a proposition makes belief the uniquely rational attitude toward it, having on-balance support of the negation makes disbelief uniquely rational, and in the absence of either sort of support, suspension of judgment is the uniquely rational attitude. It is reasonable to conclude that a unique attitude toward a proposition is rational for all topical peers. That is what (RU) asserts.

Objections to (RU) will be presented here. The objections use the perspectival nature of rational support against the thesis. More specifically, the objections use the thought that among the propositions that can be rational from someone's perspective are views that differentiate the rationality of attitudes among one's epistemic peers. It will be contended that sharing such views rationally distinguishes among the peers. Shared reasons can imply that individual differences affect the rational force of those reasons.

2. A First Version

A first version of an objection to (RU) will get us most of the way. It begins with an assumption to the effect that a rational case can be made for a certain sort of epistemic difference between peers. The difference concerns how they are rationally affected by an experience that is had by only one of them. The case for a difference in rational impact of an experiential difference may be refutable in the end. That is compatible with the assumption.

Let us develop the objection with an example. Suppose that Smith has an intuition of the truth of a principle of mereology—namely, the universal fusions principle:

(X1) For any things, there exists a fusion of those things.

Smith's epistemic peer about mereology, Jones, considers (X1) and fails to have any such intuition. Both Smith and Jones know that Smith has this intuition and Jones does not. Furthermore, Smith and Jones have been rationally convinced by an argument about intuitions. The argument **(p.73)** concludes that intuitions offer essentially private rational support. More precisely, the argument concludes that one who is having an intuition gets strong reason from it for the intuited proposition, making the proposition rational for the person to believe in the absence of defeaters, while one who considers the proposition and does not have the intuition gets no reason to believe the proposition from knowing that another intuits it.

Many find this claim about the privacy of intuitive support plausible.⁵ What the current objection to (RU) assumes is stronger. The assumption is that a sufficiently strong case has been made to Smith and Jones for this private support claim for it to be rational for them to believe the conclusion about how an intuition had only by one of them affects rational belief.

The objection continues. Suppose that Smith and Jones know that, apart from the intuition by Smith in favor of (X1), they have no reason to regard (X1) as true or as untrue. Now the reasoning against (RU) draws an inference. Their knowledge that their differing intuitive response to (X1) is their only basis for a rational difference concerning (X1), together with their rational belief in the claim about the private support of intuitions, makes it fully reasonable for each of them to believe that they differ in what doxastic attitude it is rational to take toward (X1). They are both fully reasonable in believing this.

(D) It is rational for Smith to believe (X1) and it is not rational for Jones to believe (X1).

So each of these two topical peers about (X1) quite reasonably believes that (X1) is rational for Smith to believe and (X1) is not rational for Jones to believe. Now for the next step in the objection. Here is a credible epistemic claim. It is rational for us to have, or to lack, any attitude that we rationally regard as rational for ourselves to have, or to lack. (A more careful version of this claim will be formulated and given critical scrutiny soon.) If this plausible claim is true, then, in our example of Smith and Jones, their rational beliefs in (D) suffice to make (D) true. That is, their rational beliefs in (D) make it rational for Smith to believe (X1) and not rational for Jones to do so. **(p.74)**

Finally, suppose that Smith and Jones do what we just inferred to be rational for each to do. Smith rationally believes (X1) and Jones rationally refrains from believing (X1). We have arrived at the conclusion that there are epistemic peers about the topic of (X1) who rationally disagree about (X1). This contradicts (RU).

3. Problems with the First Version

There is trouble for this objection. Suppose that we can legitimately get as far as the inference to the rationality of Smith and Jones believing (D). The next step relies on this principle:

(RR) Rationality from Rationality. If S is rational in believing that X is rational for S to believe, then X is rational for S to believe, and if S is rational in believing that X is not rational for S to believe, then X is not rational for S to believe.

(RR) is the principle by which we went from Smith and Jones having a justifying argument for a rational difference to Smith and Jones actually rationally differing.

(RR) is doubtful. One source of doubt arises from the apparent possibility of radical mistakes about what is rational. The doubt can be developed as follows. There are cases in which someone makes a reasonable mistake about what is required for a concept to apply. For instance, someone might reasonably think that gold has to be yellow, or that an ordered sequence has to have a first member. If some such reasonable mistake is used in an argument for an application of the concept, it can justify some more or less radical mistake about the application of the concept.

For an example involving the concept of rationality, let us suppose that one of our epistemic peers, Smith, is told by a seemingly quite reliable source that a rational belief is any belief that is not irrational. Smith is as credibly told that an irrational belief is any belief that is silly to hold. Using these rational beliefs, Smith judges that it is rational for her to believe a certain proposition that from her viewpoint is merely as likely as not, though she has guessed that it is true. For instance:

(E) An even number of stars went nova in our galaxy in the first millennium AD.

(p.75)

(E) is not downright silly for Smith to believe. It is a guess with even odds of being correct. If we assume that Smith rationally believes that not being silly is sufficient for a belief to be rational, then we can infer that it is rational for Smith to believe that it is rational for her to believe (E). From this, by (RR), we can further infer that it is rational for Smith to believe (E). Yet clearly (E) is not rational for Smith to believe under the circumstances. (E) is a guess on Smith's part.

Smith is seriously wrong about rationality. Being rational is being fully reasonable. It is not nearly enough for a belief to be fully reasonable that it is not silly. The belief must at least be favored by reason. Yet for Smith (E) is merely not silly. This argument seems to show that (RR) is untrue.

At best, this argument against (RR) needs tightening. It is a mild assumption that it is possible for some mistakes about the application of concepts like rationality to be reasonable. For the case to be a counterexample to (RR), more than that is required. It is required that, even though Smith is making severe errors about rationality, she still might be rational in the thought that her belief in (E) is rational. Yet Smith is so far off about what makes for rationality that it is difficult to be sure that Smith's thought is really about rationality. She is inferring the thought from the mere fact that (E) is not silly for her to believe. In actuality, that fact is far from enough to imply that (E) is rational for her to believe. To carry conviction as a counterexample to (RR), Smith's reasons would have to be bolstered. It would have to be argued that Smith is getting fully justifying reason to think that fully reasonable belief is nothing more than belief that is not silly. It is not clear that this could be done.

Exactly why not, though? What upper limit on conceptual mistakes would Smith have to exceed?

In any event, there are other grounds to doubt (RR). The rationality of a belief is an all-things-considered evaluation. Having strong reason to accept a proposition is not enough. The reason might be defeated. When someone, Jones, has a fully reasonable belief that Jones's belief in X has the status of rationality, this may occur because Jones sees a strongly supporting reason that Jones has for X, and misses the fact that Jones also has some less-than-obvious defeater for the reason. Not only might Jones fail to notice a defeater, but also Jones might fail to appreciate a noticed defeater as such. An argument that rationally persuades Jones of the rationality of Jones's believing X may make erroneous though reasonable **(p.76)** claims either about the role of defeat in general or about the epistemic relevance of a particular defeater. Whether from such an argument, or from a failure to recognize a defeater, it could be rational for Jones to believe that Jones's believing X is rational, when the support Jones has for X is in fact defeated and hence Jones's believing X is not in fact rational.

To highlight this problem for (RR), it is helpful to contrast (RR) with the following principle about evidence:

(EE) Evidence for Evidence. If S has evidence for the proposition that evidence exists in support of X, then S has evidence for X.⁶

The idea behind (EE) is that any evidence that supports to us the proposition that some evidence exists for X in effect tells us that X has something to be said in its favor. When a person has evidence that something favors the truth of X, the person has reason to think that X has something epistemic going for it. This reason is intuitively one epistemically rational consideration in favor of X's truth. That confirms (EE).

(EE) does not imply that someone who has evidence of evidence for X is justified in believing X. A belief's being justified is a summary evaluation. Defeat can prevent the justification to someone of a proposition for which the person has evidential support. (EE) allows this. Again, it asserts only that one who has evidence of evidence for X has some evidence for X, whether or not the evidence for X is defeated.

In contrast, (RR) asserts that, when a proposition of the form—X is rational for S to believe—is rational for S, X has for S the summary epistemic status of being rational to believe. Yet neither the dependence of rationality on lack of defeat, nor the ways in which defeat can occur, are inescapably obvious. An actually defeated proposition could be fully reasonably regarded as rational to believe.

4. A Revised Version of the Objection

(RR) is stronger than is required to make a case for rational disagreement between epistemic peers. A principle closer to (EE) is sufficient. The **(p.77)** rationality of a rational difference need only differentiate the strength of the shared reasons.

Suppose that Smith and Jones heed some argument that makes it rational for them to believe:

(D) It is rational for Smith to believe (X1) and it is not rational for Jones to believe (X1).

We have lately noted that (D) could be rational to someone who overlooks, or fails to appreciate, a defeater that makes the content of (D) untrue. But it seems that whatever support makes (D) fully reasonable for Smith would at least give Smith *a reason* to believe (X1).

It seems clear that the support that Smith has for the rationality of her believing (D) gives her reason to believe (X1), whether or not it is a defeated reason. Further, in our example it is reason for Smith to believe (X1) that Jones does not have, since (D) relates Jones to (X1) only by *denying* that it is rational for Jones to believe (X1). Finally, it seems that, in the absence of a defeater, the support of that differentiating reason for (X1) to Smith, and not to Jones, would make it fully reasonable for Smith, but not Jones, to believe (X1).

This reasoning replaces (RR) with a principle about a reason. The new principle is this:

(RE) Reason from Rationality. If S is rational in believing that it is rational for S to believe X, then S thereby has a reason to believe X, and if S is rational in believing that it is not rational for S to believe X, then S does not thereby have a reason to believe X.

(RE) is more defensible than is (RR). The first conjunct of (RE) requires S to have some reason for thinking that S's belief in X is rational. S must have some such reason, if the rationality of S's believing X is to be supported by S's reasons on balance, as the rationality of the belief requires. Whatever reason does this, it indirectly bears positively on the truth of X. The reason in effect argues that S is in a good position to affirm X. Anything arguing for that would seem at least to enhance S's epistemic position for regarding X as true. Things that enhance someone's epistemic position for taking this attitude are epistemic reasons for the person.

The reason that S gets for X, when rational belief in X is supported to S, might be the supported rationality of the belief itself. The rationality of S's **(p.78)** believing X intuitively bears positively on the truth of X. It is thus plausible that a proposition of the form—X is rational for S—is itself qualified to be a reason for S to think that X is true. If S always gets a reason to believe X, either from this proposition, when it is rational for S, or from S's reason for believing it, or from both, then the first conjunct of (RE) is correct.

It is compatible with always having some such reason, as (RE) implies, that something else always defeats it. For instance, in the example of Smith and the proposition about stars going nova, (E), Smith receives testimony that apparently makes it rational for her to think that beliefs are rational when they are not silly. If the testimony does this, then (E) qualifies to Smith as one of her rational beliefs, since she knows that (E) is for her a 50-50 guess and thus not silly. By (RE), Smith thereby has a reason to think that (E) is true. Yet the basis of the rationality for her of believing (E) includes that she is just guessing that (E) is true. It is plausible that her awareness that (E) is a guess defeats the support that she has for (E) from its rationality for her. With the support defeated, the balance of Smith's reasons that bear on (E) does not favor (E). If all of this is correct, then (RR) is mistaken about this situation whereas (RE) is not.^{7,8}

In all of this, (RE) is in the spirit of (EE). Just as evidence for evidence is plausibly thought to be evidence, whatever defends the rationality of thinking a proposition to be true seems to constitute a defeasible reason in support of the proposition's truth.

(RE) says more. The rest is also broadly similar to (EE). If S is fully reasonable in thinking that X is *not* rational for S to believe, then S thereby has reason to *doubt* that one is in a good position to affirm the truth of X. Accordingly, (RE) denies that one gets reason to believe X on this basis.

(RE) thus seems reasonable. **(p.79)**

It is important to note that for S to meet the antecedent conditions in the second conjunct of (RE)—for S to have grounds on which it is rational to deny that X is rational for S to believe—is for S to have a potential defeater of any reason in support of X that S might also have. To see this, suppose that one peer has knowledge of the existence of an undefeated reason for X that another peer has. Suppose that this knowledge of the other's reason for X gives the one peer a reason of her own for X. Yet suppose further that it is rational for the one peer to think that her belief in X is *not* rational. From such a person's perspective, there must be something flawed or counterbalanced about the reason that she gets from the other peer's reason to believe X. It must be a reason that misfires in her own case, since she fully reasonably thinks that it fails to make X rational for her. Since that is what her perspective indicates to her, she is thereby made less reasonable in believing X. The same defeating effect occurs in the case of whatever positive reasons for X that she possesses. Each time, her reasonable doubt that X is rational for her casts doubt on the support to her of the reason. We can state this defeating effect in a further principle:

(DD) Defeat from Doubt. If S is rational in believing that X is not rational for S to believe, then S thereby has a defeater that at least weakens the support of any reason that S has to believe X.⁹

The revised reasoning begins in much the same way as the original. Again, Smith and Jones, who are epistemic peers on the topic of proposition (X1), are presented with some supporting argument for this:

(D) It is rational for Smith to believe (X1) and it is not rational for Jones to believe (X1).

Again we assume that the supporting argument for (D) is good enough to make it rational for Smith and Jones to believe (D).

Now the new principle (RE) comes into play. Given this rationality of (D) for Smith, it follows by (RE) that Smith has a reason to believe (X1). This reason is not defeated, since we are assuming that Smith and Jones have no reasons for or against (X1) apart from Smith's intuition and **(p.80)** whatever that implies. Nothing in (RE) suggests a defeater for the reason. (RE) further implies that Jones lacks this reason to believe (X1).

It is possible that Jones gets a reason to believe (X1) from Smith's intuition. Jones, being Smith's epistemic peer on the topic of (X1), knows that Smith has the intuition. This knowledge might be a reason for Jones to believe (X1).

The existence of a reason that Jones gets for (X1) from knowledge of Smith's intuition of (X1) is not clear. The supporting argument that Jones has for (D) itself asserts that another's intuition is not a reason for Jones. So, from Jones's point of view, what gives Smith a reason does not do so for Jones.

But let us suppose that Jones's knowledge of Smith's intuition is a reason for Jones to believe (X1). Still, as we noted in proposing (DD), the rationality for Jones of this proposition—(X1) is not rational for Jones—at least somewhat defeats any reason that Jones might get for (X1). Whatever reason supports that proposition to Jones creates a doubt about her having support on balance for (X1). This doubt prevents any new reason, including knowledge of Smith's intuition, of (X1) from making belief in (X1) fully reasonable for her. So, even if Jones does get a reason for believing (X1) from knowledge of Smith's reason, it is weakened, if not entirely neutralized, by the rationality for Jones of the proposition that (X1) is not rational for Jones to believe. Thus, (X1) is less well supported to Jones than it is to Smith.

If belief comes in degrees, then refuting (RU) does not require the complete defeat of any reason that Jones gets for (X1) from her knowledge of Smith's intuition. As long as Jones is less rational to any extent in believing (X1) than is Smith, Jones's fully rational degree of belief is less than that of Smith. So we have a difference in rational attitudes between epistemic peers. Whatever reason Jones gets from knowledge of Smith's intuition, it seems clear that Jones's grounds for believing that (X1) is not rational for her would at least make a more tentative belief in (X1) fully reasonable for Jones than for Smith. (RU) mistakenly implies that the same attitude is fully rational for both.¹⁰ **(p.81)**

If belief does not come in degrees, then we have only that belief in (X1) is less well supported for Jones than for Smith. Differing support among peers goes against the spirit, if not the letter, of (RU). Furthermore, we can adjust the case to induce a difference in fully reasonable attitude. We can add some other defeater that both Smith and Jones possess of the support for (X1) from Smith's intuition and Jones's knowledge of it. The other defeater can be calibrated so that the difference in support between Smith and Jones places Smith beyond whatever minimum support is needed for fully reasonable belief in (X1). The effect of this other defeater on Jones, with her support for (X1) weakened by the defeater (D) as it applies to her, is that her support for (X1) does not reach the minimum. The identity of such a supplemental defeater depends on details about when support suffices for belief and how strongly (D) defeats Jones's support for (X1) from knowledge of Smith's intuition (if indeed Jones gets any such support to defeat). The result is that (RU) is refuted by such a case.

5. Bootstrapping?

One concern about this objection is that (RE) might seem to imply an unacceptable sort of bootstrapping.¹¹ (RE) seems to imply that new reasons for believing a proposition are created merely by noticing old ones. To illustrate this concern with a different sort of example, suppose that Jones has strong perceptual support for believing some ordinary observational proposition:

(X2) A tree stands before me.

Jones thus has a perceptual reason to believe (X2). Suppose that Jones reflects on how her perceptual and background information give her reason to believe (X2), and she does not turn up anything that seems to her at all to counterbalance or undermine that support for (X2). After Jones **(p.82)** does this, it seems rational for Jones to believe that it is rational for her to believe (X2). (RE) implies that Jones thereby gets a reason to believe (X2). This is a new reason, a reason that arises from the considerations that make the belief rational that Jones's belief in (X2) is rational. Those considerations include her introspection of her perceptual support for (X2)—support that she already has—and they also include the more abstract consideration to the effect that the particular perceptual support is a reason that she has for believing (X2). Thus, the perceptual support for (X2) is one reason, and this new reflective outcome is something different. The outcome makes it rational for her to believe that (X2) is rational for her to believe. Therefore, according to (RE), the reflective outcome is also a reason for Jones to believe (X2). Thus, it seems that according to (RE) Jones can in this way get a new reason in support of (X2) that enhances (X2)'s rationality for Jones. Yet when we stand back and consider what reasons develop from the perception, it can seem that Jones has no new evidence for (X2). It can seem that her real evidence for (X2) is just the perceptual support, in the presence of suitable background information, that makes the experience good reason to believe (X2). Jones's simply taking note of these things can seem *not* to add evidence for (X2). It may thus seem that the implication of (RE) that the rationality of (X2) for Jones *would* be enhanced by the reflection shows that (RE) is incorrect.

No fault in (RE) has been located here, though. When things go as described in the present example, (RE) does imply that Jones has a new reason to believe (X2). But (RE) implies nothing about the reason's strength. The new reason may constitute a minimal additional increment of rational support. In the example, the implied new reason derives from considering other, directly supporting perceptual evidence for the proposition. It is plausible that such reflections do add something to the rationality of the belief. By reflecting as described, Jones gains some appreciation of her evidential situation concerning (X2). That makes (X2) at least a bit more rationally secure for Jones than it was prior to the reflection. Coherentists about rational belief will find this thought especially congenial. But rationality need not consist in coherence for reflective rational appreciation to add a reason. Appreciating the support of an old reason is a new reason. This bears out (RE). **(p.83)**

6. Transmutation?

Let us consider again how a defender of (RU) sees any argument about our epistemic peers Smith and Jones for the conclusion

(D) It is rational for Smith to believe (X1) and it is not rational for Jones to believe (X1).

A defender of (RU) must find some fault in any such argument. In our illustrative example, the argument for (D) depends on a claim about the essential privacy of the support given to a proposition by an intuition of its truth. A defender of (RU) will hold that this claim is mistaken. The defender is likely to point out that an epistemic peer must know that the other peer does intuit the proposition. A defender is likely to hold that this knowledge of another's intuition gives as much support to the intuition's content as does the intuition itself. Such a defender would conclude that any argument that depends on denying this is therefore unsound. Thus, from the (RU) proponent's perspective, the reasoning against (RU) that we have been considering depends on an unsound argument for (D). Yet the argument is supposed to make (D) true. That may seem fishy. It may seem to imply the transmutation of a bad argument—or, at least, an arguably bad argument that has not been defended—into a good one.

The objection to (RU) does not fail in that way. The unsoundness of the defense of (D) can be granted. The way in which a proponent of (RU) faults the argument for (D) is to contend that it relies on a falsehood. The reasoning against (RU) is compatible with that. What the reasoning needs from the defense of (D) is just that it can be rational for epistemic peers Smith and Jones to conclude (D) on the basis of some such argument. To dispute this, a proponent of (RU) would need some new consideration, beyond the falsehood of a premise. It would have to be argued that (D) could not even be made rational to epistemic peers by any such argument. The new line of argument would have to be extremely wide-ranging in its capacity to deny rationality. Peers might have highly varied bodies of evidence about what makes for rationality.

For instance, let us suppose that Smith and Jones are bright undergraduate epistemic peers on the topic of mereology who are considering (X1) in a first metaphysics course. They gain their rational belief in (D) from an **(p.84)** argument given by their professor during a classroom discussion in which they air their differing views concerning (X1). The professor argues for (D) as follows.

The rational attitude is the one that the person epistemically ought to have. Epistemic ought implies can. Smith, you cannot help but believe what you intuit, including (X1). So you epistemically ought to believe (X1) and belief is therefore your rational attitude toward the proposition. Jones, you know that Smith intuits (X1) but you do not intuit (X1) yourself, you have no other reason to believe (X1), and you are able not to believe (X1). What one fails to intuit and one is able not to believe requires for rational belief a better reason than just having knowledge of the existence of someone else's intuition in its favor. Thus, (D) is true.

This reasoning could at least make (D) rational for Smith and Jones under such circumstances. It is clear that they might have good evidence for the premises of the argument. For instance, their philosophy instructor might go on to give the premises impressive testimonial endorsement as being among the most secure of epistemic assumptions, while no doubts about the premises happen to occur to them. It is unlikely that a proponent of (RU) could establish that no such argument ever yields rational belief in the likes of (D).¹²

7. Arguing against Rational Uniqueness from (EE)

(RE) is plausible from a variety of perspectives about rationality. (RE) is not particularly evidentialist in its appeal or in its interest. (EE) is also broadly plausible. (EE) is entirely about evidence for propositions, rather than about epistemic rationality, whatever that turns out to be. So (EE) is of particular interest to evidentialists. Still, since (EE) is broadly plausible, an argument using (EE) against (RU) might be of wide interest too.

For this reasoning we need an evidentialist counterpart of (D). Here is one.

(DE) Smith has evidence for (X1) and Jones does not have evidence for (X1).

(p.85) We can suppose that both Smith and Jones know that Smith intuits (X1) and Jones tries and fails to do so. They know this by telling one another in a thoroughly credible way. Now we add that they also have good evidence for the claim that intuitions are evidence only for those who have them. One potential source of this evidence is that ever-handy source of good evidence for anything—testimony from a trustworthy source. Smith and Jones have it on the authority of their metaphysics professor that intuitions work as evidence in this way and we can suppose the professor to be a paragon of integrity as far as they can tell. For good measure we can add that Smith and Jones have other evidence for the claim. Their professor offers them a best explanation argument. He argues to them that the essentially private support that is provided by evidence from intuitions best explains why eminent philosophers can reasonable disagree about issues within their expertise. Here is his argument.

Eminent philosophers know the evidence on all sides concerning issues of their philosophical expertise, and this includes knowledge of the intuitions had by those of their peers who hold opposing views. Eminent philosophers heed the totality of their evidence about their philosophical views. Yet they retain belief in their own views. Part of the best explanation for this retention must be that knowledge of others' intuitions is not evidence for those who lack them.

This line of argument at least gives evidence for (DE) to people in the position we can suppose to be that of Smith and Jones, able undergraduates who see no weakness in the argument.

With Smith and Jones having evidence that the evidence from intuitions is private, (EE) goes to work. Since they have evidence that Smith has evidence for (X1), by (EE) Smith has such evidence. Nothing in the example implies that this evidence that Smith has for (X1) is defeated.¹³ **(p.86)**

Nothing in the example asserts that Jones has evidence for (X1). Also, (EE) does not imply this, because Jones has no evidence that she has such evidence. But it might be that Jones *does* have evidence for (X1). The evidence might consist in Jones's knowledge that Smith intuits (X1). That might be evidence for (X1) to Jones, in spite of what the reliable testimony and best explanation argument from their professor contend.

We can assume that Jones has her knowledge of Smith's intuition as evidence for (X1). Any such evidence is defeated by the testimony and the best explanation argument. They make it rational for Jones to believe that Smith's intuition is not evidence of (X1) for Jones. The basis on which that proposition is rational for her, or the rationality of proposition itself, at least weakens the support of the evidence that she has for (X1) from her knowledge of Smith's intuition.¹⁴ Nothing in our example suggests that Jones has any other evidence for (X1). So we can add that Jones does not have any such evidence.

Thus, assuming that one's evidence for a proposition is epistemically rational support for it, we have the implication that Smith has undefeated rational support for (X1) and Jones has at best weaker support for (X1). If there are degrees of belief, then differing degrees of belief are fully reasonable for the two, and (RU) is refuted.

If there are no degrees of belief, then it may be that that the evidence of each on balance supports the proposition, though to differing extents. Because of this, it may be that belief is the doxastic attitude that is rational for both. But the difference in rational support is again of considerable interest. Any such difference again runs contrary to the spirit of (RU). The underlying idea is that epistemic peerage implies a sharing of all rational considerations. The implications of (EE) for the present example show that this is not so. And, as with our first objection to (RU), we can supplement the example so as to refute (RU). We can adjust the rest of the evidence that Smith and Jones share with the net effect that Smith is beyond the **(p.87)** minimum for fully reasonable belief while Jones, with her weaker support for (X1), does not reach the minimum.

8. Conclusion

(RU) is about shared reasons on the topic of a proposition. It claims the sufficiency of these reasons for a unique rational doxastic attitude toward the proposition. We have found objections to this claim. People who have shared reasons on a topic may rationally regard themselves as differing in their basis for rational belief on the topic. The rationality of this perspective is enough to differentiate the support provided by their reasons or their evidence.

The objections to (RU) would be avoided by a principle that excluded the differences that allow the objections. The objections do not apply to a principle that is about peers who are epistemically alike in certain ways beyond sharing reasons and evidence on a topic. Suppose that a principle requires that peers also have no rational basis to think otherwise about themselves. That is, the peers must also have no reason or evidence to think that they differ in any reason or evidence that either pertains to the topic or that they have reason or evidence to think pertains to the topic, or reason or evidence to think to be such reason or evidence, and so forth, all the way up. The objections would not apply to such principles.

The intuition that underlies a uniqueness principle such as (RU) might be that our perspective on a topic settles the identity of our rational attitudes on the topic. This is not correct. At least, the epistemically relevant aspect of our perspectives is not exhausted by our reasons and shared evidence on the topic. This much leaves out our perspective on reasons and evidence. That aspect of our perspective is relevant to topical epistemic rationality too. Our reasons influence our reasons.

(EE) may be refutable. Suppose that you have evidence that I have played a hoax on Smith. You have learned that I arranged the conference room so that it looked to Smith as though the following falsehood was true.

(JR) Jones was in the room today.

(p.88)

You have learned that I placed Jones's coffee mug in the room where Smith knows that Jones usually leaves it, usually recovering it later before she departs campus. Having learned this, you have evidence that there is evidence, possessed by Smith, in support of (JR). But, since you received this information about the evidence only in learning that it exists as a result of a hoax, it seems that you do not thereby have evidence for (JR). You have no other evidence for (JR). So (EE) appears to go wrong in this sort of case. (This example is quite similar to the Clever Car Thief case that Peter Klein uses against the transitivity of the confirmation relation in *Certainty: A Refutation of Skepticism*.¹⁵)

This sort of case is not conclusive against (EE). In the example you might in fact have evidence for (JR), as (EE) implies. It might be that you have evidence for (JR) in virtue of the fact that you know, among other things, that the room appeared to Smith as though Jones had been there today. We might be inclined to overlook this evidence that you have for (JR) as we consider the example, because we see that it would not amount to your having a reason to believe (JR). This is true because the evidence is given to you embedded in thoroughly defeating counterevidence. You learn that Jones had evidence of Smith's presence in learning that it was part of a hoax about Jones's presence.

There is no need for the objection against (RU) to determine the success of the objection to (EE). Whether or not (EE) can be thus sustained, the examples used against (RU) do not include putative evidence that is acquired in combination with its own defeat.

The hoax sort of example does refute a stronger principle about evidence from evidence that more literally implements the slogan “evidence from evidence is evidence”:

(EE*) If S has evidence, E*, for the proposition that evidence exists in support of X, then S has E* as evidence for X.

In the hoax example, some of your evidence for the proposition that evidence exists for (JR) is the whole basis of your knowledge that a hoax is being perpetrated to make it appear to Smith that (JR) is true. This manifestly supports to you that (JR) is not true and thus is not evidence to you for (JR).

A principle of intermediate strength survives the example, if the original (EE) does:

(EE') If S has evidence, E', for the proposition that evidence exists in support of X, then S has some evidence that is at least included in E' as evidence for X.

(p.89) The relevant “included” evidence that you get in the (JR) example is the information that it appears to Jones that (JR) is true.

The lines of reasoning against (RU) seem most open to objection where they distinguish between the support that different peers get from certain reasons or evidence. (D) and (DE) deny that Jones gets rational support from a source that they affirm to work for Smith. It is plausible that these denials are in fact mistaken. It is plausible that in fact Jones gets support for (X1) by knowing about Smith's intuition of (X1). If this is correct, then Jones differs rationally from Smith concerning (X1) only if the support that Jones gets for (X1) from this knowledge is weaker.

One way in which it would be weaker is if it were true that knowing about another's intuition *must* be less supporting than having the intuition oneself. But it is quite mysterious why that would be true no matter how well justified was the belief in the other's intuition.

The other way the support would be weaker is the way defended in the two lines of reasoning presented here against (RU). Having a defeater can weaken the support from knowledge of another's intuition. The two lines of reasoning assert that this happens because certain propositions are made rational to Jones that defeat the support for (X1) that derives from Jones's knowledge of Smith's intuition of (X1).

A defender of (RU) might question either that the purported defeater propositions really could be made rational to Jones, or that their rationality really makes them defeaters.

It would be ill advised for a defender of (RU) to deny that Jones could get any support for any proposition denying that Jones has some reason or evidence. Testimony and argument can support anything.

Better for the defender of (RU) to focus on the supported claims. The claims deny that one's knowledge of another's intuition is reason or evidence for oneself. In order to have rational support for those claims, one must grasp them. A defender of (RU) could maintain that possessing the concepts of intuition, reason, and evidence that are used in the claims brings with it certain conceptual knowledge. Possessing the concepts gives their bearer knowledge that another's intuitions must give one who knows about them equally good reasons and evidence. The defender would infer that this conceptual information that one would have to have defeats the potential defeater of one's support from knowledge of another' intuition.¹⁶ So **(p.90)** the purported defeater, itself defeated, could not defeat the support for the intuited proposition from that knowledge. Thus, the support of a proposition, intuited by at least one peer, that is possessed by all of the epistemic peers, would turn out to be the same all the way around, just as (RU) implies.

The objections to (RU) need not have asserted rational support for the proposition that there is special private support from intuitions in particular. The supported proposition could have been about special private support from episodes of memory, perceptual experience, or any experiential source. So the conceptual claim by a defender of the principles would have to be about all of that. The conceptual claim would have to be that the concepts of experience, reason, and evidence invariably inform their bearers, concerning any experiential episode that might be reasonably thought to be a source of reasons or evidence for a proposition to the one who undergoes the episode, that it equally supports the proposition to all who know that it occurred.

That claim seems not to be justified to all who acquire the concepts of reason and evidence and consider how they bear on intuition, experience, or the like. The view that intuitions give special support to those who undergo them has significant credibility. It is easy to consider this view at length without having any sense that it gives rise to a conflict with something else for which we have justification, either from the very concepts involved or otherwise. It seems clear that there is at least room for rational doubt of the claim asserting that there is no such special support. Yet this room for doubt gives the reasoning against (RU) all the foothold that it needs. Strong enough support for the privacy claim can then render it fully reasonable.

Notes:

(1) Our epistemic reasons could determine which are our epistemically rational attitudes toward a given proposition, while the reasons allow that more than one attitude is epistemically rational. Rational uniqueness principles go farther. They imply that a single doxastic attitude is rationally allowed by the reasons. (I am grateful to Tim Williamson for emphasizing this distinction in a discussion of an earlier draft of this chapter.) This difference does not affect the arguments of the chapter. They oppose a uniqueness principle by arguing that differing attitudes are uniquely rational.

(2) “Mutual knowledge” is a convenient phrase. What is actually crucial for issues about rationality is not mutual knowledge but mutual fully reasonable belief. So that is the official assumption. On the topic of epistemic peers who may have differing but shared evidence, see Richard Feldman, “Some Epistemological Puzzles about Disagreement,” in Steven Hetherington (ed.), *Epistemology Futures* (Oxford: Oxford University Press, 2006), section IVA, “Private Evidence,” 216–36.

(3) Peter van Inwagen, “It is Wrong Everywhere, Always, and for Anyone to Believe Anything upon Insufficient Evidence,” in Jeff Jordan and Daniel Howard-Snyder (eds.), *Faith, Freedom, and Rationality: Philosophy of Religion Today* (London: Rowman and Littlefield, 1996), 137–53. There are important questions of detail about what such insights could be, and how they might provide support. I pursue these questions in “Peerage,” *Episteme: A Journal of Social Epistemology*, 7/1 (2009).

(4) This shared ability requirement blocks differing attitudes from being rational simply on the basis of an epistemic “ought-implies-can” principle. Such principles are contestable, but the contest can be harmlessly sidestepped here.

(5) For a critical discussion of the privacy claim, see Feldman, “Some Epistemological Puzzles about Disagreement,” 222–4.

(6) Richard Feldman affirms this principle in “Some Epistemological Puzzles about Disagreement,” 223. (EE) is critically discussed in Appendix 1 below.

(7) I am grateful to Peter Vranas for comments on a previous draft of this chapter. Among other helpful things, the comments called for me to sharpen my explanation of the plausibility of (RE) in comparison with (RR).

(8) It may be that Smith's awareness of her guessing not only prevents her from having rational justification for (E) itself, but also disallows the testimony as justification for the epistemic proposition that (E) is rational for her. The awareness may do this by engaging Smith's understanding of rationality and thereby excluding a guess as a rational belief. It is doubtful that grasping rationality well enough to use the concept in thought must yield this defeat, since it is doubtful that we must see that rational belief excludes guessing in order to have the concept of rationality. But, if we must, then the present objection to (RR) would not succeed because (RR)'s antecedent conditions would not be met in such cases. (I am grateful to Jim Pryor for help here.) The failure of such objections would not salvage (RR). It remains possible for there to be mistakes about defeaters that allow rational but false self-attributions of rational belief. Such cases would still go against (RR) and not (RE).

(9) The defeating effect need not be total. It would be sufficient for the doubt to make S have overall weaker support for X. My thanks to Jim Pryor for a suggestion to this effect. I am also grateful for comments from Peter Vranas that called for more attention to this issue.

(10) In fact, even if neither Smith nor Jones gets better reason to believe (X1) from the argument about the privacy of support from intuitions, Jones would have weaker overall support for (X1) than Smith. Jones would have, and Smith would lack, the doubt about the support Jones gets from Smith's intuition of (X1). That is enough to differentiate the rational status of (X1) for them. (Thanks to Jim Pryor for a suggestion to this effect.)

(11) "Bootstrapping" is used by Jonathan Vogel to describe a certain apparently illicit way in which one might get knowledge from a reliable belief-forming process of its own reliability ("Reliabilism Leveled," *Journal of Philosophy*, 97 (2000), 614-16). Here the idea is that the reasons for a belief might be illicitly supplemented by its being rationally regarded as rational.


(12) The prospects of otherwise resisting the reasoning are discussed in Appendix 2 below.

(13) A complication: Smith knows that Jones tries and fails to intuit (X1). Would knowledge of this failure to intuit by an epistemic peer invariably defeat the support provided by Smith's intuiting of (X1)? No. For one thing, Smith and Jones might have good evidence that intuitions are fluky. They might have evidence that people often get intuitions after failing to have them, and that one's evidence varies accordingly. So, a peer's failure to intuit might be reasonably thought to be a psychological accident with no epistemic bearing on the support for a proposition from one's own intuition. For another thing, even if the other's failure to intuit would be a defeater in the absence of further evidence about defeat, they may in fact have further evidence about defeat. They may have received good evidence that failure to intuit by a peer does not discredit the justifying power of one's own intuitions. With this further evidence about how defeat works, any defeat of Smith's intuition of (X1) by knowledge of Jones's failure to intuit X2 would be itself defeated.

(14) Or at least it weakens the support to Jones of this knowledge. The principle about evidence that is the exact counterpart to (DD) implies this. Again, to imply a difference in rationality, any defeat for Jones and not for Smith is sufficient. But the differential enhancement by evidence for evidence also seems to occur, and the credibility of the two epistemic effects seems to stand (or fall) together.

(15) Peter Klein, *Certainty: A Refutation of Skepticism* (Minneapolis, MN: University of Minnesota Press, 1981), 33-5.

(16) It is not at all clear that having concepts implies having *any* knowledge about them. For one expression of doubts about this, see Timothy Williamson, *The Philosophy of Philosophy* (Oxford: Wiley-Blackwell, 2008), ch. 4, "Epistemological Conceptions of Analyticity."



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

You Can't Trust a Philosopher

Richard Fumerton (Contributor Webpage)

DOI:10.1093/acprof:oso/9780199226078.003.0006

Abstract and Keywords

This chapter discusses disagreements in philosophy. It acknowledges that learning of peer disagreement provides a reason against a philosophical belief one holds, but argues that one can discount this sort of reason against one's view on the grounds that most philosophers are unreliable on the sort of issue in question, a fact that makes it likely that the philosopher with whom one disagrees is also unreliable. However, this way of out the sceptical puzzle posed by disagreements leads to another problem, since the reliability of a philosopher making such an argument is also called into question.

Keywords: disagreement, justification, defeat, cognitive defects

And also, considering how many conflicting opinions there may be regarding the self-same matter, all supported by learned people, while there can never be more than one which is true, I esteemed it as well-nigh false all that went only so far as being probable.

(Descartes, *Discourse on Method*)

1. Introduction

On what* is surely the classical approach to epistemology, each of us must build all of our knowledge and justified belief on a foundation of evidence to which we have a privileged access.¹ Still, even within such a framework, setting aside certain skeptical concerns, we can reason legitimately *from our egocentric perspective* that there are others who disagree with us concerning conclusions we have reached. Under what circumstances can such discoveries defeat whatever justification we might otherwise have had for believing some proposition? That knowledge of disagreement (conjoined with certain critical background evidence) does *sometimes* defeat prior justification seems obvious to me, and I begin this chapter by detailing what I take to be uncontroversial examples of such defeat. **(p.92)** It seems equally obvious, however, that discovering other sorts of disagreement leaves my epistemic position with respect to what I believe relatively untouched. I try in this chapter to make a principled distinction between the cases in which the discovery of epistemic disagreement is, and the cases in which it is not, epistemically significant. I then try to apply the lessons learned to the question of whether the discovery of disagreement in such fields as philosophy and politics defeats whatever other justification one might have had for one's philosophical and political views.

2. Unproblematic Cases of Disagreement Leading to Epistemic Defeat

Case I: I carefully add up a column of figures, check my addition once, and reach the conclusion that the sum is 5,432. I surely justifiably believe this conclusion. I then discover that you just added up the same column, checked your addition, and reached the conclusion that the sum is 5,433. I further have every reason to believe that you are at least as good at elementary math as I am and are just as careful as I am. With this background knowledge, my discovery that you reached a different conclusion than I surely weakens my justification—probably defeats it.

Case II: I remember fondly my days as a graduate student at Brown and, in particular, I sometimes think about the statue outside Maxcy Hall (once the home of the Philosophy Department), a statue I seemed to remember being of Mark Antony. I think I had reasonably good justification for believing that the statue was of Mark Antony. Reminiscing with Ernie Sosa, I am told by him that the statue is actually of Marcus Aurelius. I surely just lost at least a great deal of the justification I might have had for thinking that the statue was of Mark Antony. Again, I am no doubt relying on all sorts of relevant background information—that in general Ernie has at least as good a memory as I do, that he knows Brown's campus at least as well as I do, and so on.

We must be careful in describing the way in which the above facts involving differing opinion defeat my justification. In particular, it should be obvious that it would be highly misleading to suggest that there is anything in the above examples that casts doubt on the traditional **(p.93)** egocentric conception of justification. The epistemic status of my beliefs, before and after the discovery of disagreement, is a function of *my* evidence and what *it* supports. In Case I I had at *t* justification E1 for believing a proposition about the sum of the numbers in the column (call that P). At a later time *t*2 I added to E1 another body of evidence E2 (the evidence that gave me justification for believing the relevant propositions describing the nature and existence of disagreement) where my total body of evidence no longer justified me in believing P. In general, there is nothing odd about the fact that through accumulation of evidence the epistemic status of a belief changes. As soon as I find out that someone else came to a different conclusion about the sum of the numbers, someone I have every reason to trust as much as I trust myself, I then have reason to think I might well have made a mistake.

3. Discovery of Disagreement but no Defeat

Not all discovery of disagreement leads to defeat of prior justification. The most unproblematic of such cases involve background knowledge that allows me to understand how the person with whom I disagree has reached a false conclusion. I have been told that our next department meeting will be this Friday at 3:00 p.m., and on the basis of this information take myself to have good reason to believe that the meeting will be at 3:00. Diane believes that the meeting is scheduled for Thursday at 7:00 a.m. (something about which she is vociferously complaining). But I also have evidence that another of my colleagues has played a practical joke on Diane and has deliberately misinformed her as to the time of the meeting. Diane and I disagree, and I know this, but my total body of evidence allows me to ignore the disagreement as epistemically irrelevant. It is not, of course, that I have reason to believe that Diane's belief is unjustified. Indeed, I am justified in believing that she has perfectly good reason to believe what she does. But I have evidence that she lacks, and my additional evidence allows me to see the way in which Diane's evidence is, in a sense, defective. My total body of evidence contains information that would defeat Diane's justification were it added to her evidence base. Diane herself would regard her evidence as defeated should she acquire the additional information that I possess. **(p.94)**

Or consider a slightly more subtle example. You have probably heard of the Monty Hall Puzzle. As I have heard the story, Hall himself was genuinely puzzled by a phenomenon he reported. In his game show, contestants hoping for a prize were asked to choose from three doors (call them 1, 2, and 3), only one of which hid a prize. After making a choice, the contestant was typically shown a door (say 3) behind which there was no prize. The contestant was then given the opportunity either to stay with his or her original choice or switch. Which course of action is most likely to lead to success—stay or switch? When the question was first posed to me, I was absolutely sure that it did not make any difference—that, relative to the contestant's new epistemic position, there is a 0.5 probability that the prize is behind door 1 and a 0.5 probability that it is behind door 2. The person presenting the puzzle to me assured me that I was wrong. Monty Hall himself, while sharing my intuitions, told various probability experts that “switchers” won more often than “stayers.” Eventually, I figured out how and why my strong “intuitions” had led me astray. But it took a while. When I subsequently explain the puzzle to others (who have not heard of it), the vast majority vehemently disagree with the conclusion that switching doubles the chances of winning. They are as sure as I was that that is a false, almost absurdly false, conclusion. But their vehement disagreement with me does nothing to weaken my justification for believing what I do. I have very good reason to believe that I have improved on the epistemic position in which they find themselves. This case is interestingly different from the earlier one, because it is not as if there is *available* to me evidence that was not available to those who disagree with me. Rather, there is a process that I now understand involving the *appreciation* of available evidence, a process that I have gone through and that I have good reason to believe (based on analogy) they have not gone through. Further, I have good reason to believe that, should those who disagree with me go through the process, they would end up agreeing with my conclusions.

So we have at least two general sorts of cases in which the discovery of disagreement poses no particular threat to the justification I have for believing a given proposition. One involves cases where I know that I have quite different and, importantly, better evidence upon which to base my conclusions. The other, subtly different, involves cases where I know (or have good reason to believe) that I have taken into account available **(p.95)** evidence in ways in which my critic has not. But there are still other cases, I think, in which my justification can withstand the discovery of disagreement.

Consider the following cases, superficially similar to I and II above, situations in which I am not the least bit inclined to think that the discovery of apparent disagreement defeats my justification. If I am justified in believing anything, I am justified in believing that $2 + 2 = 4$. My hitherto trusted colleague, a person I always respected, assures me today, however, that $2 + 2$ does not equal 4. Does this rather surprising discovery of my colleague's odd assertion defeat my justification for believing that $2 + 2 = 4$? Hardly. But this time we must be careful how we describe the relevant situation. When confronted by my colleague, my first (and probably last) reaction will be that he is not serious, that he does not believe what he says, and thus that there is no real disagreement between him and me. He can swear up and down on a stack of bibles that he is serious, and I will still probably conclude that he is lying. I will think that it is some kind of weird experiment or joke.

Alternatively, I might eventually conclude that he does believe what he says, but that there is some sort of verbal dispute interfering with communication.² My colleague is a philosopher, after all, and perhaps he is advancing some controversial thesis about the meaning of the identity sign. He might think that numbers are properties and that the property of being $2 + 2$ is not identical with the property of being 4 (though there might be some sort of synthetic necessary connection between the two properties). But it will be almost impossible to convince me that he really believes a contrary of what I believe. *Almost*. To be sure, the more crazily my colleague begins to behave in general, the more likely it is that I will start entertaining the hypothesis that he really was serious in denying that $2 + 2 = 4$ (in the ordinary sense in which people make such claims). But **(p.96)** that is just the point. To convince myself that he really is disagreeing with me, I would have to convince myself that he is crazy. And, as soon as I become convinced that he is crazy, I will not and should not pay any attention to what he believes. My justification for believing that he has lost his mind neutralizes whatever epistemic significance his disagreement with me might otherwise have had.

This last case is a bit different from the Monty Hall example we considered earlier. There, I had reason (based on analogy) to believe that the person with whom I was arguing had not successfully taken into account available evidence. I understood, or at least had good reason to believe that I understood, the reasons for his cognitive failure. In this last example, I do not understand what is up with my colleague. To be sure, the hypothesis that someone has gone mad is a *kind* of explanation of odd behavior, but it is a bit like explaining the ease with which an object shattered by pointing out that it was highly fragile. I do not know or understand what *in particular* is going through my colleague's mind—his mind has become a kind of mystery to me. But my general reason for thinking that it is a *defective* mind is a good enough reason for discounting the epistemic significance of his beliefs.

And I would probably say just the same thing about a friend who assures me that I have not existed for more than a day or two—that I just popped into existence *ex nihilo* replete with inexplicable vivid and detailed memories of a long past. When asked to explain this odd view, he tells me that he cannot—it is top secret, he says, and he has sworn an oath not to disclose his evidence. Again, initially, I almost certainly would not believe that there is genuine disagreement between him and me and I would retain that position until I become convinced that he is nuts. And when I become justified in believing that he is insane, I will also be justified in discounting the epistemic significance of beliefs he has that contradict mine.

Both of these examples invoke the possibility of an extreme cognitive defect. But, as I shall point out later, there are continua of cognitive defects. Bias, wishful thinking, stubbornness, intellectual competitiveness, all can affect one's ability to assess properly one's evidence, and it may be possible to reject the significance of another's belief when there is reason to suspect that the belief in question results from one of these. I will eventually argue that, whether or not one can reasonably believe that **(p.97)** one's philosophical and political opponents have some *specific* cognitive defect, there is almost always available a *prima facie* powerful reason to think that they are at least unreliable and, in that sense, defective when it comes to arriving at philosophical and political truth. The good news is that appreciating this fact blunts the discovered disagreement as a defeater for one's justification. The bad news is that the very reason for discounting the epistemic relevance of the disagreement is potentially a different sort of defeater for one's justification.

4. Some Tentative Preliminary Conclusions

My justification gets defeated in Cases I and II because I add to my initial evidence for reaching the respective conclusions new evidence that justifies me in believing that other people probably have evidence that would give them good reason to believe their respective conclusions. Furthermore, (and crucially) I have no more reason to think that their evidence is any worse than the evidence upon which I relied in believing my initial conclusion, *nor is their ability to process the relevant evidence*. I also realize, in effect, that there is a perfect *symmetry* in our epistemic situations with respect to one another. In Case I, by hypothesis, my careful addition gives me the same *sort* of evidence (no better and no worse) than your careful addition gives you. To be sure, the results of my attempt at addition cast doubt on the success of your attempt at addition. But then, by parity of reasoning, the result of your attempt at addition equally casts doubt on the success of my attempt. *Indeed, if I really do have good reason to believe that you are in general just as reliable as I am when it comes to adding columns of numbers, discovering the results of your addition would have precisely the same significance as doing the addition again myself and coming to a different conclusion*. We have all done just that. We check our figures and come to a different sum. At that point, we have no more reason to trust our present self than our prior self. All we can do is check a few more times in an effort to break the epistemic stalemate.

It is precisely the same in Case II. My apparent memory (at least when it used to be half-decent) might cast doubt on the veridicality of Sosa's apparent memory, but no more than his apparent memory casts doubt on the veridicality of my apparent memory. Unless I have some reason to **(p.98)** believe that one of us has a better memory than the other, the discovery that there is disconfirming evidence of equal strength will defeat our respective justification. Again, it is just as if I myself had conflicting memories. Such inconsistent memories would deprive me of whatever justification I might otherwise have had for believing some proposition about the past.

In discussing Cases I and II, I did ignore some very real complications, complications to which I shall return later in this chapter. I have presupposed that there is no real difficulty getting myself justification for believing the relevant propositions describing the fact that there is someone who disagrees with me, who has evidence just as good as mine, and is just as reliable as I am in processing that evidence. When thinking about such matters we would do well to keep in mind traditional epistemological problems. There really are genuine epistemological problems concerned with knowledge and justified belief about other minds. We really do have better access to what goes on in our own minds than we do to what goes on in the minds of others. I will almost always have better knowledge of my thought processes than I will of yours. It was probably too hasty to conclude that my justification would automatically get defeated by accumulation of the additional evidence described in Cases I and II. In my case, the defeat *would* probably occur, but that is only because I seem to remember being pretty bad at adding long columns of figures. I have some reason to believe that there are all kinds of people who are better, who are more reliable, at this than I am. And, sadly, I now also seem to remember seeming to remember all sorts of things that did not happen. My memory is turning on itself, leaving me in a precarious position with respect to the character of statues encountered long ago. The truth is that I trust Sosa's memory about such matters more than I trust my own. Were it not for these apparent defects in my own cognitive structure, I suspect that the disagreements I encountered in Cases I and II would leave me with a weakened justification for believing what I do, but still with more reason to retain my belief than to abandon it. By the time I have very carefully added the figures in the column three, four, five, or six times, it will start approaching the case in which my crazed colleague starts ranting about $2 + 2$ not equaling 4, and I will be unmoved by the fact that there is another who disagrees with me about the sum. Again, the relevant thought experiment involves imaginatively adding to one's own evidence the evidence that the other person possesses to see whether or not that would defeat my justification. (p. 99)

Let me emphasize again that, in order for my discovery of the results of your addition to defeat my justification, I must have good reason to believe that you are at least as reliable at addition as I am. Of course, it is often not that easy to reconstruct the evidence that would allow me to reach such a conclusion. When it comes to relatively simple arithmetic, however, it is probably nothing more exotic than an inductive generalization upon which I rely. Most educated people are fairly good at summing numbers—at least as good as I am. I infer from this that you are just as likely to be coming up with the truth as I am. And most people have relatively decent memory and are fairly reliable when it comes to arriving at true conclusions about the past based on that memory.

We can encounter disagreement without losing justification when (1) we have good reason to believe that we have a different and better evidence base than the person with whom we disagree, (2) we have good reason to believe that we have engaged a common evidence base more successfully than the person with whom we disagree, or (3) we have good reason to believe that the person with whom we disagree is cognitively defective.³

5. Philosophical and Political Disagreement

There are a host of cases that are particularly difficult and interesting for those of us in academics, particularly in fields like philosophy. When in his *Discourse on Method* Descartes remarked that “there is nothing imaginable so strange or so little credible that it has not been maintained by some philosopher or other” (1960: 13), he did not overstate his case much. Famous, respected, apparently intelligent and sane philosophers have taken diametrically opposed positions with respect to a host of issues with which they are concerned. I have thought long and hard about issues in the philosophy of mind and am a confirmed property dualist. Most of the philosophers I know reject the view. Indeed, most of the philosophers I know reject most of my views, and I nevertheless think quite highly of many of those philosophers. What epistemic significance, if any, should my knowledge of the relevant disagreement have for the epistemic status of my philosophical beliefs? **(p.100)**

The existence of radical disagreement among philosophers is, of course, hardly unique to our field. There is just as much disagreement among economists, religious theorists, and political theorists, to consider just a few. Consider the last. Most academics would view my political views as slightly to the right of Attila the Hun. For example, I think that the foreign policy of the United States over the last hundred years or so has been something of which we should be very proud. In general, the wars we fought were the right wars to fight, and, even when they were not, we fought for admirable reasons. I believe, if anything, we ought to be far more aggressive in confronting hostile nations in the Middle East and elsewhere. I know, of course, that many, indeed most, well-educated and intelligent people disagree with me. I have seen Noam Chomsky, for example, ranting on television about the evil of American foreign policy, and people in linguistics seem to think that he is, in general, a knowledgeable and intelligent person. What is the epistemic significance, if any, of my knowledge that Chomsky and his ilk vehemently disagree with me? Should I take whatever justification I might have had for my beliefs to be seriously threatened by knowledge of our difference of opinion?

Well, can I discount the relevance of philosophical or political disagreement in any of the ways that we discussed above? First, can I legitimately conclude that I have access to better or more complete evidence for my philosophical or political views than those with whom I disagree? It is obviously a difficult question. It is more difficult in the case of philosophy, probably, than in the case of politics. Part of the difficulty in philosophy stems from the fact that it is not all that easy to characterize the evidence upon which we do rely in reaching our philosophical conclusions. But it is surely the case that, on a superficial characterization of our respective evidence, most of the philosophers with whom I disagree on various issues have *available* to them the same, or better, evidence as I have. They have certainly typically read just as much as I have. They have carefully considered the same sorts of arguments as I have. They probably have more empirical knowledge than I have on a host of issues. To be sure, I have argued elsewhere (1999) that one almost never settles a genuinely philosophical controversy through the accumulation of empirical evidence. That, however, is yet another point on which I disagree with many of my colleagues. In any event, it is going to be an uphill climb to convince myself or others that I am in a privileged **(p.101)** position with respect to access to evidence that bears on philosophical problems.

The case is not much different with respect to, for example, political disagreement. Here, however, it should be fairly obvious that the rationality of political means-ends calculations is often held hostage to empirical information. But, if we are trying to decide whether or not we acted rationally in going to war with Iraq, say, it is highly doubtful that reasonable, informed people are led to different conclusions by possessing interestingly and importantly different evidence. Every educated person knows the sad history of Chamberlain's appeasement of Hitler. Every educated person knows the sad history of our futile intervention in Vietnam. Almost every educated person knows full well that some wars succeed in accomplishing noble ends and that some wars have devastatingly bad results. Just about everyone (but the odd conspiracy theorist) knows that the USA had fairly good reason to believe that Iraq had, or could easily develop again, an assortment of chemical weapons, and had long-term ambitions to develop nuclear weapons. Just about everyone knows that it was always likely that successfully promoting democracy in the Middle East is, at best, a long shot. It is hard to believe that there is a significant difference in the kind of evidence available to reasonable, well-educated people who nevertheless dramatically disagree about the wisdom of going to war with Iraq.

So it is going to be a hard sell to convince myself, let alone others, that I have reached different conclusions on philosophical and political matters from many others because there is available to me evidence that is hidden from them. Nor is it, at least *initially*, much more plausible to suppose that I can invoke the third of the strategies discussed above to discount the epistemic relevance of disagreement. While I am sorely tempted, occasionally, to view my philosophical and political opponents as suffering from some sort of madness, in my more cautious moments I am disinclined to embrace that conclusion. There are exceptions. Eliminative materialists—philosophers who seriously maintain that there are no such states as belief, being in pain, being in fear, and so on—really do present a puzzle for me. They really do strike me a bit the same way as my hypothetical colleague did who professed to have discovered that $2 + 2$ does not equal 4. When Paul Churchland (1981) appears to embrace eliminative materialism, my first instinct is to suspect that he is not really serious—that he is just messing around a bit trying to provoke an **(p.102)** interesting discussion. When I begin to suspect that Churchland and other eliminativists are serious, I am genuinely puzzled as to what is going on in their minds. They become a mystery to me. They become the kind of enigma about which I can make neither heads nor tails, and at such time I discount completely the epistemic significance of what they apparently believe (beliefs they are officially committed to disavowing—disavowals they are officially committed to disavowing . . .). But the proponents of eliminative materialism are in a world of their own. In general, I do not take such extreme and pessimistic views about the cognitive abilities of my colleagues.

In our earlier discussion of having reason to believe that others are cognitively defective, however, I focused on extreme cases. On all such matters there is, of course, a continuum. Do I have reason to suspect that some of my colleagues are plagued by more subtle defects? Perhaps I have some reason to believe, for example, that they are the victims of various biases that cause them to believe what they want to believe, or ignore evidence or arguments that they find inconvenient. Indeed, I suspect that I *do* have reason to believe that others are afflicted in such ways, though at this point I am going to stop identifying particular philosophers whose particular alleged problems cause their mistakes—many of these people are, after all, my friends. What kind of cognitive defects do I seem to find in people whose intellectual abilities I generally respect? Well, it does not take a genius to notice that many, if not most, philosophers are heavily influenced by their philosophical environment. It is surely not a coincidence that many of Sellars's students are very sympathetic to Sellars's views, that many of Bergmann's students are Bergmannians, that Harvard philosophers are not all that fond of the analytic/synthetic distinction, that the East and West coasts are awash with externalists in the philosophy of mind. The connection between intellectual environment and political views is even more pronounced. A significant majority of registered Republicans are children of registered Republicans and a significant majority of registered Democrats are children of registered Democrats.

But so what? Why would that even suggest a cognitive defect on the part of people who were influenced by others whom they respect? One can be initially caused to believe a view by factors that may not have much connection to epistemic justification, but, as we have been taught in our first logic course, the genesis of a belief must surely be distinguished from **(p.103)** its epistemic status. Before I let my suspicion that a colleague has a belief that was causally influenced by his or her intellectual environment cast doubt on the epistemic rationality of that belief, I would surely need to know a whole lot more about my colleague's present epistemic situation. Furthermore, why should I think that *I* am any better at detecting and fighting *my* philosophical and political biases than the others upon whom I am casting aspersions? I am a confirmed foundationalist and I studied at Brown—just a coincidence?

Well, here it is easy to sound a bit like an egomaniac. I do, in fact, think that I have got more self-knowledge than a great many other academics I know, and I think that self-knowledge gives me a better and more neutral perspective on a host of philosophical and political issues. I suspect that it is in part the fact that I take this belief of mine to be justified that I do think that I can sometimes discount to some extent the fact that well-known and respected intellectuals disagree with me. But I would also stress that it seems to me that I should take each controversy on a case-by-case basis. I am subjectively confident, for example, that space and time are not finite—indeed, that the hypothesis that they are is essentially unintelligible, that Euclidean geometry is not only true, but necessarily true. I am inclined to believe that there are no universals, substances, bare particulars, or objective values. I have got arguments supporting these sundry beliefs, but I would not bet huge amounts of money on any of them being sound. *One* reason for this is the suspicion that equally rational people reflecting on the relevant evidence could reach quite different conclusions. But I am not sure that it is the existence of disagreement that is, in the final analysis, doing much work. It seems to me that the real justification for a kind of epistemic modesty on these matters lies no further than my own realization that the arguments I put forth in support of my views are hardly conclusive. I can often see the attraction of alternative positions, and I understand that often I defend a position primarily to see how far I can get defending it. I am inclined to think I can get an ontology with which I am comfortable relying heavily on tropes, for example, but I often get confused thinking about the issue. I often start wondering if I even truly understand all of the terms of the debate.

Philosophy is by its very nature difficult. As I indicated earlier, I am committed to a radical foundationalism that puts me at true epistemic ease only when I have traced my justification back to a secure foundation. **(p.104)** And that is very hard to do. Despite my commitment to foundationalism, however, it is not hard to see philosophers proceeding as if their primary goal was a kind of grand coherence among their philosophical views. In practice, I suspect, we often simply start somewhere that seems halfway plausible and see whether we can embed this starting point in a consistent big picture that incorporates other things that seem plausible.⁴ By the time we have published a few articles or a book, we have a reputation to defend, and we sound as though we are willing to die in the trenches defending our positions. Reflecting on all this might well incline one to the view that, in philosophy at least, something like a coherence theory of justification provides the standards by which philosophical ingenuity, skill, and success are judged. Indeed, I think there is more than a grain of truth in all of this. We tell our students, for example, that two classmates can provide diametrically opposed critical evaluations of a position and each get an A+ for their efforts. Something like an emphasis on the value of presenting a plausible coherent story must be part of that upon which we base such positive evaluations. But it is important to realize that at least some versions of the coherence theory of justification are anathema to the idea that we should give weight to disagreement. As long as we view the justification of a person's belief as a function of that belief's coherence with the rest of *that person's* beliefs,⁵ it should become immediately obvious that the existence of another (perhaps rational) person with whom I disagree is no real threat to the justification I possess for my beliefs.

Again, let me stress that I do not subscribe to a coherence theory of justification. Coherence is neither a necessary nor a sufficient condition for the justification of a belief. I am merely pointing out that, insofar as philosophers often proceed the way I described above, I would not take the fact that a perfectly rational philosopher incorporates into his or her theory elements that contradict my beliefs to be an indicator that the philosopher has any real justification for those beliefs. And, without that conclusion, of course, my knowledge that the philosopher in question disagrees with me has no epistemic significance. Of course, to the extent that I can conclude **(p.105)** that I myself just choose positions and run with them, I should be equally cautious about claiming a positive epistemic status for my own beliefs. But, again, that has nothing to do with the existence of disagreement. The appropriate modest epistemic conclusion should be derived from the problematic nature of the data upon which my philosophical "system" is built.

If it is often reasonable to conclude that others often have biases that interfere with their reasoning in philosophy, it is even more obvious, I think, that biases corrupt reasoning in politics. The trouble, of course, is that I cannot illustrate the claim without making controversial claims that most of you will reject. But consider for a moment the overheated rhetoric employed by intelligent people taking opposed positions on controversial political issues. I have heard intelligent people—people who *obviously* know better—claim that Bush lied to get us into a war with Iraq and offer as their *evidence* that he made claims that turned out to be false. Bush may have lied, of course, but even a reasonably intelligent child can see that asserting a falsehood and lying are entirely different matters. Or consider the many who praise or criticize foreign-policy decisions based on actual consequences. If intervention in Iraq goes badly, this is taken to be proof that such intervention was a mistake. Again, I have heard many intelligent people make this kind of argument, despite the fact that in other contexts they would be perfectly capable of making the common-sense distinction between wise decisions and successful decisions. Everyone understands that a rational gamble can cost one dearly. Almost everyone surely understands that a rational gamble can even be such that it is *likely* to cost one dearly. To be clear, I am not arguing that it is obvious that the decision to use force in Iraq was correct. I am pointing out only that people who clearly are in a position to know better use arguments that they would never endorse but for a desire to reach a certain conclusion.

When I argue this way, I again risk sounding like a bit of a jerk. Do I really suppose that I am justified in thinking that there is an asymmetry between myself and others when it comes to various epistemic defects? Am *I* any less likely to be blinded to what is reasonable to believe by antecedent views or desires? Well, to be honest I suppose that I think that I am. And this brings me to the sort of considerations we considered in discussing the Monty Hall puzzle and the way in which reflection on my own thought processes lead me to dismiss the epistemic significance of those who disagree with me when I have reason to believe that they have **(p.106)** not gone through the same progression of thought. One of the things that moves me strongly to ignore the relevant disagreement in the example of the Monty Hall puzzle is my confidence that, if I lead those who reject my conclusion through the progression of thought I went through, they will eventually end up agreeing with me. And, when I am confident of my philosophical or political views, I believe I have good reason to think that with enough time and patience I can bring my opponents around—at least those I take to be genuinely reasonable. But notice how careful I was to restrict my claim to those philosophical or political views about which I am reasonably confident. I stress again that I lack such confidence with respect to a great many of the views I assert and defend. Williamson aside, knowledge is not even close to being the norm of assertion for philosophy and politics. I am not even sure that *belief* is the norm of assertion for philosophy and politics. Much of what we are trying to do is to get as clear as we can about issues that concern us, and often the best way to arrive at the truth is to get argument started with assertion.

In those situations in which I retain the confidence that with enough time and energy I can turn my opponents into allies, I probably do rely on a point made earlier. I do know how I reason better than I know how others reason. It is important to keep firmly in mind that in the final analysis there really is no alternative to the egocentric perspective. Even when my discoveries about what others believe defeat the justification I had prior to those discoveries, it is *my* discoveries that are doing the defeating. I can use the discovery of disagreement to weaken my justification only insofar as I trust *my* reasoning. Without such trust, there is no access even to what others believe. That is not to deny that trust in my reasoning ability can turn on itself—can lead me to doubt the very faculties that I trust. But when that has not happened, and when I cannot understand exactly what is going on in the minds of others, I will always turn back to the reasoning I understand best—my own.

6. A Global Reason for Suspecting that Intellectuals Suffer from Cognitive Defects

In the discussion above I focused primarily on reasons to suspect that my philosophical and political opponents might have some specific cognitive **(p. 107)** defect—that they might suffer from bias, or stubbornness, for example. I want to conclude, however, by discussing a more abstract reason for suspecting that intellectuals with whom I disagree suffer from a cognitive defect. I will focus on philosophy, but my comments will apply, *mutatis mutandis*, to political theory, and indeed a host of other fields in which intellectuals struggle to retain justified beliefs in the face of disagreement.

In Cases I and II, cases where the discovery of disagreement clearly defeated my justification, I pointed out that I needed the background evidence that the person with whom I disagreed not only had the same kind of evidence as I, but was just as good at processing that evidence. In short, I needed a justified belief that the person with whom I disagreed was just as reliable as I am when it comes to addition in Case I and memory in Case II. I also know that reliability is very much relative to a field or even a subfield. There are people in physics I trust to give me information about the physical constitution of various kinds of things. I think that I have good reason to believe that they are more or less reliable when it comes to information in their fields. But it does not take long to discover that even brilliant physicists are often hopeless philosophers. When they stray beyond the boundaries of their expertise, they are not to be trusted. So, before I take the fact that another philosopher disagrees with me to be counterevidence to what I believe, I would need good reason to believe that the philosopher in question is reliable when it comes to the discovery of philosophical truth. And how would I get that evidence?

I suppose I could try an induction of the sort I discussed with respect to reliability at addition. I could employ the premise that most philosophers are reliable when it comes to arriving at philosophical truth. But the premise is obviously false. We need only remind ourselves of Descartes's observation about philosophers. If you get ten philosophers in a room discussing any of the fundamental issues in philosophy, you are likely to get ten different and incompatible positions. If there is one thing I can be virtually certain of, it is that most philosophers are *not* reliable when it comes to arriving at interesting philosophical truth. And it does not help much to turn to “brilliant” philosophers. I would readily admit that many of the philosophers whose work I respect disagree with me. Surely, I cannot think of these philosophers as exceptionally good without thinking of them as reliable. But obviously the problem noted above has not gone **(p.108)** away. The philosophers I respect also disagree with each other, often quite radically. So it cannot even be true that most of them are reliable when it comes to the subject matter upon which they disagree. My respect cannot rationally be based on a rational judgment about their reliability. It has more to do with the considerations of coherence that I discussed earlier. Whether we are grading students or evaluating colleagues, we obviously do not do so by trying to determine what percentage of their arguments are sound.

When we cannot rely on a generalization to reach a conclusion about the reliability of someone with respect to a given subject, how should we proceed? Well one way, of course, is to ascertain independently various truths in the field, and see how often the person in question is able to discover them. But how would I do this in philosophy? My only way of discovering independently the relevant philosophical truths is to figure out myself what is true. But then the only philosophers I would deem reliable employing this method are those who end up agreeing with me (at least most of the time). And, since we can divide philosophy into subfields, and those subfields into even smaller fields, there is nothing to stop me from reaching the conclusion that a philosopher who is reliable when it comes to matters epistemological is decidedly unreliable when it comes to ethical theory. And a philosopher who is reliable when it comes to normative ethical theory might be unreliable when it comes to metaethical theory. Again, insofar as I think I am getting at the truth in my philosophical inquiry, I shall be reaching the conclusion that others (perhaps very bright people) are unreliable when they often disagree with me.

So, in the final analysis, there does seem to be a really significant difference between Cases I and II and the kind of disagreement I discover between my philosophical views and the philosophical views of others I respect. Without some basis for thinking that other philosophers are reliable when it comes to reaching philosophical conclusions, my discovery that they disagree with me cannot defeat my justification. But I have strong evidence to believe that philosophers are in general unreliable. There are so many different and incompatible answers posed to philosophical questions by even really intelligent philosophers that we can deduce that most philosophers, even most intelligent philosophers, have false beliefs about the correct answers to any interesting philosophical question. If **(p.109)** I try to check the reliability of a philosopher with respect to a given area of philosophy by presupposing my own reliability, I will obviously reach the conclusion that the only philosophers who are reliable are those who generally agree with me. Again, I will not need to worry about the discovery that others disagree with me. The fact that they do is evidence of their unreliability.

7. Out of the Frying Pan and into the Fire

It does not take long to see that we have traded one problem for another. We have discovered a plausible defeater for the potential defeater presented by the discovery of disagreement over philosophical (and other highly contested intellectual) propositions. I can discount the fact that another philosopher, for example, disagrees with me by reasonably concluding that that philosopher is probably unreliable when it comes to philosophical truth. If I have reason to believe that you are unreliable when it comes to adding numbers, I will not take the fact that you came to a different sum to present much counterevidence to my belief that I have added the figures correctly. I have overwhelming evidence that most philosophers, even most really good philosophers, are unreliable when it comes to arriving at philosophical truth. I know that most of them have false beliefs. I know that because I know that, at least typically, most philosophical views are minority opinions. Each positive philosophical view is usually such that most philosophers think that it is false, and there is typically nothing approaching a consensus on the correct alternative. I can, therefore, infer that most philosophical views are false. This is a strong reason for me to think that philosophers are not reliable, even if I am not sure precisely what the cognitive defect is that leads to their unreliability. The difficulty, of course, is that I also know that I am one of those philosophers whose reliability is under attack. This reason for thinking that my opponents are probably cognitively defective is also a reason for thinking that I am probably cognitively defective. And now I face again the task of trying to argue plausibly that I am an exception to the rule. To do so, I am back to the task of trying to convince you that you (and others) suffer from specific defects that explain your unreliability, defects I have somehow managed to avoid. **(p.110)**

References

Bibliography references:

Churchland, Paul (1981). "Eliminative Materialism and the Propositional Attitudes," *Journal of Philosophy*, 78: 67-90.

Descartes, René (1960). *Discourse on Method and Meditations*, trans. Laurence LaFleur. Indianapolis: Bobbs-Merrill.

Fumerton, Richard (1999). "A Priori Philosophy after an A Posteriori Turn," *Midwest Studies in Philosophy*, 23: 21-33.

— (2006). "Epistemic Conservatism," in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, ii. Oxford: Oxford University Press.

Goldman, Alvin (1988). "Strong and Weak Justification," in James Toberlin (ed.), *Philosophical Perspectives 2: Epistemology*. Atascadero, CA: Ridgeview Publishing Co., 51-69.

Notes:

* I would like to thank faculty and students at the University of Western Ontario, Wooster College, and the University of Iowa for helpful comments and criticisms made on earlier drafts of this chapter.


(1) The privileged access might not involve infallible justification—it might just be truths about which we have better justification than anyone else could have.

(2) Indeed, I sometimes worry that verbal dispute might be more common than philosophers like to realize. The internalism/externalism debate in epistemology, for example, sometimes strikes me as partially verbal. I understand perfectly well, I think, the concept the externalist is interested in analyzing—I can even help with suggestions concerning how to avoid counterexamples. But I am also convinced that there is another, different concept that is, and always has been, of paramount importance to many epistemologists. And I suspect that at least some externalists are willing to admit that this other concept might exist. Goldman (1988), for example, was finally willing to distinguish between what he called strong and weak justification (though I am not suggesting that I would agree with his analysis of weak justification).

(3) Where, as we shall see, cognitive defects come in degrees.

(4) If some version of epistemic conservatism were plausible, the fact that something seems to me to be true might give me foundational justification for believing it. I do not believe, however, that epistemic conservatism is plausible. See Fumerton (2006).

(5) Where the contrast is to some sort of “social” coherence theory that requires coherence of beliefs among members of a community. A social coherence theory applied to the community of philosophers pretty much precludes any justified beliefs.



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Peer Disagreement and Higher-Order Evidence

Thomas Kelly

DOI:10.1093/acprof:oso/9780199226078.003.0007

Abstract and Keywords

This chapter distinguishes and assesses a set of views about the rational response to peer disagreement. It raises a series of objections to views according to which one should assign 'equal weight' to one's own view and the views of those with which one disagrees. It concludes by defending 'The Total Evidence View', according to which what one is justified in believing in cases of peer disagreement depends upon what is supported by one's total evidence.

Keywords: peer disagreement, higher order evidence, equal weight view, total evidence, justification

1. Introduction

My aim in this chapter* is to develop and defend a novel answer to a question that has recently generated a considerable amount of controversy. The question concerns the normative significance of *peer disagreement*. Suppose that you and I have been exposed to the same evidence and arguments that bear on some proposition: there is no relevant consideration that is available to you but not to me, or vice versa. For the sake of concreteness, we might picture

You and I are attentive members of a jury charged with determining whether the accused is guilty. The prosecution, following the defense, has just rested its case.

You and I are weather forecasters attempting to determine whether it will rain tomorrow. We both have access to the same meteorological data.

You and I are professional philosophers interested in the question of whether free will is compatible with determinism. Each of us **(p.112)** is thoroughly acquainted with all of the extant arguments, thought experiments, and intuition pumps that the literature has to offer.

Suppose further that neither of us has any particular reason to think that he or she enjoys some advantage over the other when it comes to assessing considerations of the relevant kind, or that he or she is more or less reliable about the relevant domain. Indeed, let us suppose that, to the extent that we do possess evidence about who is more reliable—evidence afforded, perhaps, by a comparison of our past track records—such evidence suggests that we are more or less equally reliable when it comes to making judgments about the domain in question.¹ Nevertheless, despite being peers in these respects, you and I arrive at different views about the question on the basis of our common evidence. For example, perhaps I find myself quite confident that the accused is guilty, or that it will rain tomorrow, or that free will and determinism are compatible, while you find yourself equally confident of the opposite. Question: once you and I learn that the other has arrived at a different conclusion despite having been exposed to the same evidence and arguments, how (if at all) should we revise our original views?

Some philosophers hold that, in such circumstances, you and I are rationally required to *split the difference*. According to this line of thought, it would be unreasonable for either of us simply to retain his or her original opinion. Indeed, given the relevant symmetries, each of us should give *equal* weight to his or her opinion and to the opinion of the other in arriving at a revised view. Thus, given that I am confident that the accused is guilty while you are equally confident that he is not, both of us should retreat to a state of agnosticism in which we suspend judgment about the question. This is:

The Equal Weight View. In cases of peer disagreement, one should give equal weight to the opinion of a peer and to one's own opinion.

Recently, the Equal Weight View has been endorsed by a number of philosophers. Here, for example, is Richard Feldman (2006: 235): **(p.113)**

Consider those cases in which the reasonable thing to think is that another person, every bit as sensible, serious, and careful as oneself, has reviewed the same information as oneself and has come to a contrary conclusion to one's own . . . An honest description of the situation acknowledges its symmetry . . . In those cases, I think, the skeptical conclusion is the reasonable one: it is not the case that both points of view are reasonable, and it is not the case that one's own point of view is somehow privileged. Rather, suspension of judgment is called for.²

It is no surprise that the Equal Weight View has found sophisticated advocates; it is in many respects an appealing view. Indeed, reflection on certain kinds of cases can make it seem almost trivial or obviously true. Consider, for example, cases involving conflicting perceptual judgments such as the following:

Case 1. You and I, two equally attentive and well-sighted individuals, stand side by side at the finish line of a horse race. The race is extremely close. At time t_0 , just as the first horses cross the finish line, it looks to me as though Horse A has won the race in virtue of finishing slightly ahead of Horse B; on the other hand, it looks to you as though Horse B has won in virtue of finishing slightly ahead of Horse A. At time t_1 , an instant later, we discover that we disagree about which horse has won the race. How, if at all, should we revise our original judgments on the basis of this new information?

Many find it obvious that, in such circumstances, I should abandon my original view that Horse A won the race and you should abandon your original view that Horse B won the race. For each of us, suspension of judgment is now the uniquely reasonable attitude. We should become agnostics about which horse won the race until further evidence becomes available. This, of course, is exactly what the Equal Weight View enjoins. But one might expect that what holds for perceptual judgments holds also for judgments of other kinds, and thus, in general. **(p.114)**

Further evidence for the Equal Weight View seems to be afforded by certain natural analogies involving inanimate measuring devices. Consider for example:

Case 2. You and I are each attempting to determine the current temperature by consulting our own personal thermometers. In the past, the two thermometers have been equally reliable. At time t_0 , I consult my thermometer, find that it reads '68 degrees,' and so immediately take up the corresponding belief. Meanwhile, you consult your thermometer, find that it reads '72 degrees,' and so immediately take up that belief. At time t_1 , you and I compare notes and discover that our thermometers have disagreed. How, if at all, should we revise our original opinions about the temperature in the light of this new information?³

I take it as obvious that in these circumstances I should abandon my belief that it is 68 degrees and you should abandon your belief that it is 72 degrees. In particular, it would be unreasonable for me to retain my original belief simply because this was what *my* thermometer indicated. Indeed, inasmuch as the relevant evidence available to us is exhausted by the readings of the two thermometers, neither of us should be any more confident of what his thermometer says than of what the other person's thermometer says. In these circumstances, we should treat the conflicting thermometer readings as equally strong pieces of evidence. But—one might naturally conclude—what holds for the conflicting readings of equally reliable thermometers holds also for the conflicting judgments of individuals who are peers in the relevant respects. The mere fact that I originally judged that the accused is guilty is no reason for me to retain that view once I learn that you originally judged that he is innocent. Just as I should retreat to a state of agnosticism about whether the temperature is 68 or 72 degrees once I learn what your thermometer indicates, so too I should retreat to a state of agnosticism about whether the accused is guilty or innocent once I learn your opinion about the matter.

In view of considerations such as these and others that have been offered on its behalf, the Equal Weight View can seem quite compelling. Nevertheless, I believe that here appearances are misleading: the Equal **(p.115)** Weight View is false. The main negative burden of what follows is to show that (and why) this is so. After offering a critique of the Equal Weight View, I will use that critique as a point of departure for the development of an alternative proposal about how we should respond to peer disagreement. For reasons that will emerge, I call this alternative proposal **the Total Evidence View**.

I begin with some taxonomy.

Philosophers who hold views inconsistent with the Equal Weight View maintain that, in at least some cases of peer disagreement, it can be reasonable to *stick to one's guns*.⁴ A particularly radical alternative is this:

The No Independent Weight View. In at least some cases of peer disagreement, it can be perfectly reasonable to give no weight at all to the opinion of the other party.

That is, even if one retains one's original opinion with wholly undiminished confidence upon learning that a peer thinks otherwise, one's doing so might be perfectly reasonable.

According to more moderate alternatives, while one is always rationally required to give at least some weight to the opinion of a peer, one is not always required to split the difference. That is, even if one's new opinion is closer to one's original opinion than to the original opinion of one's peer, one's new opinion might nevertheless be perfectly reasonable. Of course, there are many possible views of this kind. We might picture these possibilities as constituting a spectrum: at one end of the spectrum sits the Equal Weight View; at the other end, the No Independent Weight View; in between, the more moderate alternatives, arranged by how much weight they would have one give to the opinion of a peer relative to one's own. The more weight one is required to give to a peer's opinion relative to one's own, the more the view in question will resemble the Equal Weight View; the less weight one is required to give, the more it will resemble the No Independent Weight View.

Among alternatives to the Equal Weight View, another distinction is worth marking. Suppose that, upon learning that we hold different opinions about some issue, neither you nor I splits the difference: each **(p.116)** of us either simply retains his or her original opinion, or else moves to a new opinion that is closer to that opinion than to the original opinion of the other. Again, according to the Equal Weight View, both you and I are unreasonable for responding to our disagreement in this way. Among views inconsistent with the Equal Weight View, distinguish between those according to which you and I might *both* be reasonable in responding in this way and those according to which *at most one* of us is being reasonable. As an example of the former, consider a view according to which everyone is rationally entitled to give some special, presumptive weight to his or her own judgment.⁵ If such a view is true, then both you and I might be perfectly reasonable, even though neither one of us splits the difference. As an example of the latter kind of view, consider a view according to which how far you and I should move in response to our disagreement depends on whose original opinion better reflects our original evidence (Kelly 2005a). Given such a view, and given certain further assumptions, it might be that, when you and I fail to split the difference, at most one of us is being reasonable.

Taking these two distinctions together, the view most radically at odds with the Equal Weight View would seem to be the following:

The Symmetrical No Independent Weight View. In at least some cases of peer disagreement, both parties to the dispute might be perfectly reasonable even if neither gives any weight at all to the opinion of the other party.

Thus, according to the Symmetrical No Independent Weight View, even if both you and I remain utterly unmoved upon learning that the other holds a different opinion, it might be that neither one of us is responding unreasonably.

It is not my purpose to defend the Symmetrical No Independent Weight View. Indeed, the view about peer disagreement that I will ultimately endorse is consistent with both it and its negation. That having been said, I am inclined to think that the Symmetrical No Independent Weight View is true. Moreover, I also believe that, precisely because it contrasts so sharply with the Equal Weight View, considering it can help to illuminate the latter by making plain some of the less obvious dialectical commitments **(p.117)** incurred by proponents of the Equal Weight View. For these reasons, I want briefly to explore what might be said on its behalf.

2. Cases in which Both You and I are Perfectly Reasonable, Despite Giving No Weight to the Other's Point of View

First, a preliminary remark about the Equal Weight View. The Equal Weight View is sometimes defended in contexts in which the propositional attitude of belief is treated as an all-or-nothing matter: for any proposition that one considers, one has in effect three doxastic options—one either believes the proposition, disbelieves the proposition, or suspends judgment as to its truth.⁶ However, in considering the Equal Weight View, it is for various reasons more natural to treat belief not as an all-or-nothing matter but rather as a matter of degree. Indeed, it does not seem that the Equal Weight View can even be applied in full generality in a framework that treats belief as an all-or-nothing matter. Thus, consider a possible world that consists of two peers, one of whom is a theist and the other of whom is an atheist. When the theist and the atheist encounter one another, the response mandated by the Equal Weight View is clear enough: the two should split the difference and become agnostics with respect to the question of whether God exists. Suppose, however, that the two-person world consists not of a theist and an atheist but rather an atheist and an agnostic. How do they split the difference? (In this case, of course, agnosticism hardly represents a suitable compromise.) In general, the simple tripartite division between belief, disbelief, and suspension of judgment does not have enough structure to capture the import of the Equal Weight View when the relevant difference in opinion is that between belief and suspension of judgment, or between suspension of judgment and disbelief. Clearly, the natural move at this point is to employ a framework that recognizes more fine-grained psychological states. Let us then adopt the standard Bayesian convention according to which the credence that one invests in a given proposition is assigned a numerical value between 0 and 1 inclusive, where 1 represents maximal confidence that the proposition **(p.118)** is true, 0 represents maximal confidence that the proposition is false, 0.5 represents a state of perfect agnosticism as to the truth of the proposition, and so on. Thus, if the agnostic gives credence 0.5 to the proposition that God exists while the atheist gives credence 0.1 to the same proposition, the import of the Equal Weight View is clear: upon learning of the other's opinion, each should give credence 0.3 to the proposition that God exists.

Moreover, even if one restricts one's attention to what are sometimes called 'strong disagreements'—that is, cases in which the relevant proposition is initially either believed or disbelieved by the parties⁷—it seems that an advocate of the Equal Weight View still has strong reasons to insist on a framework that treats belief as a matter of degree. For consider a world of three peers, two of whom are theists and one of whom is an atheist. The animating thought behind the Equal Weight View—namely, that the opinion of any peer should count for no more and no less than that of any other, would seem to be clearly violated by the suggestion that the parties to the dispute should retreat to a state of agnosticism, since that would seem to give more weight to the opinion of the atheist than to the opinion of either theist. (The atheist's opinion is in effect given as much weight as the opinions of both theists taken together in determining what should ultimately be believed by the three.) On the other hand, the suggestion that theism wins simply because the atheist finds himself outnumbered would seem to give too little weight to the atheist's original opinion, if it is understood to mean that all three should ultimately end up where the two theists begin. Once again, it seems that an advocate of the Equal Weight View should insist on a framework that treats belief as a matter of degree, since only such a framework can adequately capture what is clearly in the spirit of his or her view.

Having noted this elementary point, I will now describe a possible case in which it is plausible that you and I are both perfectly reasonable, despite giving zero weight to the other person's opinion.

Case 3. How things stand with me:

At time t_0 , my total evidence with respect to some hypothesis H consists of E . My credence for H stands at 0.7. Given evidence E , this credence is perfectly reasonable. Moreover, if I was slightly less confident that H (p.119) is true, I would also be perfectly reasonable. Indeed, I recognize that this is so: if I met someone who shared my evidence but was slightly less confident that H was true, I would not consider that person unreasonable for believing as she does.

How things stand with you:

At time t_0 , your total evidence with respect to H is also E . Your credence for H is slightly lower than 0.7. Given evidence E , this credence is perfectly reasonable. Moreover, you recognize that, if your credence was slightly higher (say, 0.7), you would still be perfectly reasonable. If you met someone who shared your evidence but was slightly more confident that H was true, you would not consider that person unreasonable for believing as she does.

At time t_1 , we meet and compare notes. How, if at all, should we revise our opinions?

According to the Equal Weight View, you are rationally required to increase your credence while I am rationally required to decrease mine. But that seems wrong. After all, *ex hypothesi*, the opinion that I hold about H is within the range of perfectly reasonable opinion, as is the opinion that you hold. Moreover, both of us have recognized this all along. Why then would we be rationally required to change? One sympathetic to the Equal Weight View might attempt heroically to defend the idea that you and I are rationally required to revise our original credences in these circumstances. However, a more promising line of resistance, I think, is to deny that Case 3 is possible at all. That is, an adherent of the Equal Weight View should endorse

The Uniqueness Thesis. For a given body of evidence and a given proposition, there is some one level of confidence that it is uniquely rational to have in that proposition given that evidence.⁸

Suppose that the Uniqueness Thesis is true. Then, if it is in fact reasonable for me to give credence 0.7 to the hypothesis, it follows that you are guilty (**p.120**) of unreasonable diffidence for being even slightly less confident. On the other hand, if you are reasonable in being slightly less confident than I am, then I am guilty of being unreasonably overconfident. Hence, the description of Case 3 offered above is incoherent; Case 3 is not in fact a possible case.

Clearly, the Uniqueness Thesis is an extremely strong claim: for any given batch of evidence, there is some one correct way of responding to that evidence, any slight departure from which already constitutes a departure from perfect rationality. How plausible is the Uniqueness Thesis? For my part, I find that its intuitive plausibility depends a great deal on how we think of the psychological states to which it is taken to apply. The Uniqueness Thesis seems most plausible when we think of belief in a maximally coarse-grained way, as an all-or-nothing matter.⁹ On the other hand, as we think of belief in an increasingly fine-grained way, the more counterintuitive it seems. But, as we have seen, the advocate of the Equal Weight View has strong reasons to insist on a framework that employs a fine-grained notion of belief.

Some philosophers find it pre-theoretically obvious that the Uniqueness Thesis is false.¹⁰ Many others accept substantive epistemological views from which its falsity follows.¹¹ Although the Uniqueness Thesis is inconsistent with many popular views in epistemology and the philosophy of science, its extreme character is perhaps best appreciated in a Bayesian framework. In Bayesian terms, the Uniqueness Thesis is equivalent to the suggestion that there is some single prior probability distribution that it is rational for one to have, any slight deviation from which already constitutes a departure from perfect rationality. This contrasts most strongly with so-called orthodox Bayesianism, according to which any prior probability **(p.121)** distribution is reasonable so long as it is probabilistically coherent. Of course, many Bayesians think that orthodoxy is in this respect overly permissive. But notably, even Bayesians who are considered Hard Liners for holding that there are substantive constraints on rational prior probability distributions other than mere probabilistic coherence typically want nothing to do with the suggestion there is some uniquely rational distribution. With respect to this long-running debate, then, commitment to the Uniqueness Thesis yields a view that would be considered by many to be beyond the pale, too Hard Line even for the taste of most Hard Liners themselves.

Of course, despite its radical character, the Uniqueness Thesis might nevertheless be true. In fact, some formidable arguments have been offered on its behalf.¹² Because I believe that the Uniqueness Thesis is false, I believe that the Symmetrical No Independent Weight View is true, and (therefore) that the Equal Weight View is false. However, especially in light of the fact that here I will neither address the arguments for the Uniqueness Thesis nor argue against it more directly, I will not appeal to the possibility of so-called reasonable disagreements in arguing against the Equal Weight View. Indeed, because I am convinced that we should reject the Equal Weight View in any case, I will proceed in what follows as though (what I take to be) the fiction of uniqueness is true. My dialectical purpose in emphasizing the apparent link between the Uniqueness Thesis and the Equal Weight View is a relatively modest one. As noted above, the Equal Weight View can sometimes seem to be almost *obviously* or *trivially* true, as though its truth can be established by quick and easy generalization from a few simple examples or analogies. However, if I am correct in thinking that commitment to the Equal Weight View carries with it a commitment to the Uniqueness Thesis, then this is one possibility that can be safely ruled out. Even if it turns out to be true, the Uniqueness Thesis is an extremely strong and unobvious claim. Inasmuch as the ultimate tenability of the Equal Weight View is bound up with its ultimate tenability, the Equal Weight View is similarly an extremely strong and unobvious claim.

I turn next to some arguments against the Equal Weight View. **(p.122)**

3. Why We Should Reject the Equal Weight View

Let us suppose for the sake of argument, then, that the Uniqueness Thesis is correct: for a given batch of evidence, there is some one way of responding to that evidence that is the maximally rational way. Consider

Case 4. Despite having access to the same substantial body of evidence E, you and I arrive at very different opinions about some hypothesis H: while I am quite confident that H is true, you are quite confident that it is false. Indeed, at time t₀, immediately before encountering one another, my credence for H stands at 0.8 while your credence stands at 0.2. At time t₁, you and I meet and compare notes. How, if at all, should we revise our respective opinions?

According to the Equal Weight View, you and I should split the difference between our original opinions and each give credence 0.5 to H. This is the reasonable level of confidence for both of us to have at time t₁. As a general prescription, this strikes me as wrong-headed, for the following reason. Notice that, in the case as it has been described thus far, nothing whatsoever has been said about the relationship between E and H, and, in particular, about the extent to which E supports or fails to support H. But it is implausible that how confident you and I should be that H is true at time t₁ is wholly independent of this fact. For example, here is a way of filling in the details of the case that makes it implausible to suppose that you are rationally required to split the difference with me:

Case 4, continued. In fact, hypothesis H is quite unlikely on evidence E. Your giving credence 0.2 to H is the reasonable response to that evidence. Moreover, you respond in this way precisely because you recognize that H is quite unlikely on E. On the other hand, my giving credence 0.8 to H is an unreasonable response and reflects the fact that I have significantly overestimated the probative force of E with respect to H.

At time t₀, then, prior to encountering the other person, things stand as follows: you hold a reasonable opinion about H on the basis of your total evidence, while I hold an unreasonable opinion about H on the basis of the same total evidence. (Again, the difference in the normative statuses of **(p.123)** our respective opinions is due to the fact that your opinion is justified by our common evidence while mine is not.) If one were to ask which one of us should revise his or her view at *this* point, the answer is clear and uncontroversial: while it is reasonable for you to retain your current level of confidence, I should significantly reduce mine, since, *ex hypothesi*, this is what a correct appreciation of my evidence would lead me to do.

For an advocate of the Equal Weight View, this seemingly important asymmetry completely washes out once we become aware of our disagreement. Each of us should split the difference between his or her original view (regardless of whether that view was reasonable or unreasonable) and the original view of the other (regardless of its status).

I take this to be an extremely dubious consequence of the Equal Weight View.¹³ We should be clear, however, about exactly which consequences of the Equal Weight View warrant suspicion and which do not. According to the Equal Weight View, after you and I have met, I should be significantly less confident that the hypothesis is true. That much is surely correct. (After all, I should have been significantly less confident even before we met.) The Equal Weight View also implies that, after we have met, you should be more confident that the hypothesis is true, despite having responded correctly to our original evidence. While less obvious, this is also—for reasons that I explore below—not implausible. What *is* quite implausible, I think, is the suggestion that you and I are rationally required to make *equally* extensive revisions in our original opinions, given that your original opinion was, while mine was not, a reasonable response to our original evidence. After all, what it is reasonable for us to believe after we have met at time t_1 presumably depends upon the total evidence that we possess at that point. Let us call the total evidence that we possess at time t_1 E^* . What does E^* include? Presumably, E^* includes the following:

Our original body of evidence E .

The fact that I responded to E by believing H to degree 0.8.

The fact that you responded to E by believing H to degree 0.2.

Notice that, on the Equal Weight View, the bearing of E on H turns out to be completely irrelevant to the bearing of E^* on H . In effect, what **(p.124)** it is reasonable for you and I to believe about H at time t_1 supervenes on how you and I respond to E at time t_0 . With respect to playing a role in determining what is reasonable for us to believe at time t_1 , E gets completely swamped by purely psychological facts about what you and I believe. (This despite the fact that, on any plausible view, it was highly relevant to determining what it was reasonable for us to believe back at time t_0 .) But why should the normative significance of E completely vanish in this way?

We can, of course, imagine a case in which it would be reasonable for one to form an opinion about H by simply splitting the difference between your opinion and mine: namely, a case in which those opinions are the only relevant evidence that one possesses. Imagine, for example, the position of a third party who lacks any direct access to E, and knows only that, of two equally well-informed parties, one gives credence 0.2 and the other gives credence 0.8 to hypothesis H. (Suppose also that the individual lacks any other relevant evidence.) For an individual so situated, assigning a probability of 0.5 to H is at least as reasonable as any other course. Perhaps the same would be true of you and me, if, at some still later time t_2 , we completely lost access to our original evidence—say, in virtue of forgetting it—while retaining our original levels of confidence. However, it is mysterious why, in cases in which we *do* have access to the original evidence, that evidence should play no role in determining what it is reasonable for us to believe but is rather completely swamped by the opinions that we form in response to it. It is a weakness of the Equal Weight View that it assimilates cases in which one does have access to the original evidence to cases in which one does not.

I find the suggestion that the original evidence makes no difference at all once we respond to it a strange one. Of course, others might not share my sense of strangeness, and even those who do might very well be prepared to live with this consequence, given that other considerations might seem to tell strongly in favor of the Equal Weight View. For this reason, I want to press the point by offering four additional arguments. I offer the first two arguments in the spirit of plausibility considerations, designed to bring out further what I take to be the counterintuitiveness of the suggestion that the original evidence gets completely swamped by psychological facts about how we respond to it. The third and fourth arguments are **(p.125)** considerably more ambitious, inasmuch as they purport to show that there is something approaching absurdity in this idea.

3.1. A Comparison: Interpersonal and Intrapersonal Conflicts

Compare the question of how it is rational to respond to interpersonal conflicts between the beliefs of different individuals with the question of how it is rational to respond to intrapersonal conflicts among one's own beliefs. Suppose that one suddenly realizes that two beliefs that one holds about some domain are inconsistent with one another. In such circumstances, one has a reason to revise one's beliefs. But how should one revise them? We can imagine a possible view according to which, whenever one is in such circumstances, one is rationally required to abandon *both* beliefs. This view about how to resolve intrapersonal conflicts is the closest analogue to the Equal Weight View. But such a view has little to recommend it. In some cases of intrapersonal conflict, the reasonable thing to do might be to abandon both beliefs until further evidence comes in. But, in other cases, it might be perfectly reasonable to resolve the conflict by dropping one of the two beliefs and retaining the other. What would be a case of the latter kind? Paradigmatically, a case in which one of the two beliefs is well supported by one's total evidence but the other is not. A normative view about how it is reasonable to resolve inconsistencies among one's beliefs that completely abstracts away from facts about which beliefs are better supported by one's evidence, and that would have one treat one's prior beliefs on a par, regardless of how well or ill supported they are by one's total evidence, would not be an attractive one. But the features that make such a view unattractive are shared by the Equal Weight View.

3.2. Implausibly Easy Bootstrapping

Consider:¹⁴

Case 5. You and I both accept the Equal Weight View as a matter of theory. Moreover, we scrupulously follow it as a matter of practice. At time t_0 , each of us has access to a substantial, fairly complicated body of evidence. On the whole this evidence tells against hypothesis H : given **(p.126)** our evidence, the uniquely rational credence for us to have in H is 0.3. However, as it happens, both of us badly mistake the import of this evidence: you give credence 0.7 to H while I give it 0.9. At time t_1 , we meet and compare notes. Because we both accept the Equal Weight View, we converge on credence 0.8.

On the Equal Weight View, our high level of confidence that H is true at time t_1 is the attitude that it is reasonable for us to take, despite the poor job that each of us has done in evaluating our original evidence. (Indeed, it would be unreasonable for us to be any less confident than we are at that point.) However, it is dubious that rational belief is so easy to come by. Suppose that when you and I meet to compare notes at time t_1 , I ask you for the evidence on the basis of which you invest such high credence in the hypothesis. You recite your evidence, evidence that in fact favors not-H over H. You then ask me for my evidence; I recite the same body of underwhelming considerations. According to the Equal Weight View, this process is sufficient to make it reasonable for both of us to have a high degree of confidence that H is true, despite the fact that, *ex hypothesi*, it was unreasonable for either of us to have a high degree of confidence before we met. But that seems mistaken.

Can the Equal Weight View be interpreted in such a way that it does not allow for such bootstrapping? A proponent might suggest the following: in response to peer disagreement, one is *rationally required* to split the difference, but it does not follow that the opinion at which one arrives by doing so is reasonable. Rather, splitting the difference is a *necessary but insufficient condition* for the reasonableness of the opinion at which one arrives. In order for that opinion to be reasonable, one must not only have arrived at it by splitting the difference, but one must have correctly responded to the original evidence as well. Thus, peers who scrupulously adhere to the Equal Weight View will wind up with reasonable opinions if they begin from reasonable opinions, but not if they begin from unreasonable opinions. In this way, the current bootstrapping objection is apparently blocked.

However, this proposed interpretation runs into serious problems elsewhere. Consider again Case 4 from above, in which you but not I respond to the original evidence E in a reasonable manner. At time t_1 , we discover our disagreement and split the difference, converging on a credence of 0.5. On the present proposal, your credence of 0.5 is perfectly reasonable, **(p.127)** since you have responded to the evidence correctly at every stage. On the other hand, my credence of 0.5 is *not* reasonable, since I misjudged the original evidence; the mere fact that I respond appropriately to your opinion by splitting the difference is not sufficient to render the opinion at which I thereby arrive reasonable. But here something seems to have gone wrong. After all: notice that, at time t_1 , you and I have exactly the same evidence that bears on H (namely, E, plus our knowledge of how each of us originally responded to that evidence), and we invest exactly the same credence in H on the basis of that evidence (namely, 0.5), yet your credence is reasonable on the evidence while mine is not. That seems wrong.¹⁵ Thus, although this interpretation of the Equal Weight View manages to avoid the charge of bootstrapping, it is untenable on other grounds. I therefore set it aside.¹⁶

It is often noted that, at least on first inspection, the Equal Weight View would seem to have relatively radical implications for our actual practice.¹⁷ After all, many of us persist in retaining views that are explicitly rejected by those over whom we possess no discernible epistemic advantage. It seems that, if the Equal Weight View is true, then many of us should give up (or, at least, become significantly less confident of) some of our deepest convictions **(p.128)** about philosophy, politics, morality, history, religion, and other subjects in which there is substantial controversy among intelligent, thoughtful, and well-informed people. Indeed, advocates of the Equal Weight View sometimes devote substantial labor attempting to show that this *prima facie* consequence is not an *ultima facie* one.¹⁸ The operative idea, it seems, is that it is surely not so difficult for intelligent, thoughtful, and well-informed people rationally to hold confident beliefs about such matters. Although I reject the Equal Weight View, I myself do not think that this consequence (if it is indeed such) should be counted as a significant theoretical cost. On the contrary, the suggestion that many or most of us tend to be too confident of our controversial philosophical, political, historical (and so on) opinions strikes me as having considerable independent plausibility. What has thus far not been adequately appreciated about the Equal Weight View is to my mind a much more damning consequence—namely, that, if the Equal Weight View is true, then there will be cases in which rational belief is *too easy* to come by. That is, views for which there is in fact little good evidence or reason to think true can bootstrap their way into being rationally held simply because two irrationally overconfident peers encounter one another and confirm each other's previously baseless opinions.

Indeed, I believe that there is significantly worse trouble for the Equal Weight View on this front.

3.3. Even Easier, and More Implausible, Bootstrapping: Single-Person Cases

On the Equal Weight View, the evidence that determines what it is reasonable for us to believe in cases of peer disagreement consists in facts about the distribution of opinion among the peers. Let us call such evidence **psychological evidence**. Let us call the original evidence on which the peers base their opinions **non-psychological evidence**.¹⁹ Above, we noted that there is at least one special case in which—as the advocate of the Equal Weight View would have it—it is highly plausible that what it is reasonable to believe is entirely fixed by the psychological **(p.129)** evidence—namely, a case in which the psychological evidence is all the evidence that one has to go on. When one is aware of nothing relevant to some issue other than facts about the distribution of opinion, it is unsurprising that such facts suffice to fix what it is reasonable for one to believe about that question. In the even more special case in which one is aware of nothing relevant other than the distribution of opinion among a group of one's peers, one should give equal weight to each of their opinions. (Crucially, these thoughts are not the exclusive property of the Equal Weight View, a point to which we will return below.)

At one end of the spectrum, then, are cases in which one's evidence is exhausted by psychological evidence concerning facts about the distribution of opinion (that is, cases in which one's non-psychological evidence has dwindled to nothing). At the other end of the spectrum are cases in which all of one's evidence is non-psychological (that is, cases in which one's psychological evidence has dwindled to nothing). Consider a case of the latter kind: at time t_0 , one possesses a body of non-psychological evidence E that bears on some question, but one is completely ignorant of what anyone else thinks about that question, nor has one yet formed an opinion about the issue oneself. Presumably, at this point a proponent of the Equal Weight View will agree that what it is reasonable for one to believe is wholly fixed by the non-psychological evidence (to the extent that what is reasonable to believe is fixed by the evidence at all). At time t_1 , one first forms an opinion about the hypothesis on the basis of this non-psychological evidence; let us suppose that one gives credence 0.7 to the hypothesis on the basis of the evidence. Assuming that one has access to facts about one's own confidence *via* introspection, one thus acquires one's first piece of psychological evidence that bears on the question. For one can now adopt a third-person perspective on one's own opinion and treat the fact that one believes as one does as evidence that bears on the truth of the hypothesis. At time t_1 then, one's total evidence consists of one's original body of non-psychological evidence E , plus a single piece of psychological evidence—namely, the fact that one believes as one does. Call this new body of total evidence $E+$:

$E+$ (one's evidence at time t_1)

The original body of non-psychological evidence E

The fact that one believes the hypothesis to degree 0.7

(p.130) Suppose that at time t_2 one gains an additional piece of psychological evidence: one learns the opinion of a peer. Suppose that the peer gives credence 0.3 to the hypothesis. At time t_2 , then, one's total evidence—call it $E++$ —consists of the following:

$E++$ (one's evidence at time t_2)

The original non-psychological evidence E

The fact that one believes the hypothesis to degree 0.7

The fact that one's peer believes the hypothesis to degree 0.3

According to the Equal Weight View, one should split the difference with one's peer and believe the hypothesis to degree 0.5 at time t_2 ; we have criticized the view at some length on the grounds that it implausibly suggests that the psychological evidence swamps the non-psychological evidence in these circumstances. At present, however, I want to inquire about what a proponent of the Equal Weight View should say about what one is rationally required to believe back at time t_1 , when one knows one's own opinion about the hypothesis but no one else's. Does the psychological evidence swamp the non-psychological evidence *even then*? It would seem that the only principled answer for the proponent of the Equal Weight View to give to this question is 'Yes.' For the proponent of the Equal Weight View will insist that, at time t_2 , what one is rationally required to believe is determined by averaging the original opinions of the two peers; moreover, if, at an even later time t_3 , one becomes aware of the opinion of a third peer, then what one is rationally required to believe will be determined by averaging the original opinions of the three peers; and if, at some still later time t_4 , one becomes aware of the opinion of a fourth peer . . . In general, for any time t_n , a proponent of the Equal Weight View will hold that what one is rationally required to believe is entirely fixed by averaging the opinions of the n peers. Why then should things be any different back at time t_1 , when the number of peers = 1? It seems as though the only principled, not ad hoc stand for the proponent of the Equal Weight View to take is to hold that the psychological evidence swamps the non-psychological evidence, even when the psychological evidence is exhausted by what you yourself believe. On this view, before one forms some opinion about the hypothesis, how confident one should be that the hypothesis is true is determined by the non-psychological evidence; after one arrives at some level of confidence—in the present example, a degree **(p.131)** of belief of 0.7—how confident one should be given the evidence that one then possesses is . . . 0.7. Of course, if one had responded to the original evidence in some alternative way—say, by giving credence 0.6 or 0.8 to the hypothesis—then the rationally required credence would be 0.6 or 0.8. On the picture of evidence suggested by the Equal Weight View, the distinction between believing and believing rationally seems to collapse in cases in which one is aware of what one believes but unaware of what others believe.

Here we note an interesting general feature of the Equal Weight View and how it makes for trouble in the present case. On the operative conception of peerhood, peers resemble each other in possessing a similar general competence for assessing relevant evidence and arguments. If you regard someone as incompetent compared to yourself with respect to his or her ability to assess relevant considerations, then you do not regard that person as your peer. (As a relatively extreme case, we might think here of the relationship that the qualified teacher of philosophy stands in to those of her students who have not yet developed any sophistication in evaluating arguments.) Of course, in order to respond correctly to one's evidence on a given occasion, it is not sufficient that one is competent to do so; one must actually manifest one's competence. Even against a general background of competence, one might still over- or underestimate one's evidence on a given occasion: one commits a performance error, as it were. Notice that it is characteristic of the Equal Weight View to credit the views of others in proportion to their general competence while abstracting away from facts about actual performance. What it is reasonable to believe in cases of peer disagreement is determined by giving equal weight to the opinions of the peers; crucially, in this calculation, the opinions that have been arrived at via the commission of performance errors will count for just as much as those opinions that are appropriate responses to the shared evidence.²⁰ Bare truths about who has in fact manifested their underlying competence and who **(p.132)** has not make no difference in cases of peer disagreement. However, once facts about general competence are privileged in this way in multi-person cases, it seems arbitrary and unmotivated to continue to maintain that actual performance makes a significant difference in single person cases (that is, cases in which a single individual arrives at an opinion on the basis of the non-psychological evidence that he possesses). Rather, on the suggested picture, if I am generally competent in the way that I respond to evidence (and I know that I am), then this should be enough to guarantee that I am reasonable in responding to my evidence in whatever way that I do. But this contradicts our initial assumption—namely, that one way of ending up with an unreasonable belief is to respond incorrectly to one's evidence, despite possessing the ability to respond to that evidence correctly.

3.4. The Litmus Paper Objection

Let us set aside, for the moment, the special case of disagreement among peers, and reflect on a much more general question: in what circumstances does it make sense for me to treat the fact that someone else believes as she does as evidence for the truth of that which she believes? A true (although perhaps not especially informative) answer: exactly when I take her belief to be a reliable indication of how things stand in the relevant part of reality. Thus, suppose that I know, on the basis of extensive past experience, that, when my weather forecaster judges that it will rain the next day, it tends to rain 80 percent of the time. In that case, I will treat her judgments to the effect that it will rain as evidence that it will rain, inasmuch as I take there to be a positive correlation between the two. Notice that, in this respect, there is absolutely nothing special about the way in which the judgments of another person come to count as evidence. Compare: I treat the fact that the litmus paper turns red as evidence that the liquid in which it is immersed is an acid, because, on the theories that I accept, the former is a reliable indication of the latter. This seems perfectly parallel to the reason why I treat the fact that my weather forecaster expects it to rain tomorrow as evidence that it will rain tomorrow. In general, the way in which the judgments of some other mind come to play the role of evidence does not differ from the way in which other states of the world do.

I believe that this observation, while elementary, is already enough to cast significant doubt on the Equal Weight View. For consider your perspective as one attempting to determine what to believe about some proposition. **(p.133)** You carefully survey what you take to be your evidence: various states of the world, the obtaining of which you take to provide clues as to whether the proposition is true or false. Some of these states of the world are bits of psychological reality, the beliefs of others—that Smith is highly confident that the proposition is true, that Jones is less so, and so on. Others of these states of the world are bits of non-psychological reality—for example, the fact that the litmus paper turned a given color in such-and-such circumstances. Insofar as you think it relatively unlikely that some part of psychological reality would be as it is unless the proposition were true, you regard the fact that things are arranged thus and so as evidence that speaks in favor of the proposition. But by the same token, insofar as you think it relatively unlikely that some piece of non-psychological reality would be as it is unless the proposition were true, you regard the fact that things are arranged that way as evidence that speaks in favor of the proposition. Now consider the special case in which you possess a considerable amount of non-psychological evidence, but where your psychological evidence is exhausted by the fact that (i) you yourself are confident that the proposition is true, and (ii) some peer is equally confident that the proposition is false. Again, on the Equal Weight View, you should split the difference with your peer and retreat to a state of agnosticism; in effect, one ought to give no weight to the non-psychological evidence in the presence of the psychological evidence. But what could be the rationale for such a policy of invidious discrimination? Why should the psychological evidence count for everything, and the non-psychological evidence for nothing, *given that the way in which the two kinds of evidence qualify as such is exactly the same?*

The idea that the non-psychological evidence is wholly swamped by the psychological evidence will look increasingly implausible the more the former is made to resemble the latter. Suppose that, in addition to my keen interest in having accurate beliefs about the future states of the weather, I am also a skilled engineer; I thus set out to construct a machine that will be of use in weather forecasting. My intention in constructing the machine is that its states will track future states of the weather: ideally, the machine will be in State A on a given day if and only if it will rain at some point during the following day, and so on. Unfortunately, the machine turns out to be a highly imperfect indicator of impending rain: I know, on the basis of significant past experience, that, 70 percent of the time that the machine enters into State A, it rains the next day, while 30 percent of the time it **(p.134)** does not. Suppose further that this is roughly how reliable *I* am when it comes to predicting rain, and, therefore, how reliable my peers are (and that I am aware of this). Given the characteristic tendency of proponents of the Equal Weight View to privilege the psychological evidence over the non-psychological evidence, it seems that they are committed to holding that there are circumstances in which I should treat the fact that a peer believes that it will rain tomorrow differently from the way I should treat the fact that the machine is in State A. (For example, the latter piece of evidence is susceptible to being swamped when enough of the other non-psychological evidence points in the opposite direction, while the former is not.) But this is odd, given that the way in which my peer's opinion is related to future states of the weather and the way in which the state of the machine is related to future states of the weather would seem to be very much the same.²¹

Perhaps the Equal Weight Theorist will insist instead that I should treat the machine as a peer with respect to weather forecasting. On this line of thought, in building the machine I am in effect constructing a robot peer; the machine's believing that it will rain tomorrow consists in its being in State A. However, this maneuver seems somewhat desperate and will not always be available. (The fact that the litmus paper turns red does not constitute its coming to believe that the liquid in which it is immersed is an acid.)

Although I will offer a final argument against the Equal Weight View in the context of discussing the views of Elga (2007),²² this concludes my primary brief against it. Even if one thinks that there is some merit to the arguments of this section, however, one might still be reluctant to abandon the Equal Weight View. For one might think that it is strongly supported by certain kinds of cases—for example, cases involving conflicting perceptual judgments. I will address such concerns directly in Section 5 below. To anticipate: while I agree that the Equal Weight View returns the correct verdicts about such cases, I do not believe that this fact is ultimately of much dialectical significance. The reason for this is straightforward: what I take to be the best alternative to the Equal Weight View yields the same **(p.135)** verdicts when applied to the cases in question. Before explicitly addressing those considerations that seem to favor the Equal Weight View then, I will put my preferred alternative on the table.

4. The Total Evidence View

Against the Equal Weight View, I have argued that there is at least one type of situation in which one is not required to split the difference with a peer who disagrees. Specifically, if you and I have arrived at our opinions in response to a substantial body of evidence, and your opinion is a reasonable response to that evidence while mine is not, then you are not required to give equal weight to my opinion and to your own. Indeed, one might wonder whether you are required to give *any* weight to my opinion in such circumstances. Suppose that, when we first discover that we hold different opinions about the issue in question, neither one of us is aware of why the other believes as he does. Naturally enough, you inquire as to why I believe as I do: because you regard me as in general a competent evaluator of evidence, you suspect that the fact that we believe differently about the issue is underwritten by my having access to some crucial piece of evidence to which you lack access (or vice versa). However, this suspicion turns out to be disappointed: I myself insist that I have no additional evidence, apart from considerations of which you yourself were already aware. In arriving at your opinion about the issue, you have thus taken into account every consideration that I would cite as a reason to justify my own opinion. Given this, why would you be rationally required to give some additional weight to my opinion (an opinion that is in fact unreasonable) and move from your own, an opinion that is in fact a perfectly reasonable response to the only evidence that either one of us claims to have?

Recall from above

The No Independent Weight View. In some cases of peer disagreement, one might be perfectly reasonable even if one gives no weight at all to the opinion of one's peer.

and

The Symmetrical No Independent Weight View. In some cases of peer disagreement, both parties to the dispute might be perfectly **(p. 136)** reasonable even if neither gives any weight at all to the opinion of the other party.

In Section 2, I tentatively suggested that these views might be defended by appeal to the apparent possibility of ‘reasonable disagreements:’ cases in which there is a range of rationally permissible attitudes to take toward some proposition given one’s evidence. However, for the sake of argument, we are currently supposing that the Uniqueness Thesis is true, and, thus, that the apparent possibility of reasonable disagreements is merely apparent: in fact, whenever two individuals believe different things on the basis of the same evidence, at least one of the two is being less than fully reasonable. Assuming that the Uniqueness Thesis is true, the Symmetrical No Independent Weight View is false. However, even if the Symmetrical No Independent Weight View is false, the No Independent Weight View might still be true. For even if it cannot be reasonable for both you and me to give no weight to the other’s opinion, perhaps it is nevertheless reasonable for you to give no weight to my opinion in the envisaged circumstances. As formulated above, the No Independent Weight View states that it might be perfectly reasonable to give no weight to the opinion of one’s peer ‘in some cases.’ We have now arrived at a proposal for what the relevant class of cases is—namely, the class of cases in which one’s original opinion correctly reflects the evidence that one shares with one’s peer but his opinion does not. Consider then:

The Asymmetrical No Independent Weight View. In cases of peer disagreement, it is reasonable to give no weight to the opinion of a peer as long as one’s own opinion is the reasonable response to the original evidence.

On this view, if either of the two peers engaged in a disagreement has in fact evaluated her shared evidence correctly, then that peer should stick to her guns, and the other peer should convert, since the opinion in question is the one that is in fact best supported by their evidence.

However, the Asymmetrical No Independent Weight View is false. Even if one responds to the original evidence in an impeccable manner and one’s peer does not, the fact that one’s peer responds as he does will typically make it rationally incumbent upon one to move at least some way **(p.137)** in his direction. First let us satisfy ourselves that this is so; we will then inquire as to why it is so.

Consider:

Case 6. You are a professional mathematician. Within the mathematics community, there is substantial and long-standing interest in a certain mathematical conjecture. (Call it **the Conjecture**.) If forced to guess, some members of the community would guess that the Conjecture is true, others that it is false; all agree that there is no basis that would justify a firm opinion one way or the other. Then, one day, the unexpected happens: alone in your study, you succeed in proving the Conjecture. On the basis of your proof, you become extremely confident, indeed practically certain, that the Conjecture is true. Because your high degree of confidence is based on a genuine proof that you correctly recognize as such, it is fully justified. Later, you show the proof to a colleague whose judgment you respect. Much to your surprise, the colleague, after examining the proof with great care, declares that it is unsound. Subsequently, you show the proof to another colleague, and then to a third, and then to a fourth. You approach the colleagues independently and take pains to ensure that they are not influenced by one another in arriving at their judgments about the status of your proof. In each case, however, the judgment is the same: the proof is unsound. Ultimately, your proof convinces no one: the entire mathematical community is united in its conviction that it is unsound, and thus, that the status of the Conjecture remains very much an open question.

In the face of this consensus, it would be unreasonable for you to remain practically certain that the Conjecture is true. You should be less confident of the Conjecture after your proof has been deemed unsound by the mathematical community than you were immediately after you first proved the Conjecture, back when you were alone in your study. Of course, because the proof is in fact sound, the judgment of the community to the contrary is misleading evidence, evidence that points in the wrong direction. But misleading evidence is evidence nonetheless, and the acquisition of such evidence will typically make a difference to what it is reasonable for one to believe. Moreover, if you are rationally required to be less confident after *all* of your peers have disagreed with you, then it would seem that you are also required to be at least somewhat less **(p.138)** confident after even *one* of your peers disagrees with you. For suppose that it was rationally permissible to give zero weight to the opinion of the first colleague. In that case, you could have left her office as rationally confident as when you entered, in which case you would have been in the same state of practical certainty upon entering the office of the second colleague that you consulted. Indeed, in that case it seems that you might as well simply forget about the fact that the whole unpleasant business with the first colleague occurred at all before visiting the second colleague, in which case you would be in more or less exactly the same position upon entering the office of the second colleague. And, if it is rationally permissible to give zero weight to *his* opinion . . .

Moral: the fact that a peer believes differently can make it rationally incumbent upon you to change what you currently believe, even if, had the peer responded to the evidence in a reasonable manner, he too would believe exactly as you believe. One should give some weight to one's peer's opinion, even when from the God's-eye point of view one has evaluated the evidence correctly and he has not. But why? Exactly because one does not occupy the God's-eye point of view with respect to the question of who has evaluated the evidence correctly and who has not.²³ Typically, when one responds reasonably to a body of evidence, one is not utterly blind to the fact that one has done so; on the other hand, such facts are not perfectly transparent either. Even if one has in fact responded to the evidence impeccably on a given occasion, one might still have reason to doubt that one's performance was impeccable. Such a reason is provided when a peer responds to that same evidence differently. To give no weight to the fact that a peer responds to the evidence differently is in effect to treat it as certain that one's peer is the one who has misjudged the evidence. But it would be unreasonable to be certain of this, even when it is true.²⁴ **(p.139)**

Rationality consists in responding appropriately to one's evidence. But one's evidence includes evidence to the effect that one does not always respond appropriately to one's evidence (that is, evidence to the effect that one is fallible in responding appropriately to one's evidence), as well as evidence to the effect that one is more likely to have responded inappropriately when one finds oneself in certain circumstances. When one possesses higher-order evidence to the effect that one is currently in circumstances in which one is more likely than usual to have made a mistake in responding to one's first-order evidence, one has a reason to temper one's confidence—even if that confidence is in fact an impeccable response to the first-order evidence. When one finds oneself in the position of a minority of one in the way that one has responded to the evidence, one should temper one's confidence, for one now possesses higher-order evidence that suggests that the bearing of the original, first-order evidence is something other than what one initially took it to be. Moreover, this is so, even if the higher-order evidence is misleading, as when one has in fact responded appropriately to the first-order evidence and one's peers have not.

The tendency of higher-order considerations to affect how much confidence one should invest in the deliverances of one's first-order reasoning is a quite general phenomenon; it will be helpful briefly to explore the general phenomenon in order to gain some perspective on the case at hand. Consider first a case of practical deliberation: one is attempting to determine which of two attractive job offers to accept. Some considerations point in one direction; other considerations point in the opposite direction. One deliberates and concludes that, on balance, one's practical reasons favor accepting one of the two offers over the other. How confident should one be that this is what one should do? This can depend, not only on the strength of the reasons for and against accepting that offer but also on (what one knows about) the circumstances in which one is deliberating. If one knows that one is deliberating in circumstances in which one's ability to weigh practical reasons tends to be compromised (for example, one is **(p.140)** in a state of inebriation), then one should be less confident than if one is deliberating in more ideal circumstances. Moreover, this is so even if one has in fact deliberated impeccably, despite being inebriated, and has assigned the correct weights to all of the practical reasons that bear on the choice. Of course, higher-order considerations having to do with whether one is deliberating in the bad case (one is inebriated) or in the good case (one is sober) are not themselves reasons for or against accepting either job offer in the way that the practical reasons that bear directly on the choice are; nevertheless, they too make a difference to what it is appropriate to conclude on the basis of one's deliberations. On the present picture, even if one's practical reasoning in the good case is identical to one's practical reasoning in the bad case, what one is justified in concluding on the basis of that reasoning might differ.

The same point holds for theoretical reasoning. I am entitled to have more confidence in the conclusion of a given piece of mathematical reasoning when I have performed the relevant calculation in the morning when I am wide awake than if I arrive at the same conclusion by reasoning in exactly the same way late in the evening, when I know that I am prone to making mistakes because of fatigue.

On the present view, cases in which one in fact responds impeccably to one's evidence, but one's peer responds inappropriately, are much like cases in which one engages in a flawless piece of practical reasoning despite being inebriated. The fact that a peer has responded to the evidence differently should lead one to temper one's confidence in one's own response, just as the fact that one is inebriated should lead one to temper one's confidence in the conclusion of one's practical reasoning, despite the actual flawlessness of one's performance. Again, in both cases, it is the fact that the status of one's performance is not perfectly transparent that opens the door for higher-order considerations to make a difference.

Of course, to acknowledge that higher-order considerations make *some* difference is not to fall back into the mistake of thinking that they make *all* the difference. After all, even when one's current level of inebriation makes it significantly more likely that one will over- or underestimate the strength of one's practical reasons (and one knows that this is so), one can still make more or less rational decisions, and the status of a given decision will typically depend a great deal on the overall disposition of those practical reasons. Similarly for the theoretical case: although you should be **(p.141)** somewhat less confident that the Conjecture is true upon finding that a colleague remains unconvinced despite having been presented with your proof, it is a mistake to think that at that point the only evidence that makes a difference is the respective psychological reactions of you and your colleague. When one possesses what is in fact a genuine proof that one correctly recognizes as such, one possesses an extremely strong piece of evidence. (Indeed, it would perhaps be difficult to imagine a stronger single piece of evidence for anything.) The justification afforded by such a piece of evidence has a certain robustness in the face of challenge: it is not easily washed away by the fact that another mistakenly fails to appreciate it on a given occasion. Of course, your colleague might feel just as confident that your proof is unsound as you feel that it is sound. Indeed, all of the psychological accompaniments of the two judgments might be the same. But, in any case, we have independent reason to be skeptical of the idea that phenomenology is that on which epistemic status supervenes. In general, when one reasons badly, one's phenomenology might be indistinguishable from one's phenomenology when one reasons impeccably (in both cases, one has the same feelings of subjective certainty, and so on). We should not thereby be driven to the conclusion that the deliverances of good reasoning and bad reasoning have the same epistemic status.²⁵

Where does this leave us?

In the previous section, I argued that, in cases of peer disagreement, getting the original, first-order evidence right typically counts for *something* (pace the Equal Weight View). In this section, I have argued that doing so does not count for *everything* (pace the No Independent Weight View). Indeed, from the present perspective, there is a sense in which the Equal Weight View and the No Independent Weight View both suffer from the same fault: they embody overly simple models of how one's first-order evidence and one's higher-order evidence interact in determining facts about what it is reasonable to believe all things considered. On the Equal Weight View, what it is reasonable to believe in cases of peer disagreement in effect supervenes on facts about the distribution of peer opinion. On the No Independent Weight View, what it is reasonable to believe in such cases supervenes on facts about the first-order evidence possessed by the peers. **(p.142)** On the present view, both of these supervenience claims are false: neither class of facts suffices on its own to fix the facts about what it is reasonable to believe. Rather, what it is reasonable to believe depends on both the original, first-order evidence as well as on the higher-order evidence that is afforded by the fact that one's peers believe as they do. For this reason, it seems appropriate to call the view on offer **the Total Evidence View**.

Even if both the Equal Weight View and the No Independent Weight View are unsatisfactory, we might still wonder: which is closer to the truth? Granted that, on the Total Evidence View, both the first-order evidence and the higher-order evidence count for something, which kind of evidence plays a greater role in fixing facts about what it is reasonable to believe?

It is a mistake, I believe, to think that there is some general answer to this question. In some cases, the first-order evidence might be extremely substantial compared to the higher-order evidence; in such cases, the former tends to swamp the latter. In other cases, the first-order evidence might be quite insubstantial compared to the higher-order evidence; in such cases, the latter tends to swamp the former. (We will consider plausible examples of each of these types of case below.) In still other cases, the two kinds of evidence might play a more or less equal role in fixing facts about what it is reasonable to believe. So the question of which counts for *more*—peer opinion, or the evidence on which the peers base their opinion?—is not, I think, a good question when it is posed at such a high level of abstraction.

Nevertheless, we can offer some general observations that bear on this issue here. Consider again the kind of case that we have employed in attempting to undermine the Equal Weight View: initially, you and I have access to the same substantial body of evidence *E*, evidence that in fact strongly favors *H* over not-*H*; you respond reasonably and so are quite confident that *H* is true; I, on the other hand, respond unreasonably and am equally confident that *H* is false. Once we compare notes, our new total evidence consists of *E**:

- (i) Our original evidence *E*
- (ii) The fact that you are quite confident that *H* is true
- (iii) The fact that I am quite confident that *H* is false

What is it reasonable for us to believe about *H* on total evidence *E**? Given that you and I are peers, it is plausible to suppose that the two **(p.143)** pieces of higher-order psychological evidence ((ii) and (iii)) are more or less equally strong pieces of evidence that point in opposite directions. All else being equal, then, one would expect *E** to favor *H* over not-*H* inasmuch as it is composed of a substantial body of evidence that strongly favors *H* over not-*H*, supplemented by two additional pieces of evidence of approximately equal strength, one of which tends to confirm *H*, the other of which tends to disconfirm *H*.

Indeed, it is tempting to think that, if in fact our respective psychological reactions count as more or less equally strong pieces of evidence that point in opposite directions, then they in effect cancel each other out and leave what it is reasonable for us to believe unchanged. According to this line of thought, what it is reasonable for us to believe about *H* on *E** is identical to whatever it was reasonable for us to believe about *H* on *E*, inasmuch as the net effect of adding the two new pieces of evidence comes to zero. Here the Asymmetrical No Independent Weight View threatens to return via the back door, at least in a special class of cases—namely, cases in which peer opinion is evenly divided. For in such cases, the evidence afforded by peer opinion is perfectly counterbalanced.

However, this tempting line of thought is mistaken. The addition of the counterbalanced psychological evidence *does* make a difference to what it is reasonable for us to believe. For, once the counterbalanced evidence is added to our original evidence, a greater proportion of our total evidence supports an attitude of agnosticism than was previously the case; the evidence available to us now is on the whole less supportive of *H* than before. The addition of (ii) and (iii) thus has a moderating impact and tends to push what it is reasonable for us to believe about the hypothesis in the direction of agnosticism. Therefore, given that *E* is a substantial body of evidence that strongly favors *H* over not-*H*, we would expect that *E** will also favor *H* over not-*H*, although not to as a great a degree as *E* does. (That is, all else being equal, the reasonable level of confidence to have in hypothesis *H* on evidence *E** will be greater than 0.5 but less than whatever it was reasonable to have on evidence *E*.)

Significantly, the point generalizes beyond the two-person case. As more and more peers weigh in on a given issue, the proportion of the total evidence that consists of higher-order psychological evidence increases, and the proportion of the total evidence that consists of first-order evidence decreases. As the number of peers increases, peer opinion counts for **(p.144)** progressively more in determining what it is reasonable for the peers to believe, and first-order considerations count for less and less. At some point, when the number of peers grows large enough, the higher-order psychological evidence will swamp the first-order evidence into virtual insignificance. In such cases, the Total Evidence View becomes more or less extensionally equivalent to the Equal Weight View with respect to what it requires the peers to believe. Moreover, this holds regardless of the particular way in which opinion is distributed among the peers. That is, it holds for cases in which peer opinion is evenly divided, for cases in which peer opinion is unanimous, as well as for intermediate cases.

Imagine an infinite number of peers confronted with a finite amount of evidence that bears on some issue. Each of the peers inspects the evidence and independently arrives at a view. When the peers compare notes, they find that opinion among them is perfectly divided: every peer on one side of the issue has one and only one counterpart on the other side. In these circumstances, the peers should suspend judgment about the issue, even if that response is not the most rational response to the original, first-order evidence. With respect to this case, the Equal Weight View returns the correct verdict from the perspective of one who holds the Total Evidence View. This is so *not* because the higher-order evidence trumps the first-order evidence in general, as the proponent of the Equal Weight View maintains. Rather, it is because, in sufficiently extreme cases, the higher-order psychological evidence might be so substantial compared to the first-order non-psychological evidence that the former in effect swamps the latter into virtual insignificance.

The same holds true for cases in which the peers find that they agree. Above, we looked askance at the idea that two peers, both of whom irrationally hold some view that is not in fact supported by their evidence, might bootstrap their way into rationally holding that view simply by encountering one another and comparing notes. Indeed, we took the fact that the Equal Weight View licenses such two-person bootstrapping as a consideration that counts against it (see Section 3.2 above). However, as the number of generally reliable peers who independently respond to their evidence in the same mistaken manner increases, such bootstrapping seems less and less objectionable. At some point, it becomes, I believe, unobjectionable. If I hold some belief on the basis of fallacious reasoning, then it will typically not be reasonable for me to hold that belief. However, **(p.145)** in the unlikely but possible situation in which a large number of generally reliable peers mistakenly arrive at the same conclusion by independently committing the same fallacy, it will typically be reasonable for them to believe that conclusion upon comparing notes, even if there is no legitimate first-order reasoning by which they could have arrived at the conclusion. Again, in this case the Equal Weight View yields the correct verdict from the perspective of the Total Evidence View. As before, this is not due to some general tendency of higher-order evidence to trump first-order evidence. Rather, it is due to the fact that, in this case, the higher-order evidence that has been amassed is sufficiently substantial compared to the first-order evidence that it effectively determines the bearing of the overall evidence.

Does this in effect give the game away to someone who takes the diversity of opinion with respect to various controversial issues to mandate an attitude of agnosticism about those issues? That is, even if the Equal Weight View is false and the Total Evidence View is true, will not all of the interesting/threatening/radical consequences that seemed to follow from the Equal Weight View still be true, at least if one is sufficiently generous in attributing the status of 'peer' to other people? Is not agnosticism the only reasonable stance to take toward all of those controversial issues on which peer opinion is heavily divided, as the proponent of the Equal Weight View has insisted all along?

Consider also those philosophical questions with respect to which there is consensus, or near consensus. Suppose, plausibly, that there are very few if any genuine skeptics about other minds: informed philosophical opinion is (close to) unanimous in holding that one is typically in a position to know that there are minds other than one's own. In Kelly (2005a), I took a dim view of the suggestion that this fact would suffice to make it unreasonable to embrace skepticism about other minds: rather, whether it is reasonable or unreasonable to embrace skepticism about other minds is primarily a matter of the quality of the first-order arguments for and against such skepticism, arguments that do not make reference to empirical, sociological facts about the number of skeptics and non-skeptics. However, in light of the present view, a reversal of this judgment might seem to be in order. Could it really be that the unreasonableness of skepticism about other minds consists in the *unpopularity* of such skepticism among the relevant class of people? **(p.146)**

Before acquiescing in this line of thought, we should note an important element of idealization in our discussion to this point, an element that looms large in the present context. Throughout, we have been concerned with the probative force of peer opinion in cases in which the peers arrive at their opinions *independently* of one another. This assumption of independence tends to maximize the probative force of peer opinion relative to the probative force of first-order evidence. Impressive evidence that a given answer to a question is the correct answer is afforded when a large number of generally reliable peers independently converge on that answer. On the other hand, the less their convergence is an independent matter, the less weight such convergence possesses as evidence.²⁶ Similarly, evidence that strongly favored agnosticism with respect to some question would be a more or less even distribution of opinion among a substantial number of peers, where each of the peers has arrived at his or her own opinion independently of the others. Again, the less such independence is present, the weaker the higher-order evidence will be relative to the first-order evidence.

Consider, as an especially extreme illustration of the importance of independence, the venerable 'Common Consent' Argument for the Existence of God. In its simplest and most straightforward form, the argument runs as follows:

- (Premise) Everyone believes that God exists.
- (Conclusion) Therefore, God exists.

(In a slightly less crude form, the premise of the argument is that *almost* everyone, or the great majority of humankind, believes that God exists.²⁷) **(p.147)**

As arguments go, the Common Consent Argument for the Existence of God is not exactly an overwhelming one, possessing as it does the twin defects of transparent invalidity and the having of an obviously false claim as its sole premise. Nevertheless, even though *God exists* does not follow from *Everyone believes that God exists*, we can ask: if it were true that everyone, or almost everyone, believed that God exists, how much support would that lend (if any) to the proposition that God exists?

This is a complicated question about which much could be said; here we note the following. Whatever evidence is afforded for a given claim by the fact that several billion people confidently believe that that claim is true, that evidence is less impressive to the extent that the individuals in question have not arrived at that belief independently. That is, the evidence provided by the fact that a large number of individuals hold a belief in common is weaker to the extent that the individuals who share that belief do so because they have influenced one another, or because they have been influenced by common sources. (I assume that both of these conditions play a large role in the case of religious belief.) In principle, the fact that a small handful of people arrive at the same belief independently of one another might be better evidence that that belief is true than if many millions of people arrive at the same belief non-independently. The intellectual case for Islam would not be any stronger today if birth rates in Muslim countries had been twice as high in past decades as they actually were; nor would the case be any weaker if such birth rates had been significantly lower.

The same holds for cases in which there is widespread disagreement but where the members of the contending factions have not arrived at their opinions independently. In an interesting recent essay, G. A. Cohen (2000) notes that the Oxford-trained philosophers of his generation are almost unanimously of the opinion that there is a philosophically important distinction between analytic and synthetic truths. But, on the other hand,

people of my generation who studied philosophy at Harvard rather than at Oxford for the most part *reject* the analytic/synthetic distinction. And I can't believe that this is an accident. That is, I can't believe that Harvard just *happened* to be a place where both its leading thinker rejected that distinction and its graduate students, for independent reasons—merely, for example, in the independent light of reason itself—also came to reject it. And vice versa, of course, for Oxford. I believe, rather, that in each case students were especially impressed by the reasons **(p.148)** respectively for and against believing in the distinction, because in each case the reasons came with all the added persuasiveness of personal presentation, personal relationship, and so forth. (Cohen 2000: 18; emphasis in original)

Consider Cohen's position as one attempting to determine what to believe about this issue. On the one hand, there are the first-order considerations that have been offered for and against the existence of a philosophically significant analytic/synthetic distinction. In addition, Cohen is also aware of the views of other individuals who are similarly acquainted with those first-order considerations and whom he regards as his peers in other relevant respects. In weighing evidence of the latter kind, Cohen should sharply discount for the fact that (as he sees it) many individuals on both sides of the issue hold the views that they do because those views were held by their teachers. That is, in the counterfactual situation in which the distribution of peer opinion is exactly as it is, but in which each of the peers arrived at his or her view in response to 'the independent light of reason itself,' the higher-order evidence possessed by Cohen would be much more substantial than it is as things actually stand. The point is not that individuals who believe what their teachers believe are less reliable than they would be if they made up their own minds. Indeed, as a general matter, this is not even true. (If your teacher is better at assessing the arguments than you are, then you will be more reliable if you simply believe as she does than if you arrive at a view on the basis of your own assessment of the arguments.) The point, rather, is that, insofar as one believes as one does because this is what one's teacher believes, the fact that one believes as one does is not an *additional* piece of psychological evidence, over and above the psychological evidence afforded by the teacher's belief.

The general moral: even in cases in which opinion is sharply divided among a large number of generally reliable individuals, it would be a mistake to be impressed by the sheer number of such individuals on both sides of the issue. For numbers mean little in the absence of independence. (It is, of course, an empirical question—one that belongs, presumably, to psychology and sociology—how independently people arrive at their views about various issues.) If one uncritically assumes that the members of the contending factions have arrived at their views independently, then one will tend to overestimate the importance of other people's opinions as evidence and underestimate the importance of the first-order evidence (**p.149**) and arguments. One will be too quick to conclude that agnosticism is the reasonable stance in cases in which opinion is sharply divided, and too quick to conclude that deference to the majority is the reasonable course in cases in which opinion is not sharply divided.²⁸

Nevertheless, it is true that, on the Total Evidence View, there will be possible cases in which the higher-order evidence is sufficiently substantial compared to the first-order evidence that the latter counts for (almost) nothing. By the same token, however, there will be possible cases in which the opposite is true. What is a case in which peer opinion effectively counts for nothing in virtue of being overwhelmed by the first-order considerations? Consider a case discussed by both Christensen (2007: 199–203) and Elga (2007: 490–1). You and I go to dinner with several friends; at the end of the meal we independently calculate what an individual share of the total bill comes to (imagine that the group has agreed to split the bill evenly among its members). You judge that an individual share is \$43 per person, a perfectly plausible (and, let us suppose, correct) answer to the question of what each of us owes. I, however, arrive at an absurd answer of \$450, an amount that significantly surpasses the total bill. Both Christensen and Elga think that, in *these* circumstances, you are *not* required to treat my answer and your answer with equal respect; indeed they think that you are entitled more or less to dismiss my answer entirely. The difficulty is how to account for this on a picture according to which splitting the difference is typically the appropriate response to peer disagreement. In general, it is at least a *prima facie* embarrassment for the Equal Weight View that the following is possible: a person for whom one has arbitrarily strong evidence that he or she is a peer might nevertheless give a patently absurd answer on a given occasion. For it seems incredible that, in such circumstances, one would be unreasonable if one failed to treat the peer's patently absurd answer and one's own non-absurd answer even-handedly.

Unsurprisingly, both Christensen and Elga have interesting and detailed stories to tell about why, in these but not in otherwise similar cases, one **(p.150)** need not give any weight to the view of one's peer.²⁹ We will not pause to evaluate the specifics of their respective proposals; here we note only how the Total Evidence View offers an extremely straightforward and compelling explanation of why you are entitled effectively to discount my absurd opinion. Quite simply: given the totality of considerations available to you that bear on the question at issue (for example, your knowledge that the total bill is n , a number that is less than \$450), it would be completely unreasonable for you to give any significant credence to the proposition that a share of the total bill is \$450, despite the fact that this is what I, your peer, believe. In this case, it is the non-psychological considerations that swamp the psychological considerations into epistemic insignificance.

5. Considerations that Seem to Favor the Equal Weight View

5.1. Perceptual Judgments

As mentioned above, I believe that much of the appeal of the Equal Weight View derives from reflection on certain kinds of examples. In particular, the Equal Weight View can seem almost obviously or trivially correct when one reflects upon examples involving the conflicting perceptual judgments of individuals equally well suited to make those judgments. Recall Case 1 from above: you and I, two equally attentive and well-sighted individuals, watch the horses cross the finish line from equally good vantage points. It looks to me as though Horse A finishes slightly ahead of Horse B, while it looks to you as though Horse B finishes slightly ahead of Horse A. The intuitive verdict: once we find that our initial judgments conflict, the uniquely reasonable course is for us to split the difference and retreat to a state of agnosticism about which of the two horses actually won the race.

I do not contest the intuitive verdict; indeed, I take it to be correct. What I do contest is the idea that the intuitive verdict has any tendency to support the Equal Weight View over the Total Evidence View. For, when the Total Evidence View is correctly applied to Case 1, it too returns **(p.151)** the intuitively correct verdict that you and I should abandon our original opinions and retreat to a state of agnosticism.

First, note that there are at least some cases in which the Total Evidence View will rationally require two individuals who begin with conflicting opinions to adopt a new opinion that is perfectly intermediate between their original opinions. Here is one such case:

Case 7. At time t_0 , you and I possess different evidence which bears on some hypothesis H . Your evidence suggests that H is true; my evidence suggests that it is false. Moreover, each of us responds to his evidence in a reasonable manner: you believe that H is true while I believe that it is false. At time t_1 , we encounter one another and pool our evidence. After doing so, our new total evidence does not favor H over not- H ; nor does it favor not- H over H .

Given that the total evidence available to us at time t_1 favors neither alternative over the other, an advocate of the Total Evidence View will maintain that we should suspend judgment. You should abandon your belief that the hypothesis is true while I should abandon my belief that it is false. In the light of our new total evidence, we should converge on the point that is intermediate between our original opinions. With respect to Case 7, then, the Total Evidence View will require us to respond in a way that is extensionally equivalent to the way that we would respond if we were both following a norm of 'Split the Difference.'

Notice, however, that Case 7 is simply Case 1, abstractly described. As you and I watch the horses cross the finish line, it appears to me as though Horse A finishes just ahead of Horse B. To the extent that I have evidence for my judgment that Horse A finished ahead of Horse B, that evidence consists of my perceptual evidence: the fact that it *looks* or *appears* to me that Horse A finishes ahead, or that my visual experience represents Horse A as having finished ahead. In the absence of other evidence that bears on the question, it is at that point reasonable for me to believe that Horse A finished ahead of Horse B, since this is what my total evidence supports. Similarly, your initial judgment that Horse B finished just ahead of Horse A is a reasonable response to the evidence that you possess at time t_0 —namely, the fact that it looked or seemed to you as though Horse B finished just ahead of Horse A. At time t_1 , we compare notes: you learn that I think that Horse A won because that is how it looked to me; I **(p. 152)** learn that you think that Horse B won because that is how it looked to you. At this point, the total evidence that is available to each of us has changed in a rather dramatic way: I have gained evidence that suggests that Horse B won the race, while you have gained evidence that Horse A won the race. Moreover, given the relevant background assumptions and symmetries, it is natural to think that the total evidence that we now share favors neither the proposition that Horse A finished ahead of Horse B nor the proposition that Horse B finished ahead of Horse A. Thus, given our new total evidence, you and I should abandon our initial opinions about which horse won the race. The Total Evidence View, no less than the Equal Weight View, requires us to suspend judgment and retreat to a state of agnosticism in Case 1 and in cases of relevantly similar structure. Thus, it is a mistake to think that such cases favor the Equal Weight View over the Total Evidence View.³⁰

5.2. The Analogy with Inanimate Measuring Devices

Recall Case 2 from above: you and I arrive at different views about the temperature by consulting our hitherto equally reliable thermometers; subsequently, we discover that our thermometers disagree. The intuitive verdict: we should abandon our original opinions about the temperature. In particular, it would be patently unreasonable for me to retain my original belief simply because that was what *my* thermometer indicated about the temperature. Indeed, I should give no more credence to what my thermometer says about the temperature than to what yours says, and vice versa. But what holds for the conflicting readings of equally reliable thermometers holds also for the conflicting judgments of individuals who are peers. **(p.153)**

Response: I embrace the intuitive verdict about Case 2 but deny that this has any tendency to support the Equal Weight View over the Total Evidence View. Of course, the mere fact that a given thermometer is *mine* is no reason for me to think that it is more trustworthy than your thermometer, or for me to favor its indications over the indications of yours. But, similarly, it is no part of the Total Evidence View that it is permissible for me to favor my original opinion over yours simply because the opinion in question belongs to me. On the Total Evidence View, once I discover that our original opinions differ, it might very well be reasonable for me to adopt an opinion that is closer to my original opinion than to yours. But, if so, that is because the opinion in question is best supported by what is now our total evidence. (By the same token, it might also be reasonable for me to adopt a new opinion that is closer to your original opinion than to mine.) Again, in such cases, whether the opinion that it is ultimately reasonable for us to hold is closer to your original opinion or to mine will typically depend on which one of us (if either) did a better job of responding to the first-order evidence in arriving at his or her original opinion. When it is construed in this way, the thermometer analogy at best tells against the view that Elga (2007) refers to as ‘the Extra Weight View.’

The objector might insist that the thermometer analogy tells against the Total Evidence View as well. For consider: even if my thermometer is in fact functioning perfectly on a particular occasion and yours is malfunctioning, it would still be unreasonable for me to favor what my thermometer says as long as I am in no position to appreciate these facts. Similarly, even if I have in fact evaluated the evidence correctly in arriving at a given belief and you have not, it would be unreasonable for me to favor my view over yours if I am in no position to appreciate that this is so.

Response: first, note that, when the thermometer analogy is construed in this way, there are significant disanalogies between it and a case of peer disagreement. In the thermometer case, the two devices function as black boxes that simply output a given number; what takes place ‘underneath the hood’ (so to speak) is completely opaque to us. In terms of interpersonal disagreement, the closest analogue to this would be something like the following. I regard the two of us as generally reliable about some domain. I then discover that you hold a different view than I do about some issue in that domain, but I have no idea about how you arrived at your view or on what basis you currently hold it: for all I know, you might hold **(p.154)** the view on the basis of relevant considerations of which I am unaware, considerations that, if presented to me, would result in a change in my view. On the other hand, it is also possible, for all I know, that you are unaware of relevant considerations that I possess (considerations that, if presented to you, would result in a change in your view). To the extent that I am ignorant of how you arrived at your view, or why you currently hold it, it is, I think, reasonable for me to treat our views even-handedly until learning more. (And of course, a proponent of the Total Evidence View need not say otherwise.) In contrast, in cases of peer disagreement, the peers have access to the same body of evidence and are aware that this is so; as we have sometimes put it, they have ‘compared notes.’ Thus, suppose that I confidently hold a philosophical thesis on the basis of arguments and considerations that are inadequate to support it.³¹ I cite these arguments and considerations in defense of my thesis, attributing my confidence in the thesis to them. Recognizing my reasons as inadequate, you come away unimpressed. In these circumstances, there is a sense in which you have witnessed the malfunction occur. In terms of the thermometer analogy, it is as though one has the opportunity to open up the other person's thermometer (as well as one's own) and inspect how the temperature reading was arrived at. In cases of peer disagreement, one gets to go underneath the hood, as it were.³²

Still, one might think that this does not yet get to the heart of the matter. Even if the two peers are both fully aware of why the other believes as he or she does, and one of the two has in fact done a better job of evaluating their shared evidence, who is to say which of the two has done a better job? The peers themselves, one might think, are not in a position justifiably to make such judgments. Perhaps then the situation of the peers is analogous to the following: although one gets to open up the two thermometers and inspect how their conflicting readings are generated, one lacks the wherewithal reliably to discriminate a malfunctioning thermometer from one that is functioning properly. Similarly, a proponent of the Equal Weight View might claim that, even in a case in which one has in fact done a better job **(p.155)** of evaluating the relevant considerations than one's peer, one has no way of discriminating such a case from a case in which things are the other way around. Even if things are asymmetrical at the bottom level (one's belief reflects the evidence better than one's peer's belief; one's thermometer is in fact functioning better than the other person's thermometer), things are symmetrical one level up: one has no justification for thinking that one's belief better reflects the evidence, or that one's thermometer is the thermometer that is working properly. One is thus no more justified in thinking that one's own belief accurately reflects the evidence than one's peer is in thinking that *his* belief accurately reflects the evidence. Therefore, given the higher level normative symmetry, it would be unreasonable to favor one's own belief over the belief of one's peer. Favoring one's own belief would be reasonable only if one had some independent evidence that one's belief is more likely than one's peer's belief to be an accurate reflection of the evidence: for example, knowledge that one has outperformed the other person in relevant ways in the past. But, *ex hypothesi*, no such independent evidence is available in a case in which the disagreement is a disagreement between peers. Thus—so the argument runs—even if one is in fact the person who has better evaluated the evidence on this particular occasion, one has no justification for thinking that this is so. Hence, one should split the difference.³³

However, the proponent of the Equal Weight View is not entitled simply to assume that things *are* symmetrical between us at the higher level. That is, the proponent of the Equal Weight View is not entitled simply to assume that you and I are equally well justified in thinking that we have correctly responded to the evidence in a case in which you have done so and I have failed to do so. Of course, given that we are peers, neither of us possesses independent evidence that suggests that he is the one who has responded correctly. However, even in the absence of *independent* evidence, there is another possibility: namely, that, when one correctly responds to a body of evidence, one is typically better justified in thinking that one has **(p.156)** responded correctly than one is when one responds incorrectly. It is this thought that I will now explore and defend.

First, let us observe the unobvious point that, when one responds correctly to a body of evidence, one typically has some justification for thinking that one has responded correctly. In paradigmatic cases in which one takes up the view that is best supported by one's evidence, it is no mere accident that one has done so (although lucky accidents are of course possible, they are atypical). Rather, one takes up the belief in question precisely because it is supported by one's evidence. Indeed, in a given case, one might very well take up the belief because one *recognizes* that this is what one's evidence supports. Plausibly, *recognizing that p entails knowing that p* . Assuming that that is so, then any case in which one recognizes that one's evidence supports a given belief is a case in which one knows that one's evidence supports that belief. Clearly, if one knows that one's evidence supports a given belief, then one is *justified* in thinking that one's evidence supports that belief; if one were not justified, one would not know. But, even if recognizing that p does *not* entail knowing that p , one would in any case not be able to recognize that p if one were unjustified in thinking that p . It follows immediately from this that, whenever one recognizes that one's evidence supports such-and-such a conclusion, one is justified in thinking that one's evidence supports that conclusion.

Thus, in any case in which you hold a given belief because you recognize that this is what your evidence supports, not only is your belief a reasonable one, but you are also justified in believing an epistemic proposition to the effect that it is reasonable for you to hold that belief. Indeed, given that *recognizing that p* entails *being justified in believing that p*, you will be justified in believing that the evidence supports your view for as long as you continue to recognize that it does. Consider then a case in which you take up a belief in virtue of recognizing that that belief is what the evidence supports; I am not yet on the scene. Subsequently, you discover that I believe otherwise despite having been exposed to the same evidence. Presumably, the proponent of the Equal Weight View will claim that, once I arrive on the scene with my conflicting belief, you are no longer justified in believing that the original evidence supports the relevant proposition because you are now no longer in a position to recognize that it does: in effect, encountering a peer who thinks otherwise strips one of one's prior ability to recognize the bearing of the first-order evidence for what **(p. 157)** it is. This suggestion is perhaps not wholly implausible: the proponent of the Equal Weight View envisions the situation as one in which one loses knowledge that one previously possessed in virtue of acquiring misleading evidence. In general, this is perfectly possible. However, in order to preserve the higher-level normative symmetry that lends the Equal Weight View its plausibility, it seems that the proponent of the Equal Weight View will have to maintain the following: when we met, the justification that I possessed for thinking that my original response to the evidence is reasonable was just as strong as the justification that you possessed for thinking that your original response to the evidence is reasonable. For, if you were better justified in thinking that your response was reasonable than I was in thinking that my response was reasonable, then this would break the putative higher-level symmetry and provide a basis for favoring your original belief over mine. (Compare a situation in which you are better justified in thinking that your thermometer is functioning properly than I am in thinking that my thermometer is functioning properly.) Of course, because you originally recognized that the evidence supported your belief, there is, on the assumption that the Uniqueness Thesis is true, no possibility that I similarly recognized that the evidence supported my belief. At best, I mistakenly took the evidence to support my belief when it did not. In order to preserve normative symmetry at the higher level then, the proponent of the Equal Weight View will insist that:

When you correctly recognize that the evidence supports *p*, you are no more justified in thinking that the evidence supports *p* than I am in thinking that the evidence supports not-*p* when I mistakenly take the evidence to support not-*p*.

Presumably, the reference to two people is inessential here. So the proponent of the Equal Weight View will also endorse:

When you correctly recognize that the evidence supports p , you are no more justified in thinking that this is what the evidence supports than you would have been had you mistakenly taken the evidence to support $\text{not-}p$ instead.

But these assumptions are quite dubious. In any case, they are not ones that a proponent of the Total Evidence View need or should accept. On its most formidable construal, then, the argument from the analogy with inanimate **(p.158)** measuring devices depends on assumptions that are at best controversial and for which no argument has been provided.

5.3. Downward Epistemic Push

Although the most recent objection is ultimately not compelling, it proceeds from a genuine insight that is worth making fully explicit. The insight in question is one that we have already briefly touched on above. It might be put like this: in general, what it is reasonable to believe about the world on the basis of one's evidence is constrained by what it is reasonable for one to believe *about* one's evidence.³⁴ (Put otherwise: what it is reasonable for one to believe about the world is not wholly independent of what it is reasonable for one to believe about what it is reasonable for one to believe about the world.)

Because this insight might naturally be taken to support the Equal Weight View, let us explore it a bit further here. Once again, let E represent one's total first-order evidence with respect to H . Consider then the epistemic proposition:

E is good evidence for H .

In Section 4, it was argued that a proponent of the Total Evidence View should agree that higher-order evidence about the quality of one's first-order evidence typically makes some difference to what one should believe on the basis of that first-order evidence. Inasmuch as this is so, she will agree with the proponent of the Equal Weight View that any evidence that bears on this epistemic proposition is also evidence that bears on H itself. In general, how confident one should be that H is true is tied to how confident one should be that the corresponding epistemic proposition is true. When one acquires reasons to increase one's confidence in the epistemic proposition, one acquires reasons to increase one's confidence in H ; on the other hand, when one acquires reasons to decrease one's confidence in the epistemic proposition, one acquires reasons to decrease one's confidence in H . That this relationship exists might seem strongly to favor the Equal Weight View, or, more generally, any view on which **(p.159)** higher-order evidence trumps lower-order evidence. For, even in a case in which E genuinely supports H , one's justification for believing H on the basis of E will tend to be undermined by evidence against the epistemic proposition that *E is good evidence for H* . And one's justification for believing this epistemic proposition would seem to be hostage to what one's peers think.

I think that this phenomenon of *downward epistemic push* is a genuine one. (If it were not, the Asymmetrical No Independent Weight View would be compelling.) However, a proponent of the Total Evidence View will insist that another point deserves equal emphasis, especially as it is apt to be given short shrift by one sympathetic to the Equal Weight View: there is also the opposite phenomenon, that of *upward epistemic push*. That is, a proponent of the Total Evidence View will insist (for the reasons given in Section 5.2 above) that, when E is genuinely good evidence for H, this very fact will contribute to the justification for believing the epistemic proposition that *E is good evidence for H* that is available for those with the relevant competence. It is not only that one's higher-order evidence typically makes a difference to what one is justified in believing about the world; it is also the case that one's first-order evidence makes a difference to what one is justified in believing about higher-level epistemic matters.

Again, the phenomenon of upward epistemic push will be most visible in single-person cases, where distracting complications are at a minimum. As argued above, any case in which one takes up a belief upon recognizing that that belief is supported by one's evidence is *ipso facto* a case in which one is justified in believing a corresponding epistemic proposition to the effect that *one's evidence supports that belief*. Moreover, it is implausible that every case in which one recognizes that a given belief is supported by one's first-order evidence is a case in which one's recognition depends on one's having some independent, higher-order evidence to the effect that one's evidence supports that belief. Rather, in some cases, one's recognition that one's evidence supports a given belief is based on an unmediated appreciation of that evidence itself.³⁵ Thus, in such cases, **(p.160)** one's evidence not only confirms the belief in question, it also confirms a proposition to the effect that it is reasonable for one to hold that belief.

I take the dialectical upshot of this picture to be the following: the proponent of the Total Evidence View can agree with the proponent of the Equal Weight View that facts about what one is justified in believing about the world are constrained by higher-level facts about what one is justified in believing about one's evidence, while denying that this favors the Equal Weight View. (Indeed, the proponent of the Total Evidence will insist that, when these matters are understood correctly, the picture that emerges will positively favor her own view.) In a case of disagreement in which one of two peers evaluates the first-order evidence correctly, the proponent of the Equal Weight View will see a higher-level symmetry and appeal to the link between levels in order to argue that neither peer is justified in favoring his original view once they compare notes. (Symmetry at the higher level creates a symmetry at the lower level that otherwise would not have existed.) In contrast, a proponent of the Total Evidence View will contend that, in such a case, the peer whose view more accurately reflects the evidence will typically be better justified in thinking that his view is the one that is favored by the first-order evidence. (Asymmetry at the lower level tends to create an asymmetry at the higher level, an asymmetry that otherwise would not have existed.)

No doubt, some will be extremely suspicious of the idea that the peer who gets the evidence right is typically better justified in thinking that he has done so than the peer who gets the evidence wrong. In fact, the attempt to exploit such suspicions is central to an argument for the Equal Weight View that has recently been offered by Adam Elga. The final argument for the Equal Weight View that I will consider then, is his.

5.4. A (No) Bootstrapping Argument for the Equal Weight View?

Elga (2007:487) argues as follows:

Suppose that . . . you and your friend are to judge the truth of a claim, based on the same batch of evidence. Initially, you count your friend as an epistemic **(p.161)** peer—you think that she is about as good as you at judging the claim. In other words, you think that, conditional on a disagreement arising, the two of you are equally likely to be mistaken. Then the two of you perform your evaluations. As it happens, you become confident that the claim is true, and your friend becomes equally confident that it is false.

When you learn of your friend's opposing judgment, you should think that the two of you are equally likely to be correct. The reason is [this]. If it were reasonable for you to give your own evaluation extra weight—if it were reasonable to be more than 50% confident that you are right—then you would have gotten some evidence that you are a better evaluator than your friend. But that is absurd.

. . . the absurdity is made more apparent if we imagine that you and your friend evaluate the same long series of claims. Suppose for *reductio* that whenever the two of you disagree, you should be, say, 70% confident that your friend is the mistaken one. It follows that over the course of many disagreements, you should end up extremely confident that you have a better track record than your friend. As a result, you should end up extremely confident that you are a better evaluator. But that is absurd. Without some antecedent reason to think that you are a better evaluator, the disagreements between you and your friend are no evidence that she has made most of the mistakes.

Elga takes the argument of this passage to successfully undermine any alternative to the Equal Weight View. In particular, he takes the argument offered here to undermine both the Extra Weight View—according to which each party to the dispute is permitted to give some special, presumptive weight to his or her own judgment—as well as views akin to the Total Evidence View, on which it matters which of the parties has in fact done a better job evaluating the evidence.³⁶ However, I believe that, while Elga's bootstrapping argument has considerable force against the Extra Weight View, it has little to none against the Total Evidence View.

In order to see this, let us focus our attention directly on the situation in which Elga claims the absurdity of any alternative to the Equal Weight View is most apparent—namely, the situation in which you and your friend each evaluate a long series of claims. Elga formulates the argument as **(p.162)** a *reductio ad absurdum*. The supposition from which the absurd consequences are alleged to follow is this:

whenever you and your friend disagree, you should be, say, 70% confident that your friend is the mistaken one.

Crucially, however, this supposition is *not* something to which the proponent of the Total Evidence View is committed. That is, the proponent of the Total Evidence View is not committed to the idea that, whenever you and your friend disagree, you should be $n\%$ confident that your friend is the one who has made the mistake (where n is some number greater than 50). Indeed, on the contrary: the proponent of the Total Evidence View will stand with Elga in rejecting any such general policy as an unreasonable one. On the Total Evidence View, it is not true, in general, that you should be more confident that your friend has made the mistake whenever the two of you disagree. In *some* cases, it might be reasonable for you to be more confident that your friend is the one who has made the mistake. But, in other cases, it might be reasonable, given the total evidence available to you, to be more confident that you are the one who has made the mistake. On the Total Evidence View, it is not true that there is some general answer to the question of how confident you should be that it is your friend who has made the mistake (as there is on both the Extra Weight View and the Equal Weight View). And this is because how confident it is reasonable to be that your friend has made a mistake is not something that floats entirely free of the evidence on which he bases his opinion. Thus, since the proponent of the Total Evidence View would not accept the supposition from which Elga derives the absurd consequence, the *reductio ad absurdum* on offer cannot show that her view is false.

Consider another view rejected by Elga, the Extra Weight View. As interpreted by Elga, the Extra Weight View would license you in being extremely confident that you are a better evaluator than your friend simply by noting the many cases in which the two of you disagree. In a parallel manner, the Extra Weight View would license your friend in being extremely confident that *he* is the better evaluator by appeal to the very same disagreements. This seems odd (to say the least): the very same events are legitimately treated by you as confirming evidence for the claim that you are a better evaluator than your friend and by your friend as confirming evidence that he is a better evaluator than you. Moreover, even **(p.163)** if you are in fact the inferior evaluator, and you consistently do a worse job evaluating the evidence on particular occasions, it will nevertheless be reasonable for you to conclude that you are superior to your friend *on the basis of those very cases*. (That is, it will be reasonable for you to conclude that you are a better evaluator of evidence on the basis of disagreements whose existence is underwritten by the fact that you have done a *worse* job than your friend has with respect to evaluating the evidence.) Here I agree with Elga: such a view makes it absurdly easy to arrive at evidence that one is a better evaluator. However, no similar absurdity follows from the Total Evidence View. It is true that the proponent of the Total Evidence View is committed to the following possibility: over time, you reasonably become quite confident that someone whom you initially regarded as your peer is not your peer, on the basis of a large number of cases in which the two of you disagree. Consider, for example:

Case 8. At the outset you regard your friend as your peer. Subsequently, however, many disagreements emerge. With respect to the vast majority of these disagreements, the position which you hold is in fact better supported by the available evidence than the position held by your friend. In these cases, your conviction that your friend's position is not adequately supported by his evidence is based on your own appreciation of that evidence, an appreciation which is more accurate than his. Over time, you thus become increasingly confident that you are a better evaluator of the evidence than your friend. You thus cease to regard your friend as your peer and conclude that your initial judgment to that effect was mistaken.

As Elga would have it, the proponent of the Total Evidence View is indeed committed to the possibility that such a change in view is reasonable in the envisaged circumstances. However, there is no absurdity here.

Elga's bootstrapping argument purports to establish that any view other than the Equal Weight View makes it too easy to reasonably conclude that you are a better evaluator than your friend. The danger in question is a real one: some views (for example, the Extra Weight View) do fall victim to it. However, there is also the opposite danger: that a given view will make it *too difficult* to reasonably conclude that another person is not, contrary to what one initially thought, one's peer. Indeed, the line of argument offered by Elga seems to suggest something like the following: once you **(p.164)** come to regard your friend as a peer about a given set of questions, it is not reasonable for you to demote him from the ranks of those to whom you accord that status on the basis of subsequent disagreements about those questions (rather, one would need to have independent evidence that you are a better evaluator than he is, evidence that is independent of the disputed issues themselves). But that seems too strong: to the extent that the argument purports to show this, the argument proves too much. For, in some cases, it might very well be rational for you to conclude that your friend is not your peer after all, where your only basis for so concluding is the lack of judgment that he displays in subsequent cases in which the two of you disagree. The possibility of rationally downgrading someone from the status of peer in this way will be especially apparent in cases in which one's initial judgment that the other person is a peer was itself based on relatively insubstantial evidence. Consider, for example:

Case 9. At the first meeting of our seminar, I strike you as a perfectly reasonable and sensible person. For the most part, we find the same arguments and considerations persuasive. Even on those few occasions when we express different views, my view seems to you to be well within the bounds of reasonable opinion, no less than your own (suppose here that you do *not* accept the Uniqueness Thesis). On the basis of this first meeting, then, you form the opinion that I am your peer.

In subsequent meetings of the seminar, however, you and I disagree often. Moreover, when we disagree, my views often seem to you to be based on relatively flimsy arguments; when I attempt to parry objections, what I say strikes you as weak and unresponsive, and so on. (Needless to say, I would dispute such assessments.) By the end of the semester, you no longer regard me as your peer.³⁷

Here, your revised estimate of my competence is based on your negative assessment of my performance in judging issues that are disputed between us. Moreover, the disputed issues are the very sorts of questions with respect to which you once reasonably took me to be a peer. Does this guarantee that it is unreasonable for you to demote me from the ranks of those to whom you accord such status? There is no such guarantee. On the other hand, there is also no guarantee that your demoting me *is* reasonable in the **(p.165)** circumstances, given only the description of Case 9 offered above. Whether your demoting me is reasonable will typically depend on such things as whether my best attempts to parry objections are weak and unresponsive as you take them to be, or whether your conviction that they are weak and unresponsive is due (for example) to your being so dogmatically committed to the opposite conclusions that you fail to appreciate the merits of what I say. The more the former is the case, the more reasonable it will be for you to revise your estimate of my competence in a downward direction; the more the latter is the case, the less reasonable such revision is. Of course, from your perspective, it might be very difficult to tell which of these is the case. From the inside, a case in which you fail to appreciate the genuine merits of what I say on behalf of my view because of dogmatic commitment on your part might seem just like a case in which my defense is indeed without merit. But the fact that it might be difficult to tell which of these is the case does not mean that it makes no difference whether your revised estimate of my competence is based on your having recognized genuine shortcomings on my part, or whether it is instead an artifact of your own shortcomings. Here, as elsewhere, there is no escape from the fact that one's judgment is fallible and subject to corruption in ways that tend to elude detection. According to Elga, (i) the relevant kind of bootstrapping is never rationally permissible, (ii) the Equal Weight View proscribes such bootstrapping, and (iii) no other plausible view does so. He thus concludes that the Equal Weight View is true. I hold that, on the contrary, because there are at least some possible cases in which such bootstrapping clearly *is* permissible, no view that generally proscribes it can be correct. Hence, on the assumption that Elga is correct in thinking that the Equal Weight View generally proscribes such bootstrapping, we have arrived at another good reason for thinking that it is false.

Indeed, Elga's blanket prohibition on the kind of bootstrapping at issue here seems to sit in at least some tension with moves that he makes elsewhere in the course of defending the Equal Weight View. Consider, for example, his argument that the Equal Weight View does not require one to suspend judgment about all controversial issues (200: 492-4), a conclusion that would be, he thinks, an absurd consequence. In attempting to block this 'problem of spinelessness,' Elga emphasizes that we should not overestimate how often we find ourselves in disagreements with those **(p.166)** whom we take to be our peers, inasmuch as I will not consider you my peer with respect to a question if you disagree with me about too many surrounding issues. (Here, it seems, Elga would permit one to appeal to one's own beliefs in order to conclude that someone who disagrees with sufficiently many of those beliefs is not a peer.)

However, there is at least a certain awkwardness in attempting to combine (i) a blanket prohibition against bootstrapping one's way to the conclusion that someone who one initially took to be a peer is not a peer, and (ii) permitting one to appeal to one's own beliefs in order to deny the status of peer to another person so long as one has not yet accorded him that status. For suppose that you and your friend disagree, not only about the moral permissibility of abortion, but also about many surrounding issues—for example, whether human beings have souls, whether it is permissible to withhold treatment from certain terminally ill infants, whether rights figure prominently in a correct ethical theory, and so on.³⁸ According to Elga, because your friend has by your lights come to the wrong conclusion about this entire cluster of closely related issues, you (reasonably) do not consider him your peer, and thus, you are not required to split the difference with him about these issues. However, even this amount of disagreement is presumably compatible with a large amount of agreement concerning moral matters. Bearing this mind, consider two cases. In Case A, you first discover all of the issues with respect to which you and your friend disagree; you thus conclude, reasonably, that your friend is not your peer with respect to difficult moral questions. You can therefore retain your original views, since you are not rationally required to split the difference and retreat to a state of agnosticism. Later, it emerges that you and your friend also agree on many moral issues; you thus view these cases as ones in which someone who is not your peer nevertheless manages to arrive at the correct conclusions. In contrast, in Case B you first happen to discover all of the moral issues with respect to which you and your friend agree; you thus (reasonably) conclude that your friend is your peer with respect to difficult moral questions. Later, all of the disagreements emerge. However, because you have already granted your friend the status of peer, it is no longer permissible, given the prohibition on bootstrapping, to appeal to these disagreements as a legitimate basis for demoting him or for discounting his **(p.167)** opinion: at this point, you are already on the normative hook, as it were. Ultimately, you are fully aware of all of your friend's opinions in both cases; the only difference is that, in Case B, you are rationally required to suspend judgment on all of the disputed questions, while in Case A, you are not. But intuitively, that seems wrong. For the only underlying difference between the two cases is the *order* in which you learn of your friend's opinions.

6. The Total Evidence View: Concluding Remarks

In the course of laying out the core ideas of the Total Evidence View, we have repeatedly compared and contrasted it with various alternatives, especially the Equal Weight View. As we have seen, the Total Evidence View and the Equal Weight View yield the same verdict when applied to some cases of disagreement; in others, the two views differ sharply in what they require of the disputing parties. In still other cases, I think that it is somewhat unclear how far the two views differ (if they differ at all), inasmuch as it is somewhat unclear how the Total Evidence View should be applied.

Consider, for example, cases in which individuals of apparently similar mathematical abilities arrive at different conclusions on the basis of temporally extended token processes of calculation.³⁹ In some cases of this sort, I believe that the Total Evidence View, when properly interpreted, will yield the verdict that the individuals should split the difference. Consider, for example:

Case 10. You and I add a series of ten three-digit numbers in our heads. A third party calls out the numbers, one after the other. Each of us keeps a running tally, adding the numbers as we go, not attempting to keep track of any particular number in the sequence after it has been added to the running total. We know that, when we have played this game in the past, we have made a more or less equal number of mistakes. This time, I arrive at the number 5,863 and you arrive at the number 5,883. Once we discover that we have arrived at different answers, how should we respond?

(p.168) One might think that, on the Total Evidence View, although both of us should be less confident of our original answers, whichever one of us has in fact performed the calculation correctly is rationally entitled to be more confident of her answer than the other person. However, in this case I believe that there is a stronger, countervailing pressure from within the Total Evidence View that militates in favor of splitting the difference. On the Total Evidence View, what it is reasonable for us to believe always depends on the total evidence that we possess. Once we learn that we have arrived at different answers, then, given that we have not retained specific information about the original numbers (and so are not in a position to reconstruct our original reasoning), it seems as though the total relevant evidence available to us consists of the fact that (i) I arrived at the number 5,863 and (ii) you (a person of apparently similar mathematical abilities) arrived at the number 5,883. Plausibly, this evidence does not favor either answer over the other. (Compare the question of what it would be reasonable for a third party to believe, a person who was not present when the original numbers were called out and who knows only (i) and (ii).) Hence, it seems as though the Total Evidence View requires us to split the difference in these circumstances.

Suppose that we subsequently go over our reasoning together, step by step. (The person who originally called out the numbers recorded them on a list.) If we are in fact both competent at arithmetic, the person who originally made the mistake will presumably correct his or her error and adopt the true view. Imagine, however, that our dispute turns out to have the kind of persistent, intractable character that many philosophical disputes seem to possess: despite explicitly rehearsing our chains of reasoning in a public manner, both of us continue to think that his or her original answer is correct. Of course, because of the quasi-algorithmic character of arithmetical reasoning, it is difficult to imagine that a dispute of this sort could persist. But, if such were to happen, I think that it is reasonable for the person whose answer is based on the sound mathematical reasoning to be relatively confident of her answer. She occupies a superior epistemic position compared to the person whose answer is based on the unsound reasoning. Yet it is far from clear how we should think about the mathematical evidence in such a case, or in cases of calculation more generally. To the extent that the relevant notion of mathematical evidence remains unclear, there will be many cases of **(p.169)** calculation in which it is unclear how exactly the Total Evidence View applies.

Still, even if there are unclarity about how the Total Evidence View applies in particular cases, the view is at least clear enough to be controverted. I close by commenting briefly on two features that are especially likely to arouse suspicion in certain quarters.

First, a central feature of the view is that the reasonableness of the parties in a case of peer disagreement will typically depend on whose opinion better reflects the first-order considerations relevant to their dispute. When one is responding correctly to the evidence, one is typically in a stronger position vis-à-vis those who think otherwise than when one merely takes oneself to be responding correctly to the evidence. Of course, there is no magic red light that illuminates when one responds to the evidence correctly, no warning bell that sounds when one does not. Indeed, as a phenomenological matter, there might be no introspectible difference between how things seem when one is responding correctly and how things seem when one is not. Given this, how can a decision to adopt the Total Evidence View (or any view that shares this central feature) reflect anything other than a meta-epistemological commitment to externalism about justification, with the Equal Weight View left as the view of choice for those with more internalist sympathies?

However, the classification of the Total Evidence View as an 'externalist' as opposed to an 'internalist' view is not a happy one. On the Total Evidence View, what it is reasonable for one to believe always depends on one's total evidence, and only considerations of which one is aware are eligible for inclusion in one's total evidence. (Relevant considerations that are known to others but of which one is unaware make no difference to what it is reasonable for one to believe.) In this crucial respect, the Total Evidence View resembles epistemological views that are paradigmatically 'internalist.' Of course, in a case in which one's view is not adequately supported by one's evidence, there might be nothing that indicates that this is so (that is, nothing *else*, beyond the evidence itself, that *ex hypothesi* one has misjudged). Because of this, when one's judgment as to the epistemic status of some belief that one holds is faulty, there is nothing that guarantees that this fact will be revealed by further reflection, no matter how conscientiously such reflection is conducted. But it is dubious that this last feature is avoided by **(p.170)** any plausible view about justification, including paradigmatically internalist ones.

A related but ultimately deeper source of resistance to the Total Evidence View is the tendency to identify *good evidence* with *potentially persuasive evidence*. Relevant here is the seductive appeal of what Timothy Williamson (2004, 2007) has dubbed 'the dialectical conception of evidence.' Recall from above the way in which a proponent of the Total Evidence View will differ from Elga regarding the legitimacy of a certain kind of bootstrapping. On the Total Evidence View, even if you reasonably take me to be your peer with respect to a given class of questions, you might later reasonably revise this judgment and conclude that I am not your peer on the basis of how I answer those very questions.⁴⁰ Consider a case in which you demote me in this way, and imagine that I subsequently learn that you no longer consider me a peer. Naturally enough, I inquire about your basis for demoting me. What evidence have you gained, since the time when you reasonably took me to be your peer, that suggests that I am not? In response to this query, you might cite one of those issues with respect to which I have by your lights misjudged the evidence. That is, for some issue about which we disagree, you might say the following:

With respect to this particular issue, the view that I hold is adequately supported by the evidence while the view that you hold is not. Hence, this case is a piece of confirming evidence for the claim that I am a better evaluator of evidence than you are.

Needless to say, I will not be impressed with this response, or agree that you have succeeded in providing any evidence of your superiority. By my lights, this is mere assertion on your part (and, indeed, false assertion at that). Moreover, it should come as no surprise to you that this will be my reaction. Given that the question of what the evidence supports is a matter that is contested between us, the question of which one of us has done a better job of evaluating that evidence will also be a contested matter. For this reason, it would be completely unreasonable on your part to expect me to treat what you say here as evidence of your superiority. Indeed, it would be pointless for you to offer these alleged facts to me as evidence **(p.171)** of your superiority, for from my perspective they are not facts at all. But, if it would be pointless for you to cite these alleged facts as evidence in response to my request, then it would be *dialectically inappropriate* for you to do so. According to the dialectical conception of evidence, only considerations that it would be dialectically appropriate to cite as evidence are genuine evidence. On the dialectical conception of evidence, then, the alleged facts that you cite as evidence of your superiority do not constitute such evidence, and thus have no tendency to justify your newfound belief in your superiority (even if what you say is true, and the alleged facts to which you appeal are genuine facts).

Moreover, the question of whether genuine evidence must consist of considerations that it would be dialectically appropriate to cite as such is of quite general relevance to the topic of peer disagreement. *Ex hypothesi*, when you and I disagree about an issue with respect to which we are peers, the fact that you do not share my view is not due to my having access to some crucial piece of evidence to which you lack access. One might think that this already renders problematic the claim that my evidence suffices to justify my belief: if my evidence were really sufficient to justify my believing as I do, wouldn't that evidence be enough to persuade you as well? Given that the evidence on which I base my belief does not persuade you, can't we conclude from this that my evidence is not sufficient to justify my belief after all? But in any case, when faced with someone who shares my evidence yet remains unconvinced, it would be pointless for me simply to recite the same considerations again and claim that that is why my view is the reasonable one to hold. Inasmuch as it would be pointless for me to cite these considerations in this way, it would be dialectically inappropriate for me to do so. Thus, if, in order for it to be the case that my evidence genuinely favors my view over yours, it must be the case that it would be dialectically appropriate for me to cite my evidence as favoring my view over yours, then my evidence does not favor my view over yours.

I believe that we have good reasons to reject the dialectical conception of evidence. One might have good evidence that some claim is true, even if one has no potentially persuasive evidence, or evidence that it would be dialectically appropriate to cite as such.⁴¹ Indeed, I believe that one can **(p.172)** have good evidence that some claim is true even if one has no evidence that it would be dialectically appropriate to offer to a person who is in general *no less reasonable than oneself*. Because lapses and blindspots are possible, the fact that a generally reasonable person fails, even repeatedly fails, to be persuaded of some conclusion by a body of evidence does not suffice to show that that evidence is inadequate to justify belief in that conclusion. Of course, once it is clear that the person does not find the considerations on offer persuasive, the dialectically appropriate course is to seek new considerations that might inspire conviction rather than simply to recite the original considerations yet again. But it does not follow from this that the original considerations were themselves inadequate to justify belief in the conclusion. The link between genuine evidence and potentially persuasive evidence is not as close as the dialectical conception of evidence suggests. At best, what is true is a relatively trivial claim: genuine evidence is evidence that will tend to persuade someone who will respond to that evidence in a fully reasonable manner. Still, it must be admitted that the dialectical conception of evidence is not wholly without its appeal. Inasmuch as this is so, a defense of the Total Evidence View more thorough than the one offered here will account for this appeal in a way that reveals it to be spurious.

References

Bibliography references:

Alston, William (1980). "Level Confusions in Epistemology," *Midwest Studies in Philosophy*, 5: 135-50.

Christensen, David (2007). "Epistemology of Disagreement: the Good News," *Philosophical Review*, 116/2: 187-217.

— (2008). "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals," in *Oxford Studies in Epistemology*, ii. Oxford: Oxford University Press, 3-31.

Cohen, G. A. (2000). *If You're an Egalitarian, How Come You're So Rich?* Cambridge, MA: Harvard University Press. **(p.173)**

Conee, Earl, and Feldman, Richard (2004). *Evidentialism*. Oxford: Oxford University Press.

Edwards, Paul (1967). "Common Consent Arguments for the Existence of God," in Paul Edwards (ed.), *The Encyclopedia of Philosophy*, ii. New York: Macmillan, 147-55.

- Elga, Adam (2007). "Reflection and Disagreement," *Noûs*, 41/3: 478-502.
- Feldman, Richard (2003). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- (2005). "Respecting the Evidence," *Philosophical Perspectives*, xix: *Epistemology*. Oxford: Blackwell, 95-119.
- (2006). "Epistemological Puzzles about Disagreement," in Stephen Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press, 216-36.
- (2007). "Reasonable Religious Disagreements," in Louise Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life*. Oxford: Oxford University Press, 194-214.
- Goldman, Alvin (2001). "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research*, 63: 85-110
- Greco, Daniel (2006). "Disagreement and the Equal Weight View." Unpublished senior thesis, Princeton University.
- Kelly, Thomas (2005a). "The Epistemic Significance of Disagreement," in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, i. Oxford: Oxford University Press, 167-96.
- (2005b). "Comments on White, 'Epistemic Permissiveness.'" Remarks delivered at the 2005 Eastern Division meetings of the American Philosophical Association. <www.princeton.edu/~tkelly/>
- Plantinga, Alvin (2000a). "Pluralism: A Defense of Religious Exclusivism," in Philip L. Quinn and Kevin Meeker (eds.), *The Philosophical Challenge of Religious Diversity*. Oxford: Oxford University Press, 172-92.
- (2000b). *Warranted Christian Belief*. Oxford: Oxford University Press.
- Pryor, James (2004). "What's Wrong with Moore's Argument?" *Philosophical Issues*, 14: 349-78
- Rosen, Gideon (2001). "Nominalism, Naturalism, Epistemic Relativism," *Philosophical Perspectives*, xv. Malden, MA: Blackwell, 69-91.
- Sosa, Ernest (1999). "Skepticism and the Internal/External Divide," in John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology*. Malden, MA: Blackwell.
- (2002). "Privileged Access," in Quentin Smith and A. Jokcik (eds.), *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press, 273-95.

— (2007). *A Virtue Epistemology*. Oxford: Oxford University Press. **(p.174)**

Van Inwagen, Peter (1996). "It is Wrong, Everywhere, Always, and for Anyone, to Believe Anything upon Insufficient Evidence," in Jeff Jordan and Daniel Howard-Snyder (eds.), *Faith, Freedom, and Rationality: Philosophy of Religion Today*. London: Rowman and Littlefield, 137-53.

Wedgwood, Ralph (2007). *The Nature of Normativity*. Oxford: Oxford University Press.

White, Roger (2005). "Epistemic Permissiveness," in *Philosophical Perspectives*, xix: *Epistemology*. Oxford: Blackwell Publishers, 445-9.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

— (2004). "Philosophical 'Intuitions' and Skepticism about Judgment," *Dialectica*, 58/1: 109-53.

— (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.

Notes:

* This chapter is something of a sequel to Kelly (2005a). While in many respects it is faithful to the position advanced there, it departs in others; significant departures are noted along the way. Earlier versions of this chapter were presented at New York University, MIT, Rutgers University, Brown University, Princeton University, and the University of California at Irvine; I am grateful to the audiences present on those occasions. In addition, I would like to thank Aaron Bronfman, David Christensen, Adam Elga, Hartry Field, Allan Gibbard, Margaret Gilbert, Daniel Greco, Aaron James, Jim Joyce, Sarah McGrath, Philip Pettit, Jim Pryor, Walter Sinnott-Armstrong, Roy Sorensen, and Ernest Sosa for helpful conversations on the topic.

(1) Of course, the kind of uncontroversial 'track-record' evidence that bears most directly on questions of comparative reliability will be much easier to come by in some domains than in others. (In this respect, contrast reliability in accurately forecasting the weather and reliability in accurately answering metaphysical questions.)

(2) Cf. Feldman (2003), who, after reviewing a number of examples of the kind at issue here, draws the conclusion that, “in the situations most plausibly thought to be cases of reasonable disagreement, suspension of judgment is the reasonable attitude to take toward the disputed proposition” (p. 189). The Equal Weight View is explicitly embraced by Adam Elga (2007), whose views I consider at some length below; David Christensen (2007) exhibits considerable sympathy for a policy of ‘splitting the difference’ throughout his own discussion of the topic. Although the view that I will put forth differs from theirs, I have learned much from each of these authors.

(3) A case of this general form was put to me by Roy Sorensen in conversation. Cf. Christensen's ‘Acme watch’ example (2007: 196) and Feldman (2006: 234).

(4) Notable here are van Inwagen (1996), Plantinga (2000a, b), and Rosen (2001); another is Kelly (2005).

(5) Cf. ‘the Extra Weight View’ discussed by Elga (2007), who argues against it.

(6) See, e.g., Feldman (2003, 2006).

(7) Again, this is characteristic of Feldman's work on the topic.

(8) ‘The Uniqueness Thesis’ is Feldman's label (2007); cf. Christensen's ‘Rational Uniqueness’ (2007). Feldman both argues for and endorses the thesis; Christensen exhibits some sympathy for the thesis and offers some considerations for thinking that it is true. White (2005) argues for the thesis at length but stops short of endorsing it.

(9) Most plausible, but still not especially plausible, I think. Again, it comes under pressure from marginal cases. Suppose that the evidence available to me is just barely sufficient to justify my belief that it will rain tomorrow: if the evidence was even slightly weaker than it is, then I would be unjustified in thinking that it will rain. Suppose further that you have the same evidence but are slightly more cautious than I am, and so do not yet believe that it will rain tomorrow. It is not that you are dogmatically averse to concluding that it will rain; indeed, we can suppose that, if the evidence for rain gets even slightly stronger, then you too will take up the relevant belief. Is there some guarantee, given what has been said so far, that you are being less reasonable than I am?—I doubt it.

(10) Here, for example, is Gideon Rosen (2001: 71): “It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable.”

(11) See, e.g., the brief survey in White (2005: 445–6).

(12) I take the most formidable case to have been made by White (2005), although he himself does not endorse the thesis. I respond to some, though not all, of White's arguments in Kelly (2005b).

(13) Is there some way of interpreting the Equal Weight View so that it does not have the consequence in question? For some variant interpretations and the difficulties that beset them, see Sect. 3.2 below.

(14) The objection raised in this section is due, in all of its essential features, to Aaron Bronfman. I utilize it here with his permission.

(15) In any case, I take it that it is not an acceptable consequence for an evidentialist like Feldman, who explicitly maintains that what one is justified in believing at any given time supervenes on what evidence one possesses at that time. See Conee and Feldman (2004). especially essay 4 and the introduction.

(16) Consider another possible interpretation of the Equal Weight View designed to avoid the charge of bootstrapping (a suggestion that is due to Jim Pryor in conversation). According to this interpretation, the Equal Weight View should be understood as a theory about how those who respond to their evidence perfectly will respond to peer disagreement. Strictly speaking, then, the view is silent on how someone who has misjudged her original evidence should respond to the discovery that a peer disagrees. Thus, when you and I encounter one another, you (who responded to the original evidence correctly) are rationally required to split the difference with me, but it is no part of the Equal Weight View that I (who responded incorrectly) am rationally required to do the same. So interpreted, the Equal Weight View is not susceptible to the bootstrapping objection; moreover, unlike the alternative interpretation just considered, it does not have the consequence that two beliefs of the same type held on exactly the same total evidence might differ with respect to epistemic status.

However, this version of the view strikes me as poorly motivated in the extreme. If the phenomenon of peer disagreement requires you to split the difference with my unreasonable opinion, why should I be spared having to split the difference with your reasonable opinion simply in virtue of having botched the evidence in the first place? Whatever normative pressure is created by the phenomenon of peer disagreement, surely one does not immunize oneself against that pressure simply in virtue of having beliefs that are *not* adequately supported by one's evidence.

(17) Interestingly, this point is emphasized both by those who are sympathetic to the Equal Weight View as well as by those who seek to resist it. Examples of the former include Feldman (2006) and Elga (2007); an example of the latter is van Inwagen (1996).

(18) See especially Elga (2007).

(19) Some might find this terminology suboptimal on the grounds that *all* of one's evidence is ultimately psychological inasmuch as it consists of one's own psychological states. I think that this complaint rests on a mistaken view about the ontology of evidence, but no matter: one who thinks that all of our evidence ultimately consists of psychological states might read 'psychological evidence' and 'non-psychological evidence' as 'doxastic evidence' and 'non-doxastic evidence' in what follows.

(20) At least, so long as one has no *independent* grounds for attributing such performance errors. Of course, it is open to a proponent of the Equal Weight View to say that, even if you and I possess similar general competence, it is permissible for you to discount my opinion when, for example, you notice that I was distracted while surveying the evidence in a way that you were not, or that I did so while under the influence of some temporarily mind-numbing drug, or so on. What the proponent of the Equal Weight View will *not* allow is that my actually having committed a performance error can make a difference when your only grounds for attributing such an error to me consists in the fact that I have arrived at (what you take to be) an incorrect answer to the question about which we disagree. It is this feature of the Equal Weight View that distinguishes it from the alternative view that I will offer and that leaves it vulnerable to the current objection.

(21) We might also imagine cases in which I am confused, or it is temporarily opaque to me, whether a given piece of evidence that favors the proposition that it will rain tomorrow consists of (i) the fact that my peer believes that it will rain tomorrow, or, alternatively, (ii) the fact that the machine is in State A. Will the proponent of the Equal Weight View insist that, once I learn the truth, a significant revision in my opinion about whether it will rain tomorrow might be in order?

(22) See Sect. 5.4 below.

(23) Cf. the lucid and illuminating discussion of this point in Christensen (2007, 2008).

(24) In Kelly (2005), I suggested that we should regard the views of a generally reasonable person as in effect providing *higher-order evidence*: that is, evidence about the normative upshot of the evidence to which she has been exposed. (See especially the discussion on pp. 185–90.) So, for example, the fact that a generally reasonable person *S* believes *p* is (defeasible) evidence in favor of the epistemic proposition that *it is reasonable to believe p given S's evidence*. I emphasized that higher-order evidence of this sort bears most directly on epistemic propositions and that acquiring such evidence will often make a straightforward difference to what it is reasonable for one to believe *about* particular bodies of evidence. On the other hand, I expressed considerable skepticism about the idea that the higher-order evidence provided by the fact that a generally reasonable person believes a given proposition will also make a difference to what it is reasonable for one to believe about that proposition in a case in which one knows that one already possesses all of the evidence on which the person bases her belief. (Foremost among my reasons for skepticism: the ‘double-counting’ argument rehearsed on pp. 187–8.) What I say here constitutes a departure from the earlier skeptical attitude: on the present view, higher-order evidence about the bearing of one's first-order evidence *is* typically relevant to what it is reasonable to believe on the basis of that evidence.

(25) Recent—and, to my mind, compelling—critiques of the idea that there is any interesting and important epistemic status that supervenes on phenomenology are provided by Ernest Sosa (1999, 2002, 2007) and Timothy Williamson (2000).

(26) On the importance and nature of independence, see especially the illuminating discussion in Goldman (2001: 150–6). In that paper Goldman is specifically concerned with the interesting question of how a non-expert should respond to disagreement among the experts, but the analysis of independence that he offers would seem to be highly relevant to a host of other important issues in social epistemology as well.

(27) Perhaps unsurprisingly, the Common Consent Argument is not taken very seriously any more, even in those circles in which arguments for the existence of God are still taken seriously. It is, for example, rarely if ever included among the usual rogue's gallery of arguments for the existence of God (the ontological argument, the cosmological argument, etc.) in anthologies or course syllabi devoted to the philosophy of religion. Historically, however, it was taken *quite* seriously. A list of prominent thinkers who endorsed some recognizable variant of it would include Cicero, Seneca, the Cambridge Platonists, Gassendi, and Grotius; in addition, it was discussed critically by (among many others) both Locke and Mill. For an overview, see the useful survey in Edwards (1967).

(28) Indeed, as Hartry Field pointed out to me, the need to discount the numbers is not limited to cases in which there is causal dependence present, as in the examples considered above. If I know that two individuals will respond to given evidence in the same manner, then I should treat their having arrived at some particular answer as *one* piece of evidence, and not two pieces of evidence, in favor of that answer (even if their both having arrived at that answer is in no way underwritten by some causal link).

(29) See Christensen (2007: 200–3) and Elga (2007: 491).

(30) In general, it is important to distinguish between (i) cases in which multiple individuals have equally strong but different bodies of evidence and (ii) cases in which multiple individuals have equally strong bodies of evidence in virtue of sharing the same evidence. Splitting the difference will often be the reasonable response in the former kind of case, but this in itself has no tendency to show that the same is true in cases of the latter kind.

Of course, a commitment to certain views about the nature of evidence might make it difficult, if not impossible, to consistently observe the distinction between (i) and (ii). For example, on a view of evidence according to which one's evidence ultimately consists of one's own private mental states, one never literally shares one's evidence with a peer; at best, one's evidence is similar in various salient respects to the evidence that one's peer possesses. Because this is the closest surrogate for genuinely sharing evidence in the literal sense, it becomes easy to conflate (i) and (ii). But such conflation should be resisted.

(31) Readers of the present chapter still unconvinced of the truth of the Total Evidence View will no doubt find *this* thought experiment an eminently manageable one.

(32) For suggesting a response along these lines, I am grateful to Roy Sorensen (who originally put the thermometer objection to me, in conversation). I am unsure to what extent he takes the response to adequately defuse the objection.

(33) In response to this objection, a proponent of the Total Evidence View might contend that it rests on a 'level confusion,' in the sense of Alston (1980): in particular, that it falsely assumes that, in order to be justified in believing *p*, one must be justified in believing that one is justified in believing *p*. In effect, such a response concedes, at least for the sake of argument, that there is a higher-level normative symmetry between the peers but denies that anything directly follows from this about the epistemic statuses of their first-order beliefs. Here I simply want to note the possibility of such a response without exploring its prospects; the response that I offer in the main text proceeds along quite different lines.

(34) This point is well emphasized by Feldman (2006) and Christensen (2007). Here is a representative quotation from the latter: “the rationality of first order beliefs cannot in general be divorced from the rationality of certain second order beliefs that bear on the epistemic status of the first order beliefs” (p. 18).

(35) Here is what I take to be a more or less conclusive reason for denying that the recognition that some body of evidence supports a given conclusion must always be based on some independent, additional evidence to that effect. Let E represent all of the evidence that you currently possess. Surely you can recognize that E supports the belief that *the sun is larger than the moon*. But this recognition is not based on some independent evidence that you possess (i.e., evidence not included in E), since, *ex hypothesi*, E exhausts what evidence you have. (For this way of making the point I am grateful to Nick Beckstead.)

(36) Elga makes the last point explicit on the same page: “Again, this absurdity is independent of who has in fact evaluated the claims properly. Even if in fact you have done a much better job than your friend at evaluating the claims, simply comparing your verdicts to those of your friend gives you no evidence that this is so” (2007: 487).


(37) This case was inspired by a similar example devised by Daniel Greco.

(38) I take this example directly from Elga (2007: 493).

(39) I take this to be among the issues raised by Christensen's ‘Restaurant Case’ (2007).

(40) Again, whether it is reasonable for you to downgrade me in this way will typically depend on whether you are correct in your assessment that my performance in evaluating our shared evidence has been inferior to your own.

(41) See especially Williamson (2004). As Williamson notes, acceptance of the dialectical conception of evidence would hand a cheap and sweeping victory to the crudest of skeptics. Thus, against a skeptic who consistently maintained that nothing is evidence for anything else, anything that one might offer as evidence would fail to qualify as such when judged by the dialectical standard; if meeting the dialectical standard was necessary for something to count as genuine evidence, one would have no genuine evidence at all when in the presence of such a skeptic. But surely this is incorrect. One can have genuine evidence, i.e., evidence that tends to justify one's beliefs, even when one has no evidence that it would be dialectically appropriate to offer. On the dangers of not recognizing the distinction in question, see also Pryor (2004).



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

How to Disagree about How to Disagree

Adam Elga

DOI:10.1093/acprof:oso/9780199226078.003.0008

Abstract and Keywords

This chapter argues for the incoherence of unrestricted conciliatory views according to which one ought always to modify one's view in the direction of one's peer in cases of peer disagreement. It contends that such views face incoherence when applied to themselves, and goes on to defend a partly-conciliatory view, recommending conciliation on topics other than disagreement itself.

Keywords: peer disagreement, conciliation, self-undermining views, incoherence

1. Introduction

Suppose* that you and a friend independently evaluate a factual claim, based on the same relevant evidence and arguments. You become confident that the claim is true. But then you find out that your friend—whose judgment you respect—has become just as confident that the claim is false. Should that news at all reduce your confidence in the disputed claim?

Conciliatory views on disagreement answer “yes.” According to such views, finding out that a respected adviser disagrees with one should move one at least a little in the direction of the adviser's view. And it should do so regardless of the subject matter under dispute. Conciliatory views are extremely natural and appealing (Christensen 2007; Elga 2007; Feldman 2007). But they seem to run into trouble when the topic under dispute is disagreement itself. Can conciliatory views accommodate disagreement about disagreement? And, if not, what does this show about what view on disagreement we should adopt instead?

I will consider two arguments that conciliatory views cannot accommodate disagreement about disagreement. Though the first argument fails, the second argument succeeds. So conciliatory views are unacceptable. But the considerations that show this make no trouble for views that **(p.176)** are partly conciliatory: views that recommend compromise in the face of disagreement about many matters, but not about disagreement itself.

2. First Argument against Conciliatory Views: Repeated Disagreements with the Stubborn

Can conciliatory views accommodate disagreement about disagreement? Here is a reason to think not.¹ Suppose that you and your friend disagree about the right response to disagreement. You have a conciliatory view, but you realize that your friend has the **stubborn view**, according to which disagreement is never cause for changing one's view on a disputed issue. It can sometimes seem as though your conciliatory nature dooms you to conceding everything to your stubborn friend, given enough discussion. Here is a representative scenario:

You think it will rain tomorrow, and your friend thinks it will not. (Here and henceforth I assume that you respect the opinions of all of your friends, and that you and your friends have the same evidence relevant to contested issues.) In response to the disagreement, you are conciliatory: you reduce your confidence that it will rain. But your friend is stubborn: he remains completely unmoved. After this first stage, a (slightly less extreme) disagreement about the weather remains. Again you are conciliatory, and further reduce your confidence that it will rain. And again, your friend stands fast. Disagreement still remains. You reduce your confidence a third time, and so on. As the discussion continues, you get pushed arbitrarily close to completely adopting your friend's view on whether it will rain.

In this case, it looks as though your conciliatory nature commits you to conceding an increasing amount, the more times you pool opinions with your stubborn friend. And this looks to be a general phenomenon. If so, that counts against conciliatory views on disagreement. For it is implausible that one should be required to give so much ground to an adviser just because the adviser is stubborn. A similar difficulty arises in the **(p.177)** case of advisers who are not completely stubborn but who have a policy of conceding very little in cases of disagreement.

That is the first argument against conciliatory views on disagreement.

3. Reply: Conciliatory Folk Need not Concede Everything to Stubborn Folk

Here is a reply: sensible conciliatory views do not entail that one should arbitrarily concede a great deal to stubborn advisers.

To see why not, imagine a cluster of advisers who you know exhibit an extreme form of groupthink: they always end up agreeing with one another. Now, you may well respect the opinions of that group. So you may well be moved if you find out that one of them disagrees with you about a particular issue. But suppose that you then find out that another member of the group also disagrees with you about that issue. That news does not call for any additional change in your view. For you knew in advance that the group members all think alike. So hearing the second dissenting opinion gives you no real new information.

In contrast, suppose that you receive an additional dissenting opinion from an adviser who formed her opinions completely independently from your first adviser. In that case, the second dissenting opinion does call for additional caution. The difference is that in this case you did not know in advance what conclusion the second adviser would reach.

The general point is that an additional outside opinion should move one only to the extent that one counts it as independent from opinions one has already taken into account.² The above example illustrates the most extreme version of this point: when one knows with certainty in advance what an adviser thinks, hearing that adviser's opinion should have no impact. But the point also holds in less extreme cases. For example, suppose that two of your friends *almost* always think alike. Then hearing that the first friend disagrees with you should have a big impact on your opinion. But suppose that you later learn that the second friend endorses the judgment of the first. That news should have only a tiny additional impact on your opinion. **(p.178)**

The above independence point is completely uncontroversial, and every sensible view on disagreement should accommodate it.³ Furthermore, conciliatory views on disagreement face no special difficulties in doing so.

Now return to the case in which you disagree about the weather with a stubborn friend. When you find out about the initial disagreement, you should indeed be significantly moved in the direction of your friend's view. But, at the second stage, news of the disagreement should not move you at all. The reason is the same as in the groupthink case: since you knew in advance about your friend's stubborn nature, his continued disagreement provides you with no additional news. Putting things another way: you count his opinion at the first stage of the dispute as completely correlated with his opinion at subsequent stages. As a result, a sensible conciliatory view will counsel you to remain unmoved at the second and subsequent stages.

A similar analysis applies in the case of an adviser who is not completely stubborn, but who has a known policy of conceding very little in cases of disagreement. The initial disagreement of such an adviser should have a big impact on your opinion. But when the adviser keeps putting forward the same view in subsequent disagreements, that should have little or no additional impact.

Moral: sensible conciliatory views do not require one to concede everything to stubborn advisers. That answers the argument.

4. Second Argument against Conciliatory Views: Such Views Undermine Themselves

Next Argument.

Just as people disagree about politics and the weather, so too people disagree about the right response to disagreement. For example, people disagree about whether a conciliatory view on disagreement is right. **(p.179)** So a view on disagreement should offer advice on how to respond to disagreement about disagreement. But conciliatory views on disagreement run into trouble in offering such advice.

The trouble is this: in many situations involving disagreement about disagreement, conciliatory views call for their own rejection. But it is incoherent for a view on disagreement to call for its own rejection. So conciliatory views on disagreement are incoherent. That is the argument.⁴

To see why conciliatory views sometimes call for their own rejection, consider an example. Suppose that you have a conciliatory view on disagreement, but you find out that your respected friend disagrees. He has arrived at a competing view (about disagreement), and tells you all about it. If your conciliatory view is correct, you should change your view. You should be pulled part way toward thinking that your friend is right. In other words, your view on disagreement requires you to give up your view on disagreement.

One might try to avoid this result by adding a special restriction to one's conciliatory view. For example, one might say that one should *in general* be moved by disagreement, but not when the disputed topic is disagreement itself. But such a restriction seems objectionably arbitrary and ad hoc.⁵ If one should be sensitive to disagreement about so many other matters, then why not about disagreement, too? (Certainly not because disagreement is an easy or uncontroversial topic, as the existence of this volume attests.)

So: conciliatory views on disagreement sometimes call for their own rejection. The next section explains why views on disagreement that call for their own rejection are incoherent. It will follow that conciliatory views on disagreement are incoherent.

5. Self-Undermining Views are Incoherent

Why is it incoherent for a view on disagreement to call for its own rejection? To see why, notice that one's view on disagreement is part of one's *inductive method*: one's fundamental method for taking evidence into **(p.180)** account. An inductive method offers recommendations on what to believe based on one's course of experience. Given a course of experience, an inductive method says what one should believe about various topics: the weather, who will win the next election, and so on. It even says how a given course of experience bears on the question "which inductive method should one use?"

Now suppose that one's view on disagreement sometimes calls for its own rejection. Then one's inductive method also sometimes calls for its own rejection. For one's view on disagreement is part of one's inductive method.⁶ So, in order to show that self-undermining views on disagreement are incoherent, it is enough to show that self-undermining inductive methods are incoherent.

That is best illustrated by the following example.⁷

The magazine *Consumer Reports* rates appliances, and gives recommendations on which ones to buy. But pretend that, in addition to rating appliances, *Consumer Reports* also rates and recommends consumer-ratings magazines. Then it cannot coherently recommend a competing magazine over itself. (By a "competing magazine" I mean a magazine that offers contrary appliance recommendations.)

To see why not, consider an example. Suppose that *Consumer Reports* says, "Buy only toaster X," while *Smart Shopper* says, "Buy only toaster Y." And suppose that *Consumer Reports* also says, "*Consumer Reports* is worthless. *Smart Shopper* magazine is the ratings magazine to follow." Then *Consumer Reports* offers inconsistent advice about toasters. For, on the one hand, it says directly to buy only Toaster X. But, on the other hand, it also says to trust *Smart Shopper*, which says to buy only Toaster Y. And it is impossible to follow both pieces of advice.

In other words:

1. *Consumer Reports* says: "Buy only toaster X."
2. *Smart Shopper* says: "Buy only toaster Y."
3. *Consumer Reports* says: "Follow the advice of Smart Shopper."

(p.181) Given what *Smart Shopper* says about toasters, items 1 and 3 offer conflicting advice. So *Consumer Reports* gives conflicting advice about toasters. And a similar conflict arises in any case in which *Consumer Reports* recommends a competing magazine over itself.

Moral: no consumer-rating magazine can coherently recommend a competing magazine over itself. For the same reason, no inductive method can coherently recommend a competing inductive method over itself. Let me explain, using an argument adapted from Field (2000: 131).

Just as a consumer-ratings magazine tells one how to shop, an inductive method tells one how to respond to various courses of experience. An inductive method says something of the form: "Given course of experience E_1 , adopt *such-and-such* belief state. Given course of experience E_2 , adopt *so-and-so* belief state. Given course of experience E_3 , adopt *blah-blah-blah* belief state . . ." In other words, an inductive method puts forward a rule for responding to possible courses of experience.

One small bit of terminology: given an initial course of experience, let us say that two inductive methods are *competitors* (and that each is a *competing* method to the other) if they offer contrary recommendations about how to respond to some possible subsequent experience.

Now: it is incoherent for an inductive method to recommend two incompatible responses to a single course of experience. But that is exactly what a method does if it ever recommends a competing method over itself.

For example, suppose that inductive methods M and N offer contrary advice on how to respond to the course of experience "see lightning, then see a rainbow." In particular, suppose:

1. Method M says: "In response to seeing lightning and then a rainbow, adopt belief state X ."
2. Method N says: "In response to seeing lightning and then a rainbow, adopt belief state Y ."

(Assume that it is impossible to adopt both belief states X and Y .) But also suppose that M sometimes calls for its own rejection:

1. Method M says: "In response to seeing lightning, stop following method M and start following method N ."

Then method M offers inconsistent advice. On the one hand, it directly recommends belief state X in response to seeing lightning and then a rainbow. **(p.182)** But, on the other hand, it also says that seeing lightning should make one follow method N , which recommends belief state Y in response to seeing lightning and then a rainbow. And it is impossible to follow both pieces of advice. So method M gives incoherent advice about how to respond to seeing lightning then a rainbow. And a similar conflict arises in any case in which an inductive method recommends a competing method over itself.⁸

So: just as a consumer-ratings magazine cannot consistently recommend a competing magazine, an inductive method cannot consistently recommend a competing method. In other words, self-undermining inductive methods are incoherent. It follows that conciliatory views on disagreement are incoherent. Call this the *self-undermining problem*.

Bottom line: the self-undermining problem shows that conciliatory views on disagreement should be rejected.

6. Reply to the Self-Undermining Problem for Conciliatory Views

There is no good reply. Conciliatory views stand refuted.

7. If Conciliatory Views are Wrong, Should we Adopt an Uncompromising View Instead?

Conciliatory views get into trouble because they require one to be conciliatory about *absolutely everything*, even their own correctness. But we have **(p.183)** seen that it is incoherent to be conciliatory about absolutely everything.⁹ So conciliatory views are no good. What view should we adopt instead? We might adopt a view that is conciliatory about many matters, but not about disagreement itself. But, as noted before, such views seem to require arbitrary and ad hoc restrictions.

Alternatively, we might adopt a view that avoids the self-undermining problem without imposing special restrictions. We have already seen one such view: the stubborn view. The stubborn view avoids the self-undermining problem because, according to the stubborn view, disagreement about disagreement should not at all affect one's views on disagreement. So there is no threat of the stubborn view ever calling for its own rejection.

A more plausible view that also avoids trouble in cases of disagreement about disagreement is the **right-reasons view**.¹⁰ The right-reasons view is best explained with an example: Dee and Dum independently assess a claim, based on the same batch of evidence E. When they later find out that they came to opposite conclusions, how should they react? According to the right-reasons view, that depends on what conclusion evidence E in fact supports. For example, suppose that E supports Dee's conclusion. Then, in reaction to the disagreement, Dee should stick to that conclusion, and Dum should switch to it.

More generally, the right-reasons view says that, in the face of disagreement, one should adopt whatever view one's original evidence in fact supports. Here one's "original evidence" is the evidence that one had before finding out about anyone else's conclusions.

The right-reasons view has no special trouble accommodating disagreement about disagreement. For example: suppose that your evidence **(p.184)** strongly supports the right-reasons view, and that, as a result, you hold the right-reasons view. And suppose that you learn that a respected adviser holds a different view about disagreement. According to the right-reasons view, this should not at all weaken your confidence in the right-reasons view. In other words, according to the right-reasons view, encountering disagreement about disagreement in this case should have no effect at all on your opinions about disagreement. Other cases are similar.

Both the stubborn view and the right-reasons views are uncompromising in the following sense: each entails that, if one has correctly judged how one's original evidence bears on a claim, then just finding out that a respected adviser disagrees should not at all change one's confidence in the claim. In other words, while conciliatory views say that disagreement should *always* move one, these uncompromising views say that disagreement should *never* do so (provided that one has correctly responded to one's original evidence).

So: we have seen two ways that a view on disagreement can coherently handle cases of disagreement about disagreement. The view can be *partially conciliatory* and say that one should be moved by disagreement about some subject matters, but not about disagreement itself. Or it can be *uncompromising* and say that one should not be moved by disagreement about any topic (provided that one has correctly responded to one's original evidence). But partially conciliatory views seem to require arbitrary and ad hoc restrictions. So the undermining problem seems to favor adopting an uncompromising view (Weatherston 2007). But this is an illusion. It is not at all arbitrary for a view on disagreement to treat disagreement about disagreement in a special way. So the self-undermining problem is no evidence for uncompromising views about disagreement. Here is why.

8. The Source of the Self-Undermining Problem

It looks arbitrary for a view to recommend that one be conciliatory about most matters, but not about disagreement itself. But in fact no arbitrariness is required, for the discussion of *Consumer Reports* and inductive methods shows that it is in the nature of giving consistent advice that one's advice be dogmatic with respect to its own correctness. And views on disagreement give advice on how to respond to evidence. So, in order to be consistent, **(p.185)** views on disagreement must be dogmatic with respect to their own correctness.

In other words, the real reason for constraining conciliatory views is not specific to disagreement. Rather, the real reason is a completely general constraint that applies to any fundamental¹¹ policy, rule, or method. In order to be consistent, a fundamental policy, rule, or method must be dogmatic with respect to its own correctness. This general constraint provides independent motivation for a view on disagreement to treat disagreement about disagreement in a special way. So partly conciliatory views need no ad hoc restrictions in order to avoid the self-undermining problem. They need only restrictions that are independently motivated.

Let me illustrate the point with a *Consumer Reports* example. Suppose that, for twenty-eight years in a row, *Consumer Reports* rates itself as the No. 1 consumer-ratings magazine. A picky reader might complain to the editors:

You are even-handed and rigorous when rating toasters and cars. But you obviously have an ad hoc exception to your standards for consumer magazines. You always rate yourself No. 1! Please apply your rigorous standards across the board in the future.

This complaint has no force. The editors should reply:

To put forward our recommendations about toasters and cars *is* to put them forward as good recommendations. And we cannot consistently do that while also claiming that contrary recommendations are superior. So our always rating ourselves No. 1 does not result from an arbitrary or ad hoc exception to our standards. We are forced to rate ourselves No. 1 in order to be consistent with our other ratings.

The same point holds for views of disagreement. Just as *Consumer Reports* has good independent motivation to avoid recommending a competing magazine, so too a view on disagreement has good independent motivation to avoid calling for its own rejection. In particular, partly conciliatory views have good independent motivation for treating the case of disagreement about disagreement differently from cases of, say, disagreement about the weather. **(p.186)**

Bottom line: partly conciliatory views need no ad hoc restrictions to avoid the self-undermining problem. So the self-undermining problem does not favor uncompromising views over partly conciliatory ones. So, even though considerations arising from disagreement about disagreement refute views that are conciliatory about *every* topic, they are no evidence against views that are conciliatory about a great many topics.

References

Bibliography references:

Christensen, David (2007). "Epistemology of Disagreement: the Good News," *Philosophical Review*, 116: 187-217.

Elga, Adam (2007). "Reflection and Disagreement," *Noûs*, 41: 478–502.

Feldman, Richard (2007). "Reasonable Religious Disagreements," in Louise M. Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life*. Oxford: Oxford University Press.

Field, Hartry (2000). "Apriority as an Evaluative Notion," in Paul Boghossian and Christopher Peacocke (eds.), *New Essays on the A Priori*. Oxford: Oxford University Press, 117–49.

Kelly, Thomas (2005). "The Epistemic Significance of Disagreement," in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, i. Oxford: Oxford University Press.

Lewis, David (1971). "Immodest Inductive Methods," *Philosophy of Science*, 38/1: 54–63.

Weatherson, Brian (2007). "Disagreeing about Disagreement." TS, January. <<http://brian.weatherson.org/DaD.pdf>>

Notes:

* Thanks to Agustín Rayo, Delia Graff Fara, John Collins, Ted Sider, Brian Weatherson, David Enoch, Jacob Ross, the Corridor group, and an audience at the University of Wisconsin-Madison.

(1) I have not seen this objection in print (though see Weatherson 2007), but have encountered it repeatedly in conversation. It deserves to be put to rest.

(2) Cf. Kelly, Ch. 6, this volume.

(3) For example, according to the Equal Weight View, it is a constraint on rationality that one's probability in a disputed claim match one's prior probability in the claim, conditional on what one has learned about the circumstances of the disagreement (see Elga 2007: n. 26). But when one is certain in advance what an adviser's reaction to the claim will be, that prior conditional probability will equal one's prior unconditional probability in the claim. So the Equal Weight View is consistent with the above observation about additional opinions (that hearing an additional opinion should move one only to the extent that one counts it as independent of information one has already taken into account).

(4) I first learned of this objection from an unpublished early draft of Kelly (2005), which discusses the objection without endorsing it. Weatherson (2007) has independently raised and developed an objection of this kind.

(5) Disclosure: it will later emerge that a similar restriction *should* be imposed. But it will take real work to explain away the seeming arbitrariness of doing so.

(6) More slowly: suppose that one has view V on disagreement, and suppose that one has inductive method M . Then view V must be part of method M . So, if (given a particular course of experience) view V says to reject view V , M must (given that same course of experience) say to reject view V . That is because M says everything V says. But to reject view V is to reject M , since V is part of M . So M says to reject M . So, if V is self-undermining, then M is also self-undermining.

(7) The *Consumer Reports* analogy is adapted from Lewis (1971: 55).


(8) It might be thought that some conciliatory views on disagreement avoid this problem because they do not entirely call for their own rejection. Rather, they merely call for their own *partial* rejection. For example, consider a case in which someone with a conciliatory view—call it C —learns about a respected friend's competing view of disagreement—call it D . The conciliatory view need not say, in this case, "Reject C and adopt D ." Instead it might say, "Become uncertain as to whether C or D is the right view on disagreement."

But even views on disagreement that call for their own partial rejection are incoherent. For notice that, when one shifts one's view about the right way to respond to disagreement, one should correspondingly shift the way one responds to subsequent disagreements. In particular, when the above subject shifts his confidence away from view C and toward view D , that should correspondingly change the inductive method he implements. It will not be as dramatic a change as if he had become completely converted to view D , but it will be a change nonetheless. In other words, even in this sort of case, view C calls for a change in inductive method. And for certain choices of view D , view C calls for a change to a *competing* inductive method. But now the argument in the main text applies. For that argument applies to any inductive method that recommends a competing method over itself.

(9) So it is a good thing that some authors who defend conciliation in a great range of cases stop short of advocating it across the board. For example, Feldman (2007) gives arguments that favor suspending judgment in symmetric cases of disagreement. But he claims only that suspension of judgment is required "at least for some range of hard cases" (Feldman 2007: 212). Similarly, Christensen (2007: 189) limits his endorsement of conciliation to a restricted range of cases: "I shall argue that in a *great many cases* [of peer disagreement] of the sort van Inwagen and others seem to have in mind, I should change my degree of confidence significantly toward that of my friend" (emphasis added). Even the Equal Weight View (Elga 2007) falls short of requiring conciliation about all topics. For that view takes the form of a constraint on conditional probabilities (see Elga: n.26). As a result, the view is compatible with thinking that agents should have probability 1 in certain propositions, and that no news of disagreement should reduce that probability.

(10) The right-reasons view is a simplified version of the view defended in Kelly (2005: 180).

(11) A fundamental method is one whose application is not governed or evaluated by any other method. See Field (2000: app.).



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Epistemic Relativism and Reasonable Disagreement

Alvin I. Goldman (Contributor Webpage)

DOI:10.1093/acprof:oso/9780199226078.003.0009

Abstract and Keywords

This chapter argues that there is a plausible though non-standard conception of epistemic relativism under which relativism is compatible with objectivism or absolutism. The key element of this conception is that people in different communities can justifiably accept different principles about reasoning. As a result, people with the same first-order ('material') evidence for a proposition can have divergent but reasonable attitudes toward it.

Keywords: disagreement, epistemic relativism, objectivism, principles of reasoning

1. Introduction

Two active* topics in current epistemology are epistemic relativism and the reasonableness of disagreement between equivalently positioned agents. These topics are usually treated separately, but I will discuss them in tandem because I wish to advance a new conception of relativism that bears on the issue of reasonable disagreement.

I begin with some familiar conceptions of epistemic relativism. One kind of epistemic relativism is *descriptive pluralism*. This is the simple, non-normative thesis that many different communities, cultures, social networks, and so on endorse different epistemic systems (E-systems)—that is, different sets of norms, standards, or principles for forming beliefs and other doxastic states.

Communities try to guide or regulate their members' credence-forming habits in a variety of different—that is, incompatible—ways. Although there may be considerable overlap across cultures in certain types of epistemic norms (for example, norms for perceptual belief), there are sharp differences across groups in other types of epistemic norms. **(p.188)**

What about the normative status of these different E-systems? Is one of them right and are the rest of them wrong from an objective or absolute point of view? Are some “more right” than others? Descriptive pluralism takes no stand on this issue, but epistemologists generally want this normative issue resolved. A second brand of relativism, *nihilistic relativism*, does take a stand. It augments descriptive pluralism in holding that there is no objective right or wrong in this matter. As the sociologists of knowledge Barry Barnes and David Bloor (1982: 27) express the matter: “For the relativist there is no sense attached to the idea that some standards or beliefs are really rational as distinct from merely locally accepted as such. [The relativist] thinks that there are no context-free or super-cultural norms of rationality . . .” Philosophers would express the view by saying that there is *no fact of the matter* about which community is (“objectively” or “absolutely”) right.

Standing in contrast with nihilistic relativism is *epistemic objectivism*. This view holds that there is objective rightness in matters of epistemic norms, standards, or principles. Epistemic objectivists characteristically hold that there is a *uniquely* correct E-system and all systems incompatible with this one are wrong. Alternatively, an objectivist might hold that E-systems can be ordered by the (objective) binary relation of *being at least as correct as*. Such an ordering does not entail the existence of a uniquely correct E-system, because two or more non-equivalent systems might tie for the most correct. Also, there might be infinitely many systems for each of which there is a more correct one. For purposes of the present chapter, however, I will make the simplifying assumption that, if objectivism is true, there is a uniquely correct E-system.¹

I shall propose a new form of relativism that strikes a compromise between nihilistic relativism and objectivism. I call it *objectivity-based relativism*. As its name suggests, this form of relativism presupposes the truth **(p.189)** of epistemic objectivism; there is nothing nihilistic about it. Nonetheless, it manages to preserve some of the pluralism associated with relativism. Moreover, I will argue that objectivity-based relativism allows the possibility that two people can reasonably disagree about a given proposition even when they have equivalent evidence vis-à-vis that proposition.

I assume a close link between epistemic objectivism and the status of being (objectively) justified or unjustified with respect to beliefs and other doxastic states. The idea is that a belief or another doxastic state is justified or unjustified so long as it conforms or fails to conform to what is prescribed by the correct E-system, given the subject's evidence. Since mainstream epistemologists generally assume that beliefs are objectively justified or unjustified, the truth or falsity of epistemic objectivism is a critical issue. The truth of objectivism is also important to the prospects for reasonable disagreement, because such prospects depend on whether two evidentially equivalent people can *each* be objectively justified if they hold conflicting doxastic attitudes. Objectivism, understood as entailing a uniquely correct E-system, seems to imply the impossibility of reasonable (that is, justified) disagreement.

Further to situate the discussion, I move to two recent pieces of epistemology: Richard Feldman's paper "Reasonable Religious Disagreements" (2007) and a chapter on epistemic relativism from Paul Boghossian's book *Fear of Knowledge* (2006).

One of the two main questions Feldman poses is the following:

(Q1) Can epistemic peers who have shared their evidence have reasonable disagreements?

To say that two people have a *disagreement*, according to Feldman, is to say that one of them believes a certain proposition and the second disbelieves it. To say that two people have a *reasonable* disagreement is to say that each is *justified* in holding his or her belief (or disbelief). To say that people are *epistemic peers* is to say that they are roughly equal with respect to intelligence, reasoning powers, background information, and so on. People have shared their evidence about a topic when they have had a full discussion of the topic and have not withheld relevant information.

Let us slightly amend Feldman's formulation. Instead of confining disagreement to cases of one person believing a proposition and another disbelieving it, let the term "disagreement" apply to any case of two **(p.190)** people holding *contrary*, or *incompatible*, credal attitudes toward the same proposition. This includes one person believing the proposition and the other suspending judgment. And, instead of restricting the range of doxastic attitudes to the tripartite categories of belief, disbelief, and withholding, let us include graded beliefs or subjective probabilities among the set of categories, either point probabilities or interval probabilities (that is, partial beliefs that are somewhat fuzzy). Finally, assume that not only belief but any doxastic attitude can exemplify the property of being justified or unjustified.

The core of Feldman's paper is his defense of a negative answer to (Q1). A crucial element in this defense is what he calls "The Uniqueness Thesis:"

This is the idea that a body of evidence justifies . . . at most one attitude toward any particular proposition. As I think of things, our options with respect to any proposition are believing, disbelieving, and suspending judgment. The Uniqueness Thesis says that, given a body of evidence, one of these attitudes is the rationally justified one. (Feldman 2007: 205)

Feldman's appeal to the Uniqueness Thesis leads directly to the question of nihilistic relativism's viability, because, if nihilistic relativism is true, there is no uniquely correct system of epistemic norms. And, if there is no uniquely correct system of norms, there is no guarantee that the Uniqueness Thesis is correct. Perhaps two or more different systems of epistemic norms are equally legitimate. One implies that a given body of evidence makes doxastic attitude D vis-à-vis proposition P rationally justified, whereas another implies that the same body of evidence makes an incompatible attitude D* vis-à-vis P rationally justified. Thus, if the Uniqueness Thesis is false, Feldman's argument for the impossibility of reasonable disagreement falls through.

However, let us examine relativism more circumspectly, looking at Boghossian's treatment of the topic in chapter 5 of *Fear of Knowledge*. Boghossian formulates epistemic relativism as the conjunction of three theses (the second of which is abridged here):

(R1) There are no absolute facts about what a particular item of information justifies. (Epistemic non-absolutism)

(p.191)

(R2) Epistemic judgments of the form "E justifies belief B" express the claim: "According to the epistemic system C, that I, S, accept, information E justifies belief B." (Epistemic relationism)

(R3) There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others. (Epistemic pluralism) (Boghossian 2006: 73)

Both R1 and R3 are good formulations of the standard version of relativism I shall consider. It is fundamentally the thesis that there are no objective or absolute facts that make an epistemic system right or correct. If the justifiedness or unjustifiedness of beliefs and other doxastic states is linked in the indicated way to the objective rightness or correctness of a unique system of epistemic norms (E-system), then, if relativism is true, no objective status (for example, truth or falsity) attaches to statements that a particular doxastic state is justified or unjustified, reasonable or unreasonable. Thus, epistemic relativism seems to be equivalent to epistemic nihilism.

I am uncertain about this interpretation of the conjunction of R1, R2, and R3, because thesis R2 proffers a construal of justification statements that seems inconsistent with nihilism. Epistemic relationism says that ordinary justification statements covertly refer to the epistemic system that the speaker accepts. It offers a relational translation of justification statements that ostensibly promises an escape from nihilism. Of course, Boghossian raises serious problems for relationism (a critique I will not undertake to assess²). But the relationist component of relativism presented by R2 seems to be non-nihilistic. For this reason, R2 does not mesh so well with R1 and R3, in my view. So I am cautious about saying, unqualifiedly, that (epistemic) relativism is a form of nihilism. However, I do not myself wish to use relationism as a partial specification of E-relativism. Under my preferred construal, *standard* E-relativism is indeed a form of nihilism.

How does epistemic relativism, as defined by R1 and R3, relate to Feldman's Uniqueness Thesis and to the dispute over reasonable disagreement (among peers with shared evidence)? As defined by R1 and R3, E-relativism clashes with the Uniqueness Thesis. Moreover, as a species of nihilism, **(p.192)** E-relativism undercuts the entire dispute about reasonable disagreement. Agents who disagree in their attitudes toward a given proposition are not objectively *unreasonable* because, without an objectively correct E-system, their attitudes cannot be assessed as objectively *unreasonable*. At the same time, there can be no assessment of their attitudes as objectively reasonable or justified. So, as would be expected under nihilism, the entire issue simply melts away.

Is there any respectable form of non-nihilistic relativism, and what would such a form of relativism imply about the dispute over reasonable disagreement? A chief aim of this chapter is to articulate a form of non-nihilistic relativism and explore its ramifications. First, however, I shall advance a brief defense of reasonable disagreement unconnected with relativism. This defense focuses on matters that need to be settled before discussing relativism—that is, how to conceive of E-systems and their connection to justifiedness or reasonability.

2. Epistemic Systems, Doxastic Categories, and Reasonable Disagreement

An epistemic (E-)system is a system of rules or norms directed at doxastic attitudes or choices. The norms in question presumably take roughly the following form: “If an agent has such-and-such evidence pertinent to proposition P, or possesses such-and-such prior beliefs, or undergoes such-and-such experiences or cognitive processes (perceptual, memorial, or reasoning processes), then doxastic attitude D is the appropriate attitude for the agent to hold vis-à-vis P.” With respect to such systems of norms, we can formulate the following linkage principle L:

(L) Agent A is justified in holding doxastic attitude D vis-à-vis proposition P iff A's total evidence vis-à-vis P (e.g., antecedent beliefs, experiences, and/or cognitive processes relevant to P) is such that the objectively right epistemic system implies that D is the appropriate attitude for A to adopt vis-à-vis P; in other words, iff A's holding D conforms to the right epistemic system.³

(p.193) The formulation of principle L links justifiedness not to any random E-system but to a *right* E-system, because there are indefinitely many possible E-systems, and conformity with an arbitrary system does not confer genuine, objective justifiedness. Only a right epistemic system has the appropriate connection with objective justifiedness or reasonability (see Goldman 1986: chs. 4-5). One question here is what “appropriate” should mean? Should it be construed as *permission* or *prescription*? A number of writers, myself included (Goldman 1986), opt for the permission construal. The prescription construal might seem to bias the landscape against the possibility of reasonable disagreement. For various theoretical reasons, however, I am going to adopt the prescription construal. As we shall see, this does not unduly prejudice the case against reasonable disagreement.

A crucial question is: what makes an E-system correct, right, or best? What is the ground, rationale, or criterion that confers such a status on an E-system? The question is not whether or how we could *tell* which E-system is right; this is an epistemological question.⁴ The question is a metaphysical one, about the constitution or ground of epistemic rightness. There are various possible approaches, and, although I will not defend any such approach in detail, we should at least get a feel for some of the alternatives to persuade ourselves that the notion of such a ground or rationale is not a mere chimera. **(p.194)**

One family of approaches is externalist, roughly reliabilist. Here is a specimen of this approach, a reliabilist criterion of system superiority or comparative goodness that might induce a uniquely correct E-system.

(RCSS) Epistemic system E is better than epistemic system E* iff belief-forming practices that conform to E would produce a higher proportion of true beliefs than belief-forming practices that conform to E*.

One obvious worry here is whether RCSS would really induce a uniquely best system. A second problem is that RCSS takes account of only one type of credal state: belief. Since a general theory of justified credal states is desirable, not merely a theory of justified *belief*, should not a truth-based criterion also make use of the truth-values of graded doxastic states in addition to flat-out belief?

The latter problem might be accommodated by moving from reliability to the related notion of degree-of-truth-possession, or veritistic value (Goldman and Shaked 1991; Goldman 1999a). Just as we say that someone “possesses” the truth categorically when she categorically believes something true, so we can associate with a graded belief a *degree of truth possession* (note, not a degree of truth) as a function of the degree of belief and the truth-value of its content. A graded belief of degree n ($0 \leq n \leq 1.0$) with respect to P is assigned a degree n of truth-possession if P is true, and a graded belief of degree n is assigned a $1-n$ degree of truth-possession if P is false. Thus, having subjective probability or credence 0.70 with respect to P yields a 0.70 degree of truth-possession if P is true and a 0.30 degree of truth-possession if P is false. And so forth. We might then propose the following truth-possessional criterion of system superiority (which, under suitable assumptions, might induce a uniquely correct system):

(TPCSS) Epistemic system E is better than epistemic system E^* iff conformity to E would produce (in the long run) a higher total amount of degrees of truth-possession than conformity to E^* would produce.

Doubtless this criterion is also open to criticism. I offer it merely as an illustration. What about internalist criteria of system superiority or system goodness? One possible internalist criterion of rightness is intuitive compellingness (**p.195**) in reflective equilibrium. If a norm is intuitively compelling, after suitable reflection, this might make it right or correct. A system of all such norms would be a uniquely correct E -system. Again I do not mean to endorse this intuition-based criterion or ground of E -system rightness. But it illustrates a species of internalist approach that some may find appealing.⁵

I turn now to the problem of reasonable disagreement, and offer a first, very simple argument for the plausibility of reasonable disagreement. In considering the relationship between psychological attitudes and epistemic prescriptions for psychological attitudes, the following mismatch can in principle arise. The minimal “width” of doxastic attitudes might be narrower, at least in many cases, than the width of the categories employed by some (correct) prescriptions. Presumably, there are psychological limits on how narrow or wide a doxastic state can be. For example, the ordinary-language category of belief does not seem to designate a maximally narrow doxastic attitude. This is why many theorists prefer to talk about gradations of belief or degrees of confidence. On the other hand, it is questionable that we can have graded beliefs as fine as point probabilities—that is too narrow in terms of psychological feasibility.⁶ But moderately fine-grained degrees of credence are certainly available.

Now it seems unlikely that correct E-norms will make doxastic prescriptions only in categories as narrow as the narrowest graded beliefs. On the contrary, for many evidential situations, correct E-norms will probably issue prescriptions in doxastic categories substantially wider than the narrowest graded beliefs. For example, the Intergovernmental Panel on Climate Change, on February 2, 2007, made projections that, they said, were “very likely,” translated as “better than 90 percent.” In its previous report, in 2001, the panel of scientists said that the confidence level for its projections was merely “likely,” translated as “66 to 90 percent.” Presumably, associated with each of these confidence levels was a (tacitly) prescribed doxastic attitude interval, an interval within which a correct doxastic attitude should fall. Such prescriptions, however, leave considerable leeway. If this is the right mold for correct epistemic norms (especially where the evidence is far from probative), different choices of doxastic states will each comply with the norms. Two people can have different (that is, contrary) graded **(p.196)** beliefs within this interval—for example, one around 70 percent and one around 85 percent—yet each would conform to the norm. The difference in graded belief would constitute genuine (albeit mild) disagreement. The disagreement would be reasonable because each of the graded beliefs would conform to the norm.⁷

Roger White (2005) offers several intriguing arguments against such permissiveness. I do not find these arguments entirely compelling, but there is not space to examine them here. Later I will offer another, quite different argument for reasonable disagreement, an argument more intimately related to the distinctive themes of this chapter. So I will not probe any further into this initial argument for reasonable disagreement.

3. A Different Conception of Relativism: Objectivity-Based Relativism

As previously indicated, I want to define a species of relativism that can coexist with epistemic objectivism. Objectivism says that there is a uniquely correct E-system such that, for any proposition P and set of evidential circumstances, it prescribes to anyone with that evidence a doxastic attitude toward P within some interval. Such a prescription holds universally for all agents, whatever their community, culture, context, historical niche, and so on. This is because the system is presumed to be objectively—hence universally—right. Let us assume the truth of objectivism and call the objectively right system “SYS.” SYS's being right does not entail that anybody in any culture or context is *justified in believing* that SYS is right. In general it does not follow from the truth of an arbitrary proposition P that everybody or anybody is justified in believing P. Some truths are hidden; they do not automatically generate evidence of their truthfulness to all populations, or any populations. Gaining epistemic access to them may be difficult and problematic. Truths concerning E-systems are likely to be in **(p.197)** this boat. Philosophers like to think of themselves as enlightened, but it is distinctly possible that even members of the philosophico-methodological subculture fail to be justified in believing, either with respect to the complete E-system SYS, or with respect to some of its individual norms, that it is the correct E-system or correct individual norm.

Failure to be justified in believing a correct norm to be correct is not the only possibility; people might be positively justified in believing some *incorrect* E-system or E-norm to be correct. In both cases, the justifiedness in question could be *objective* justifiedness. Thus, people could hold mistaken but *objectively justified* beliefs (or weaker doxastic attitudes) about E-norms. How could this occur? Would it not require a right E-system, together with suitable evidence, to undermine itself? Is this possible? Yes. Let us elaborate a few plausible examples of such a scenario. But first let us back up and say more about the contents of plausible E-systems.

We can characterize all E-norms as source authorizations. Vision is one possible source, and a vision-based norm might be: “If it looks to an agent as if P, then (in the absence of defeating conditions) the agent should believe that P.” Another possible source is memory, for which a related norm might be: “If an agent seems to remember that Q, then (in the absence of defeating conditions) the agent should believe that Q.” Many sources will be psychological sources, like vision and memory, but some sources probably will not be—for example, testimony. A testimonial norm might be: “If a random speaker or writer testifies that P, then (in the absence of defeating conditions) the agent should believe that P.” To be sure, one could not apply a testimonial norm without relying on psychological sources to decide whether a speaker has delivered testimony that P. But that does not undercut a testimonial source as an epistemic source. The first group of examples I shall give of justifiably believing an incorrect norm is drawn from the domain of testimony. In advancing these examples, I shall presuppose the correctness of a certain genre of testimonial norm that is widely accepted by epistemologists. However, I will not presuppose any highly specific testimonial norm as *the* correct one in its territory.

It is common in many cultures for children to be told by their elders that specific sources should be trusted as guides to belief.⁸ In religious **(p.198)** communities, young children are taught that a certain scripture should be trusted as a guide to the truth about religious matters and historical events, possibly including such things as the age of the Earth and when various species came into existence. The same scripture might be cited as the supreme source on moral matters. Children are in effect given E-norms with the content: “If the scripture says P, you should believe P.” In scientific educational contexts, students might be given E-norms with the content: “If scientific researchers agree on P, you should assign a high credence to P.”

Are children in such instructional contexts justified—objectively justified—in believing that such norms are correct? Whether they are so justified depends on the contents of genuinely correct E-norms. Although epistemologists do not speak with one voice about these contents, almost all believe that *generic* testimony-based norms—norms concerning testimony from arbitrary speakers—are among the right norms. Whatever the *exact* contents of generic testimonial norms, it is plausible that, when children receive religious or scientific instruction—especially early instruction, when their ability to engage in autonomous criticism is relatively weak—the instruction received from their teachers or parents renders them objectively (O-)justified in believing that the norms so transmitted belong to a correct E-system. The children are O-justified in accepting such norms. This seems especially clear if the children hear roughly the same testimony from numerous elders and no conflicting testimony, a likely scenario in many communities both historical and contemporary.

This chapter does not aim to resolve which of the specific norms just illustrated are correct and which are not. But many of the norms conflict with one another in such a way that *not all* could belong to a uniquely right E-system. For example, students in different contemporary American educational systems are exposed to different teachings about the epistemic force of evolutionary science. Those instructed in a modern biology curriculum are taught to assign high credence to whatever evolutionary science says about the world. Although this material might not be explicitly formulated in the form of E-rules, such an implication would be present. By contrast, students taught in fundamentalist schools (especially private **(p.199)** ones, with an Intelligent Design mission) are taught to be skeptical about whatever evolutionary science says. They are encouraged to accept E-norms urging low levels of credence in evidence of that kind. Given their respective exposures to the testimony of their teachers and the presumed correctness of *generic* testimonial norms, all students would be justified in believing the recommended E-norms to be correct. But the two E-norms concerning evolutionary science clearly conflict with one another, so they cannot both belong to the uniquely correct E-system. So we have at least one type of case in which an E-norm is justifiably believed to be right but does not in fact belong to a right E-system.⁹

The foregoing examples, however, feature *derivative* norms as opposed to *fundamental* ones. Even if one grants that generic testimonial norms are fundamental, surely testimonial norms prescribing trust in particular texts or authorities must be derivative norms. Readers might concede that false though justified beliefs can be held with respect to derivative norms but resist the idea that the same holds for fundamental ones.

But there is no difference here. Even a fundamental norm can have its apparent authorizing credentials amended by experience. I do not mean that the norm ceases to be **(p.200)** correct, only that some cognizer ceases to be justified in believing it to be correct. A standard example of a fundamental norm is “If it looks to you as if P, then (in the absence of defeaters) you should believe that P.” Now imagine a scenario in which someone hears credible testimony to the effect that the visible world is a sham or delusion, so that vision should not be trusted (compare the movie *The Matrix*). Is this not a case in which the epistemic status of a fundamental norm for an individual is affected by his experience? (This assumes a reliabilist or veritistic grounding of E-norm correctness.) The same thing can transpire for a norm of reasoning, which might seem to be a paradigmatic fundamental norm. Choose your favorite methodological norm: a statistical norm like the Neyman–Pearson method, the chi-square method, or even a Bayesian norm. Such norms are often the subjects of serious debate and critique—in statistical-theory circles or philosophical circles. If a novice hears a lecture from a well-certified theoretical statistician or philosopher of statistics that mounts a compelling critique of such a norm, the hearer could well be justified in reducing or moderating her credence in the correctness of the norm. So the justificational status of even a fundamental norm can be amended by application of other norms, ultimately by appeal to the ground of E-system rightness.

Let us see how this might work for probabilistic norms, even if the posited ground of rightness is intuitive appeal in reflective equilibrium. Assume that correct norms for probabilistic reasoning are associated with the standard probability calculus. One such norm is the prescription not to assign a higher probability to a conjunctive event than to one of that event's conjuncts. Thus, in Tversky and Kahneman's well-known Linda example (1983), the norm would imply that one should not believe that "Linda is a feminist and a bank teller (F&T)" has a higher probability than "Linda is a bank teller (T)." Could anybody be justified in believing (mistakenly) that a different, incompatible norm is the correct one? Yes. Tversky and Kahneman's results showed that naive intuitions tend to be driven by the "representativeness," or resemblance, heuristic. This heuristic leads subjects to judge the F&T event more probable than the T event, because Linda more closely resembles a prototypical feminist bank teller than a prototypical bank teller. When subjects were explicitly presented with two arguments, one using the conjunction rule and the other using the resemblance rule, 65 percent of the subjects found the resemblance argument more convincing than the conjunction-rule **(p.201)** argument. Apparently, resemblance considerations are intuitively more compelling for naive subjects than the conjunction rule. If we now assume that undefeated intuitions provide justification, then naive subjects who have not been tutored in probability theory may well be justified in believing that the resemblance norm is correct, even though it conflicts with the probability calculus. The ground of rightness we are considering, however, makes norm rightness a function of intuitive appeal *in reflective equilibrium*. In the present case, this ground would be applied as follows. Once people reflect systematically on matters of probability, they will come to appreciate—and find intuitively compelling—the appropriateness of norms based on the probability calculus. So we have a case in which some people—the wholly untutored ones—are justified in accepting norms that are not genuinely correct under the posited ground of correctness.

We have been considering cases in which a person is objectively justified (O-justified) in *believing* of a certain norm that it is correct or incorrect. But we might also be interested in cases in which someone is O-justified in having a graded belief rather than a full belief in a norm's correctness or incorrectness. Although her evidence might not support full belief in norm N's correctness, it might support a credence of, say, 0.60. Our framework permits justifiedness not only for full beliefs but for all grades of credence.

We can now state the central theses of the new form of relativism I wish to consider: objectivity-based relativism. These theses can be formulated as follows:

(OBR) There is a uniquely correct E-system that governs the objective justifiedness and unjustifiedness of people's doxastic attitudes. However, people occupy different evidential positions vis-à-vis this system and other candidate E-systems. Hence, the objective justificational status of different people vis-à-vis different E-systems is varied rather than uniform. Some people are objectively justified in believing certain E-norms and E-systems to be correct; others are objectively justified in believing other E-norms and E-systems to be correct. Similarly for attitudes other than full belief toward E-norm-related propositions.

(p.202) Objectivity-based relativism is very different from nihilistic relativism, and also—by my lights—fairly attractive.¹⁰ It has the virtue of accommodating an important intuition that actuates many proponents of E-relativism, the intuition that differences in intellectual procedure found in diverse cultures, communities, and historical periods do not reflect wholesale irrationality or epistemic depravity. There is something epistemically legitimate about divergent choices of procedures. Objectivity-based relativism captures this intuition by allowing members of epistemically diverse cultures to have objective justification (O-justification) for different beliefs about intellectual norms. In virtue of this norm-justification, they may also enjoy a distinct but significant justificational status for their garden-variety beliefs (beliefs about ordinary matters rather than E-norms). This is the status of being *O-justified in believing that they are O-justified in believing P*. When they use their adopted E-norms to form beliefs in garden-variety propositions, these beliefs will often fail to be O-justified. Nonetheless, they may be *iteratively O-justified*: the people are justified in believing that their beliefs are justified.

Suppose Amanda is O-justified in believing norm X to be a correct E-norm. Furthermore, given Amanda's evidential circumstances, norm X authorizes her to believe proposition P. Then she is O-justified in believing that she is O-justified in believing P.¹¹ However, second-order O-justifiedness does not entail first-order O-justifiedness.

(Non-entailment) $J_o[J_o(P)] \neq J_o(P)$

Perhaps norm X is actually incorrect, although Amanda is O-justified in believing it to be correct. Objectivity-based relativism does not imply that all norm-systems are equally right in the sense of being equally capable of conferring first-order justifiedness.

It may be helpful here to flag the distinction between *propositional* and *doxastic* justifiedness. A person is doxastically justified in having attitude D toward P if she actually has D and it is justified. A person is propositionally justified in having attitude D toward P if her epistemic position is such **(p.203)** that D is the proper attitude to adopt toward P—whether or not she actually adopts it. Arguably, iterative justifiedness makes better sense when interpreted in propositional rather than doxastic justificational terms. This is because comparatively few individuals form explicit beliefs about the justificational status of their own (first-order) attitudes. Only fairly reflective minds contemplate this sort of thing. Nonetheless, even if Jerome does not actively wonder whether he is justified in believing P, and hence does not come to any belief (or other opinion) on the subject, his evidential condition might *entitle* him to believe that he is so justified. Hence, he is propositionally justified in believing that he is justified even if he is not doxastically justified. The truth of this iterative-justificational proposition might well interest epistemologists.

4. Objectivity-Based Relativism and Reasonable Disagreement

I turn now to the implications of OBR for the reasonable disagreement controversy. The implications I will extract are fairly limited in scope, because objectivity-based relativism bears on the reasonable disagreement issue only from a single restricted angle, that of iterative justifiedness. The analytical framework presented here does not provide the resources for a full-bore attack on the reasonable disagreement problem, because it takes no stance on the contents of a right E-system. Without specifying such contents, it is hard to draw firm conclusions about the doxastic moves an agent should make if she learned various things about her peers, such as the fact that some of them disagree with her. Should she stick to her guns in believing P? Should she “split the difference” with them? These questions cannot be adequately answered without identifying the right E-system—at least the general contours of such a system. Although I will not tackle these central issues, I will use our broader analytical framework (mainly, the linkage principle) to reveal a connection between iterative justifiedness and reasonable disagreement.

An objectivist framework for E-system rightness requires any two people who have the same total evidence vis-à-vis P to take the same attitude toward P—at least if we ignore permissible differences within the prescribed attitude, as discussed in Section 2. (Henceforth I ignore such differences.) **(p.204)** For both individuals to have objectively justified attitudes toward P, their attitudes must be the same. If they differ, at most one attitude can be justified. Hence, reasonable disagreement is precluded at the first-order level of justifiedness. Even if their attitudes differ, however, each might be objectively justified in believing that her attitude is (objectively) justified. In other words, disagreement among evidentially equal agents is compatible with each agent possessing *second-order* justifiedness.

Here is a scenario by which this can transpire. Amanda and Jerome have the same evidence with respect to P but different evidence about E-system correctness. In virtue of this evidence, Amanda is O-justified in believing system E to be correct, whereas Jerome is O-justified in believing E* to be correct. Finally, the attitude required by E toward P (given the specified evidence) is incompatible with the attitude required by E*. Thus, Amanda is justified in believing that she is justified in adopting attitude D toward P, whereas Jerome is justified in thinking that he is justified in adopting attitude D* toward P, where D and D* are incompatible. At the first-order level of justification such a difference in attitude implies that at least one of them is unreasonable, but at the second-order level of justification both can be reasonable—that is, iteratively justified.

The preceding sentence incorporates a crucial step in the argument: the proposal that higher-order justifiedness can ensure, or at least make a positive contribution toward, the reasonability of a first-order belief. This is despite the fact that higher-order justifiedness does not entail first-order justifiedness. Thus, the reasonability of an agent's attitude toward P is not fixed by its first-order justificational status. This point is worth marking with a new principle:

(J1 \neq R) The first-order justificational status of an attitude does not fix its (overall) reasonability; reasonability can also be influenced by higher-order justificational status.¹²

(p.205) What considerations might support this principle? And, if second- (or higher-)order justifiedness counts in fixing reasonability, how much does it count?

A first point to make is that second-order justifiedness has *some* epistemic value, indeed, substantial value. Consider an agent who (i) forms a justified but mistaken belief that system E is correct, (ii) correctly applies E's requirements to her own evidential state, and therefore (iii) selects attitude D toward P. How well does this agent proceed in epistemic terms? She clearly proceeds well at stage (i). She justifiably forms a belief that E is correct. She cannot be faulted there in epistemic terms. Similarly, how can she be faulted for the procedures she executes at stages (ii) and (iii)? In these stages, the norms she justifiably believes to be correct are applied to her evidence, and her attitude toward P is formed on the basis of this evidence. Perhaps she can be faulted for failing to obtain a *true* belief at any of these stages. But, if we assume fallibility even for objectively right E-norms—and I do assume such fallibility throughout—then truth attainment is never guaranteed by first-order, second-order, or any order of justifiedness. So why should failure to obtain the truth imply culpability? In short, when a person's belief enjoys second-order justifiedness, there is much to be said for her epistemic conduct. If epistemic conduct can be characterized as “culpable” or “non-culpable,” a belief's second-order justifiedness entitles an agent to a respectable level of non-culpability. At a minimum it makes a *contribution* toward attainment of an overall level of positive non-culpability or reasonableness. Furthermore, it is a contribution that might *trump* the epistemic culpability associated with holding a belief (or other attitude) that is first-order *unjustified*.

Some might complain that epistemic non-culpability is a rather weak status, not strong enough to imply justifiedness or reasonability. But the argument can be rephrased in terms of “propriety” of epistemic conduct. Does not an agent engage in proper epistemic conduct if she applies the correct norm-system to her evidence at stage (i) to form a belief that system E is correct? Assume that no later evidence mandates a change in this belief. Given her proper choice of system E, does she not engage in proper epistemic conduct at stages (ii) and (iii) in applying E's requirements to her P-relevant evidence and selecting attitude D? By similar steps, a different agent might properly form a belief that system E* is correct and properly **(p.206)** select a different attitude D* with respect to P, despite having the same P-relevant evidence as the first agent. Thus, epistemic peers who share the same P-relevant evidence can reasonably disagree about P, even when this involves first-order unjustifiedness on the part of at least one of them.

Here is an additional consideration to support the significance of second-order justifiedness. When considering the reasonability of someone's belief, its truth-value does not settle the issue. A false proposition can be reasonably believed. What determines a belief's reasonability is the agent's evidence (or belief-forming methods), not the belief's truth-value. The same point holds on the topic of norm correctness. The actual rightness of an E-system does not determine the reasonability of an agent's conforming to it. What is critical is the agent's evidence about its rightness. If an agent conforms her attitude to the prescriptions of a properly chosen E-system, this should be an important—perhaps decisive—element in assessing the attitude's reasonability, even if the evidence supporting that E-system's rightness happens to be misleading.

This proposal poses a problem, however. If second-order justifiedness is relevant to reasonability, why is not every order of iterative justifiedness relevant? Indeed, given what we have said, should not each higher order of iterative justifiedness be more relevant to reasonability than its predecessor? Will this not generate a vicious infinite regress, which threatens to scotch the entire enterprise of assigning determinate justificational statuses to doxastic attitudes? Each higher-order status will trump the immediately lower-order status, and, as the orders increase, they will tend to swamp first-order justifiedness entirely. Does it not become radically unclear what overall reasonability consists in, or whether it can be determinate?

The problems in this territory are not as devastating as the foregoing portents suggest. As one ascends the hierarchy, the evidence an agent possesses vis-à-vis the preceding level of iterative justifiedness rapidly becomes negligible. In fact, it may quickly become null. If the right E-system is anything like what epistemologists suppose, an agent will typically be instructed, at the $n + 1$ st level, to *suspend judgment* about the n th level of justifiedness. While justified *belief* about a lower level of iterative justifiedness can trump lower-level justifiedness, justified agnosticism, or suspension of judgment, should not have comparable trumping power. Even an infinite series of judgment suspensions will be in the same boat. **(p.207)** So the threat of higher-order justifiedness totally swamping first-order justifiedness is not so severe.

It would be helpful, no doubt, to quantify the appropriate weightings for lower-order and higher-order justificational status. How exactly do they influence overall, or *ultima facie* reasonability? Unfortunately, I do not know how to address this issue in adequate generality. Two points should suffice for present purposes. First, it is not proposed that first-order justifiedness gets “washed out” entirely by second-order justifiedness. Forming opinions in accord with an objectively right E-system is surely worth *something*, if not *everything*, in terms of justification and reasonability. The suggestion is only that first-order justifiedness can be outweighed or superseded by higher-order justifiedness. Secondly, for present purposes, we do not have to specify the precise circumstances in which higher-order justifiedness trumps first-order justifiedness. As long as this can sometimes happen, it falsifies the general principle that two people with the same (first-order) evidence vis-à-vis P cannot reasonably adopt different attitudes toward P. The foregoing considerations adequately establish that this can sometimes happen.¹³

5. Evidence and the Peer Disagreement Controversy

Participants in the peer disagreement controversy are likely to complain that I am ignoring *their* controversy, because their controversy centers on **(p.208)** cases in which people have the same evidence vis-à-vis target proposition P. It concerns cases involving epistemic peers, where the peer relationship typically includes “evidential equality” (as Christensen (2007) calls it). By contrast, my cases are ones in which people have evidential differences—that is, differences concerning the correct E-norms. So, critics might mutter, how am I contributing to the debate?

My contribution might be viewed from the following perspective. It contributes to the debate by identifying a category of evidence that bears on the reasonableness of peer disagreement but is generally ignored in the literature. Contributors to the debate typically divide the determinants of reasonability into two sectors. The first sector consists of the agents' evidence relevant to the target proposition. This evidence is usually divided into three categories: (a) evidence “directly” concerning the target proposition, (b) evidence concerning one's own epistemic competence, and (c) evidence concerning the peer's epistemic competence. The second sector consists of rules or norms that should govern their epistemic conduct. Such rules are prescriptions or permissions, which are not, strictly speaking, propositions. Hence they are not the sorts of things for which there can be evidence; the three types of evidence in the first sector exhaust the evidence relevant to peer disagreement. However, I am pointing out an additional type of proposition with respect to which evidence might diverge. This is a proposition of the form “Norm X is a correct norm (and applies to the present doxastic choice).” Two agents can have different bodies of evidence that bear on norm correctness and are relevant to the reasonability of their respective attitudes.¹⁴ So here we highlight a species of evidence—*norm* evidence, we might call it, as contrasted with *material* evidence—that is generally ignored in the literature.

The peer disagreement literature tends to miss this point because it presumes that justifiedness or reasonability is conferred by *de facto* norm correctness. Correct norms, not an agent's evidence about the correct norms, set the standard for epistemic conduct. The issue of norm evidence simply is not raised. I am arguing that norm evidence *is* among the determinants of reasonability. Where two agents are equal with respect to material evidence but differ with respect to norm evidence—though the **(p.209)** correct norm-system stays fixed—it is legitimate for their attitudes toward a given proposition to diverge.

If this is my view, am I not siding with those espousing the maxim “No reasonable disagreement without evidential difference?” Yes, that is a fair characterization of my position—if *all* categories of evidence are included. But, if attention is restricted to material evidence (“sector 1” evidence), as it usually is, this maxim does not characterize my position. In cases where evidential equality extends only to *material* evidential equality, there is room for reasonable disagreement.¹⁵

As acknowledged at the beginning of Section 4, this chapter does not develop a comprehensive approach to peer disagreement. However, let me identify some other contours of the topic and explain why I remain silent on many of them. I will also identify some problems with existing treatments, especially concerning the nature of evidence.

My approach to peer disagreement, it might be said, embraces a *synchronic* perspective. If two agents are evidential equals with respect to P at time t, can they reasonably differ in their attitudes at t toward P? A more standard perspective is a *diachronic* one. The diachronic question is how an agent should change her opinion vis-à-vis P over time. It focuses on the following problem: at time t an agent forms an opinion vis-à-vis P in ignorance of a certain peer's opinion. At a later time t*, the agent learns that the peer, despite being an evidential equal, holds a different opinion. How (if at all) should the agent revise her opinion? Here is David Christensen's diachronic formulation of the problem, which is fairly representative: "How should I react when I discover that my friend and I have very different beliefs on some topic? . . . Should my discovery of her differing degree of belief in P lead me to revise my own confidence in P?" (Christensen (p.210) 2007: 188). We might call this the *peer responsiveness* formulation of the problem.

The framework I employ here lacks sufficient tools for a detailed analysis of the peer responsiveness problem. Ignoring the arguments of Section 2, the framework implies that under complete evidential equality (including norm-evidence equality) two people cannot reasonably have differing opinions. But suppose evidential equality is not complete. We can still ask what doxastic choices are expected of peers when they discover their disagreement? Must their degrees of belief converge? Must the mode of convergence involve "splitting the difference?" To tackle these issues we need more than framework principle L. We need a correct and detailed E-system. However, a formulation, defense, and application of such an E-system is beyond the scope of this chapter. Nonetheless, we can further contribute to the debate in two ways: first, by identifying weaknesses in certain treatments of evidence and evidential equivalence, and, secondly, by pinpointing the impact of alternative theories of evidence.

Contributors to the peer disagreement topic tend to assume that evidential equality is a readily producible scenario, that two people can make themselves evidential equals by simply "sharing" relevant evidence with one another. Feldman (2007) makes much of this procedure. The assumption seems to be that, if one person verbally communicates what he regards as his (relevant) evidence, a hearer acquires the same evidence as the speaker. If the second person reciprocates, both will have shared their total information with their opposite number and they will be evidential equals. This assumption, however, involves an unnoticed pun on the word "share." Sharing evidence in the sense of communicating the content of an evidential state does not necessarily imply that the hearer shares—in the sense of possesses—the same evidence as the speaker. This is well illustrated by the following example, provided by Apolonio Latar (2007).

Suppose Billy is accused of committing a certain crime and his friends have weighty evidence that he did it. They know, for example, that he threatened to commit a crime of that very description the day before it happened. In fact, Billy did not commit the crime. However, he was alone in his room when somebody else was committing it, so he cannot prove to his friends that he did not do it. He clearly remembers not having done it (he recalls not being near the crime scene, and so on), and this vivid memory is excellent evidence *for him*. But Billy cannot literally transmit **(p.211)** this memory to his friends (for example, by duplicating this portion of his brain state in them). All he can do is verbally report its content. Even if his friends believe his report, this does not make them evidentially equal to Billy vis-à-vis the criminal accusation. They believe he is innocent, but they do not have personal recall to support that belief. Moreover, as Latar points out, even if the friends accept Billy's report, they may not accord it as high a degree of confidence as he does. So the proposition will have less evidential power for them than it does for Billy.

Feldman rightly supposes that, if one person verbally shares his evidence with another, the latter acquires evidence *of* the speaker's evidence. He therefore articulates the principle "evidence of evidence is evidence." What he might mean by this is that, if Smith truthfully reports evidence Q concerning P, this evidence is also acquired by the hearer.¹⁶ This is not quite right. Hearing such testimony may give the hearer default justification for believing Q, but such default justification can be defeated by other information in the hearer's possession. In that case, Q does not qualify as an item of evidence for the hearer. Furthermore, even if there is no such defeat, the hearer does not necessarily acquire the same evidence possessed by Smith. Smith's saying that he had a certain visual experience, for example, does not reproduce in the hearer the same visual experience, with its full evidential load. So the hearer does not acquire the same evidence for P as Smith has.

Furthermore, it is impossible to convey to others all the subtle strands of evidence one harbors, or has harbored, for one's opinions. For example, failing to observe any counterexamples to a certain hypothesis may justify one's acceptance of it—at any rate, if there is a high likelihood that one would observe such counterexamples if the hypothesis were false. But the evidential "omissions" that collectively constitute this (past) support tend not to be stored in memory and are not readily retrieved if one is asked to defend one's belief.¹⁷ Finally, one often forgets even past observations **(p.212)** that play a lively causal role in belief acquisition. These now-forgotten observations are relevant to the current justificational status of a belief that has been preserved over time, but they are not available for "sharing" when asked for one's evidence. Thus, what speakers manage to communicate when asked for their reasons rarely approximates the whole of their relevant evidence.

The peer disagreement literature assesses the extent to which people should defer to others with similar evidence and cognitive competence but differing opinions. It asks what it would be *rational* or *reasonable* for people to do in such cases. I have offered some conclusions about what reasonability requires based on a very general conception of epistemic justifiedness. It must be emphasized, however, that such general conclusions cannot be applied to concrete cases without first settling the question of what counts as evidence. This in turn depends on the contents of the correct norms.

Let me illustrate this point with the help of Kelly's example of a mathematician who is initially confident he has a proof of a certain theorem but whose colleagues deny that it is a genuine proof (see n. 15). Kelly characterizes the case as one of evidential equality, because each mathematician has surveyed exactly the same evidence, presumably, the steps of the proof as written on paper. Another epistemologist, however, might hold that the parties are unlikely to have the same evidence in Kelly's example. Suppose that, although all the steps in the proof are correct (as Kelly supposes), each colleague, while examining the proof, seems to detect a mistaken step. The proof creator, while reviewing his proof, has no such experience of seeming to spot an error. Then under some conceptions of evidence the parties do not have the same evidence. Consider, for instance, Michael Huemer's principle of "phenomenal conservatism," which may be classified as a principle of evidence: "If it seems to S as if P, then S thereby has at least prima facie justification for believing that P" (Huemer 2001: 99). If this evidence principle is correct, the mathematicians do not all have the same evidence. Each colleague has evidence of the proof's containing a mistaken step, whereas the proof creator has no such evidence. In general, comparatively "subjective" conceptions of evidence will tend to produce **(p.213)** fewer cases of exact evidential equality than comparatively "objective" conceptions.¹⁸

Conceptions of evidence will not coincide across all E-systems. A conception of evidence is implicitly specified by the set of antecedents of an E-system's (conditional) prescriptions. These antecedents fix what the system considers to be evidence, and they will not generally be the same across E-systems. An upshot of this is that we cannot settle questions about evidential equality without settling questions about the properties of a right E-system. This task outstrips the compass of the present chapter.

6. Conclusion

This chapter has advanced two theses. The first thesis is that there is a plausible though non-standard conception of epistemic relativism under which relativism is compatible with objectivism or absolutism. The crucial point underlying this thesis is that, even if there is a uniquely right system of E-norms, people in different communities can justifiably (though not correctly) accept different E-systems as right. The second thesis consists in a moral to be drawn for the problem of reasonable disagreement. Once we distinguish first-order and second-order justifiedness, we find that two people with the same (“material”) evidence for proposition P can have contrary attitudes toward P that are both second-order justified. Since second-order justifiedness is (or can be) as important a determinant of reasonability as first-order justifiedness, these divergent attitudes can both be reasonable.

References

Bibliography references:

Barnes, B. and Bloor, D. (1982). “Relativism, Rationalism and the Sociology of Knowledge,” in M. Hollis and S. Lukes (eds.), *Rationality and Relativism*. Cambridge, MA: MIT Press, 21–47. **(p.214)**

Boghossian, P. (2006). *Fear of Knowledge*. New York: Oxford University Press.

— (2007). “The Case against Epistemic Relativism: Replies to Rosen and Neta,” *Episteme: A Journal of Social Epistemology*, 4/1: 49–65.

— (2008). “Epistemic Rules,” *Journal of Philosophy*, 105/9: 472–500.

Broad, C. D. (1985). *Ethics*, ed. C. Lewy. Dordrecht: Martinus Nijhoff.

Christensen, D. (2007). “Epistemology of Disagreement: the Good News,” *Philosophical Review*, 116: 187–218.

Feldman, R. (2007). “Reasonable Religious Disagreements,” in L. Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life*. New York: Oxford University Press, 194–214.

Goldman, A. (1979). “What Is Justified Belief?,” in G. Pappas (ed.), *Justification and Knowledge*. Dordrecht: D. Reidel, 1–23;

repr. in A. Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press, 1992, 105–26.

— (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

— (1999a). *Knowledge in a Social World*. Oxford: Oxford University Press.

— (1999b). "Internalism Exposed," *Journal of Philosophy*, 96: 271–93.

— and Shaked, M. (1991). "An Economic Model of Scientific Inquiry and Truth Acquisition," *Philosophical Studies* 63: 31–55;

repr. in A. Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press, 1992, 227–54.

Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham, MD: Rowman and Littlefield.

Kelly, T. (2005). "The Epistemic Significance of Disagreement," in T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, i. Oxford: Oxford University Press, 167–96.

Lackey, J. (2006). "It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony," in J. Lackey and E. Sosa (eds.), *The Epistemology of Testimony*. Oxford: Oxford University Press, 160–89.

Latar, A. (2007). "Shared Evidence and Reasonable Disagreement." Department of Philosophy, Rutgers University, undergraduate essay.

Neta, R. (2007). "In Defense of Epistemic Relativism," *Episteme: A Journal of Social Epistemology*, 4/1: 30–48.

Pettit, P. (2006). "When to Defer to Majority Testimony—and When Not," *Analysis*, 66: 179–87.

Rosen, G. (2007). "The Case against Epistemic Relativism: Reflections on Chapter 6 of *Fear of Knowledge*," *Episteme: A Journal of Social Epistemology*, 4/1: 10–29. **(p.215)**

Tversky, A., and Kahneman, D. (1983). "Extensional versus Intuitive Reasoning: the Conjunction Fallacy in Probability Judgment," *Psychological Review*, 90: 293–315.

White, R. (2005). "Epistemic Permissiveness," in J. Hawthorne (ed.), *Philosophical Perspectives*, xix: *Epistemology*. Malden, MA: Blackwell Publishing, 445–59.

Notes:

* Earlier versions of this chapter were presented at Northern Illinois University, Princeton University, and the Bled (Slovenia) epistemology conference of 2007. I am grateful to members of these audiences for very helpful discussions. A non-exhaustive list of those whose comments influenced the final product would include Elizabeth Harman, Gil Harman, Tom Kelly, Jennifer Lackey, Ram Neta, Baron Reed, Gideon Rosen, Bruce Russell, and Holly Smith.

(1) A plausible terminology to adopt here would distinguish *weak* and *strong* objectivism. Weak objectivism would be the view that all pairs of E-systems have an objective ranking in terms of comparative (E-) goodness, and strong objectivism would be the view that some E-system is uniquely best in terms of such a ranking. In this terminology, our focus here is strong objectivism. However, the phrase 'strong objectivism' might have unintended connotations. It might suggest a highly rigid or constraining brand of objectivism, which issues extremely fine-grained prescriptions for all evidential scenarios. This would be the opposite of what Roger White (2005) calls "epistemic permissivism." My conception of strong objectivism, however, does not have this implication. A uniquely best E-system might be a fairly loose or permissive system. See Section 2 below.

(2) For additional discussion of possible solutions to the problems of relationism, see a mini-symposium on Boghossian (2006) consisting of Boghossian (2007), Neta (2007), and Rosen (2007).

(3) Notice that the linkage principle only requires doxastic attitude D to *conform* to the right norm system in order to be justified. It does not require the agent to *follow* the system's rules in arriving at D. In particular, it does not require the agent mentally to represent the relevant norms or to be mentally guided by them. For example, assume that the correct E-system includes perceptual and memory norms, which "approve" of an agent's holding perceptual or memory beliefs under specified circumstances. A non-reflective agent, who does not mentally represent these norms, can nonetheless justifiably hold such beliefs as long as she *conforms* with the norms. It would be an excessive demand to place on justifiedness to require rule-following "all the way down." See Boghossian (2008), who argues that systematic rule-following involves a vicious regress.

(4) Boghossian (2006) examines an argument in support of nihilistic relativism based on the premise that, if there is an objectively right E-system, it is possible to be justified in believing that it is right. The argument then proceeds to deny that this is possible on the grounds that such justification would involve norm circularity, because it would have to presuppose the rightness of the system. Obviously, the argument also needs the further premise that norm circularity vitiates justifiedness. A worry I would pose for this argument (different from Boghossian's criticisms) is where any proponent of such an argument would get the last premise. What justifies the premise that norm-circular arguments are justificationaly impotent? Is this premise embedded in the right E-system? Is it embedded in *every* E-system? Is it a higher-level *constraint* on any E-system? Each of these assumptions is problematic. So it is hard to see how to construct a successful epistemological argument of this sort against the existence of a uniquely correct E-system. Such reflections tilt against the cogency of any putative epistemological constraint on the existence of an objectively right E-system comparable to the constraint that it must be possible to have non-circular justification for such a system.

(5) For an illustration of one possible ramification of an intuition-based criterion, see n. 10 below.

(6) Of course, the propositional *content* of a belief might include point probabilities. But that is not pertinent to the present discussion.

(7) Even if we adhere to the tripartite scheme of belief, suspension, and disbelief on the assumption that they are the “thinnest” doxastic states available (a very implausible assumption), a right E-system could still issue prescriptions for disjunctive categories like “belief or suspension” or “disbelief or suspension.” Single-word labels could be invented for such doxastic intervals (e.g., “belension” and “disbelension” respectively). Thus, the same argument for reasonable disagreement can be presented within the tripartite taxonomy.

(8) I focus my examples on children's receipt of testimony because children have smaller stores of real-world knowledge or belief as compared with adults. This implies, among other things, that they possess fewer evidential resources to challenge the testimony of their elders, and hence fewer evidential resources to defeat the prima justifiedness that arises from receiving such testimony.

(9) Some writers on testimony might resist my conclusion on the grounds that a hearer's justifiedness in accepting a piece of testimony from a speaker depends not only on the hearer's evidence but also on the reliability of the speaker. In the cases before us, therefore, children might not be justified in believing what their elders say about the trustworthiness of a specified source, because the elders in these communications are not reliable sources. Jennifer Lackey (2006) gives an example called "NESTED SPEAKER," in which Fred has reasons to believe that Pauline is a reliable testifier about wild birds, but in fact she is not reliable on this subject. Lackey contends that, when Fred forms a belief about albatrosses based on Pauline's testimony, his belief is not justified. That is not the result of any flaw in Fred's reasons, but rather a result of Pauline's unreliability. Similarly, as both Lackey and Baron Reed have argued to me in conversation, if the elders in my example are unreliable speakers, the children are not justified in believing in the correctness of the elders-commended norm. One response to these arguments is to question the judgment that Fred's belief is unjustified. My own intuition about this case is murky, by no means clear-cut in Lackey's direction. However, let us concede the classification of the case for purposes of further argument. It is clear in our examples that, where the norm endorsed by the speakers is incorrect (in virtue of the norm's unreliability), the speakers are unreliable *on this topic*. But that does not mean that they are unreliable speakers in general. In fact, they may well be reliable with respect to all the mundane matters on which they also testify to the children (the locations of specific rooms in the school, chalk in the cupboard, etc.). Does the NESTED SPEAKER case show that topic-specific reliability is necessary for hearer justification? No, because Pauline is unreliable in general, not just in testimony about wild birds, and it could be her general unreliability that (partly) undercuts Fred's justifiedness. Moreover, it is noteworthy that Lackey herself concedes that there are *some* concepts of justifiedness that escape some of the arguments in her paper. In particular, Lackey concedes that the kind of justification she calls "justification grounded entirely in one's subjective perspective" escape these arguments (2006: 182 n. 1). We can take ourselves here to be addressing such a conception of justification.

(10) The uniqueness requirement for a right E-system is admittedly not so attractive. It might be possible to replace this stringent condition with a weaker one.

(11) Perhaps a further condition should be added here—namely, that the agent must be justified in believing that she satisfies the evidential circumstances specified in the norm. Such a condition could easily raise questions about the nature of evidential circumstances: whether to construe them "internalistically" or "externalistically." That is a topic for a different occasion.

(12) At the outset of the chapter, I used the term 'reasonable' interchangeably with 'justified,' as do many epistemologists. With the present principle, however, these terms acquire slightly distinct uses or meanings. When speaking strictly, talk of a doxastic attitude's *justifiedness* should henceforth be qualified by reference to the order of justifiedness in question: first-order justifiedness, second-order justifiedness, etc. An attitude's *reasonability* arises from one or more of its various justificational statuses. It remains to be explored just which justificational statuses are most determinative of reasonableness and under what conditions.

(13) The issues in play here have obvious analogues in moral theory, where the operative terms of appraisal are "right," "obligated," etc., rather than "justified" or "reasonable." Some moral theorists who probe analogous issues in moral theory like to distinguish different senses of "right" or "obligated." For instance, Broad (1985: 128) discusses the question of whether a person is morally obligated to render military service if he is a citizen of a country that is at war, if he is of military age, and if his services are legally demanded of him. Broad assumes, for the sake of argument, that the situation *in fact* makes a moral demand on him. Still, there is a question of whether his obligation is to do what the situation *in fact* morally demands or whether he is obligated to do only what he *recognizes* the situation to demand. Broad says that one can go either way here. One can say that he is obligated by what the situation in fact demands or one can say that he is obligated to do only what he recognizes to be morally demanded of him. (A better analogue of what is under discussion here would be the claim that a person is only obligated by what he is *justified* in believing morality to require, not by what he does believe it to require.) This prompts Broad to speak of obligation or rightness in two different senses: an objective sense and a subjective sense. He writes: "it is futile to pretend that there is just one *right* sense of 'right' and one sense in which we *ought* to use 'ought' " (Broad 1985: 127). Obviously, a similar strategy of distinguishing senses of "justified" or "reasonable" can be adopted in the epistemological case. In effect, this is part of what I am doing. (Thanks to Holly M. Smith for the reference to Broad.)


(14) This evidence could have been acquired in the past, of course. And it may not be readily retrievable, as discussed in the text below. It is also important in this context to note the *path-dependence* properties of evidence acquisition (see, for example, Pettit 2006).

(15) Does the disagreement literature already recognize the variety of evidence I am highlighting? Thomas Kelly (Ch. 6, this volume) speaks of “higher-order evidence,” and sometimes it sounds as if he is concerned with evidence of E-norm correctness. On balance, however, Kelly's discussion of higher-order evidence has a rather different thrust. He argues (this volume) that, if you have higher-order evidence to the effect that you probably made a mistake in responding to your first-order evidence, you should temper that initial confidence in the conclusion. For example, suppose you are a professional mathematician who thinks you have proved a certain theorem, but each of several colleagues claims to find a mistake in the proof. Your colleagues' dissent is higher-order evidence about your original performance in examining the proof. This case illustrates that what Kelly means by “higher-order” evidence is not evidence about the content of any (correct) E-norm but evidence about your competence in arriving at your initial credence—hence evidence about how heavily to weight this initial credence when revising your opinion.

(16) In fact, Feldman (personal communication) reports that this was not his intended meaning. Nonetheless, it is a possible interpretation of the catchy slogan, so it is worth examining closely. That is what I undertake in the remainder of the paragraph.

(17) Another important category of evidence that does not get encoded in memory and therefore is not available for subsequent report is observed evidence whose significance is not *appreciated* at the time of observation. Such evidence is unlikely to be recalled later. Nonetheless, on a plausible approach to evidence, it is part of the total (cumulative) evidence that bears on a belief's justificational status. I assume here a historical approach to justification (cf. the historical reliabilist approach of Goldman 1979). Defenders of other approaches to justification might dissent at this juncture, but the problem of forgotten or neglected evidence is important (cf. Goldman 1999b).

(18) A reader may be surprised to find me giving respectful treatment to a strongly subjective conception of evidence, because this seems at variance with justificational externalism that I have favored in the past. Two clarifications are in order. First, I am not endorsing Huemer's phenomenal conservatism, merely citing it as an extant position. Secondly, in the present chapter's framework, the principal externalist dimension of justification or reasonability would arise from the *ground* or *criterion* of E-system rightness. That is where reliability would enter the picture. An architecture that introduces external factors at this level is entirely compatible with evidential states being highly subjective or “internal.”



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

The Moral Evil Demons

Ralph Wedgwood (Contributor Webpage)

DOI:10.1093/acprof:oso/9780199226078.003.0010

Abstract and Keywords

This chapter discusses moral disagreement and moral relativism, and places the long-standing discussion of these issues in the context of recent discussions of the epistemology of disagreement. It argues that it is rational to have a special sort of 'fundamental trust' in one's own moral intuitions, but it is not even possible to have the same sort of trust in the intuitions of others. As a result, it can be rational for both parties to a peer disagreement to have more confidence in their own views than in the incompatible views of their peers.

Keywords: moral disagreement, moral relativism, intuitions, self-trust, rationality

Moral disagreement* has long been thought to create serious problems for certain views in metaethics. More specifically, moral disagreement has been thought to pose problems for any metaethical view that rejects *relativism*—that is, for any view that implies that, whenever two thinkers disagree about a moral question, at least one of those thinkers' beliefs about the question is not correct. In this chapter I shall outline a solution to one of these problems. As I shall argue, it turns out in the end that this problem is not really a special problem about moral disagreement at all: it is a general problem about disagreement as such. For this reason, in the later sections of this chapter, I shall turn to some general questions in epistemology, about the epistemic significance of disagreement.

1. The Problem of the Moral Evil Demons

There are several different ways in which relativists have argued that moral disagreement poses a problem for their opponents. For example, relativists such as Gilbert Harman (in Harman and Thomson 1995) have argued that relativism gives a *better explanation* of the sort of moral disagreement that exists than any rival view. According to these relativists, the best explanation of this sort of disagreement involves the hypothesis that both sides of the disagreement are in their way correct—whereas no equally **(p.217)** good explanation involves the hypothesis that at least one side of the disagreement holds a belief that is not correct. Whether or not these relativists are right to argue this is a complicated empirical question. As several opponents of relativism, such as Judith Thomson (in Harman and Thomson 1995), have argued, it seems that there are in fact a great many social and psychological mechanisms that could perfectly well explain why we would end up with seriously distorted views about many moral questions.

At all events, it is not this problem that I shall focus on here. I shall suppose that the anti-relativist can give an explanation of the existence of moral disagreement that is at least as good as the explanation that is offered by the relativist. I shall also suppose that the anti-relativist has succeeded in developing a plausible moral epistemology, according to which it is rational for one to form moral beliefs on the basis of one's moral intuitions, at least so long as those moral intuitions have a reasonable degree of coherence with one's overall set of moral beliefs.¹ Even if we grant all these assumptions to the anti-relativist, a further problem seems to arise: once we *learn* about the sort of disagreement that actually exists, why does this information not remove any justification that we might previously have had for the relevant moral beliefs? Why does the information about all the moral disagreement that exists not force us into a thoroughgoing *scepticism* about our moral beliefs?

This problem arises most clearly in cases of moral disagreement between two thinkers who are *equally rational*, and *equally well informed about the non-moral facts*. Obviously, there are many cases where moral disagreement is explained by the fact that one party to the disagreement (or perhaps even both parties to the disagreement) are less well informed about the non-moral facts than they might have been: one party to the disagreement might simply be ignorant of certain non-moral facts; or one party might actually have mistaken or erroneous beliefs about these non-moral facts. Similarly, there are also many cases where moral disagreement is explained by some sort of procedural irrationality on one side or the other. For familiar reasons, bias and self-interest are particularly likely to cause self-deception about moral questions; as a result, people often persist in their **(p.218)** moral beliefs, in spite of the fact that their own stock of beliefs and other mental states would have motivated them to abandon those moral beliefs if they had reflected more rationally about the question.

However, there is no obviously compelling reason why we should deny the possibility of moral disagreements that are not of these kinds.² In the absence of any such obviously compelling reason, we should assume that it is possible for there to be moral disagreements that are not explained either by irrationality or by any lack of non-moral information on either side.

In what follows, I shall suppose that this assumption is correct: it is possible for there to be moral disagreements of this kind—that is, disagreements in which both parties to the disagreement are forming and revising their beliefs in procedurally quite rational ways, and neither side holds its belief because of any error or ignorance about the purely non-moral facts. Indeed, it may even be that some disagreements of this sort are *actual*. For example, some thinkers believe that it is morally wrong for people to eat meat (unless those people have to eat meat in order to stay alive and well), while others believe that eating meat simply for the pleasure of doing so is perfectly permissible. This disagreement may not be due to any irrationality, or to any non-moral error or ignorance, on either side.³ There may be many other such disagreements: for example, there are all the disagreements about sexual morality (for instance, about whether or not there is something morally inferior about homosexuality compared to heterosexuality); there are political disagreements about what forms of liberty or equality are important, and why; there is the disagreement about the moral status of early human fetuses and embryos, and the disagreement about whether the intrinsic value of species diversity and thriving natural ecosystems gives us any moral duties to respect this value; and there are many disagreements about how to balance different moral values or reasons—for instance, how to balance individual rights against collective security, individual autonomy against social order and cohesion, and so on. **(p.219)**

If moral disagreements of this kind are not explained either by procedurally irrational reasoning or by non-moral error or ignorance, what does explain these disagreements? In such moral disagreements, it seems that the two parties hold their incompatible beliefs, not because of procedurally irrational reasoning or non-moral error or ignorance, but simply because the two sides have sufficiently different *pre-theoretical moral intuitions*, which lead them to believe different fundamental moral principles. (For example, vegetarians may have the moral intuition that any creature with the capacity for pain and suffering has the kind of status that makes it impermissible to kill it just for the pleasure of eating it—while carnivores may think that we have much less powerful reasons to refrain from killing non-rational animals than we have to refrain from killing animals, like human beings, who have either the capacity for rational thought or at least the potential to develop this capacity.)

As I noted above, my goal here is to solve a problem that moral disagreement creates for those metaethical views that oppose relativism—that is, for those views that imply that, whenever two thinkers disagree about a moral question, it is impossible for both thinkers to be right. So I shall assume here (at least for the sake of argument) that some metaethical view of this general kind is correct. Given classical logic, it follows that, whenever two thinkers disagree, at least one of the parties has a false or mistaken belief about the question in dispute. So, if the disagreement is due to the two parties' having sufficiently different pre-theoretical intuitions, at least one of the parties must have had *misleading* pre-theoretical intuitions, which have led them to a false and mistaken belief about the question. (I assume here that there need be nothing *irrational* about having such misleading moral intuitions—just as there need be nothing irrational about undergoing a hallucination or optical illusion.)

Now, in some cases of this sort, the pre-theoretical moral intuitions of one of the two parties may contain some sort of incoherence that is not present in the intuitions of the other party. In that case, even though each of the two parties will base its thinking about this issue on its pre-theoretical moral intuitions, it may be fairly easy for the party whose belief is in fact mistaken to discover its mistake by means of this kind of thinking. In some other cases, however, the misleading pre-theoretical moral intuitions of this party to the disagreement may be relatively *systematic*: that is, although these intuitions are in fact misleading, **(p.220)** they also form an overall set that is no less coherent than the intuitions of the other party.

In a case of such systematically misleading pre-theoretical intuitions, even though one of the two parties has an incorrect or mistaken belief, it seems that ordinary moral reasoning will be incapable of leading the believer to discover this mistake. It seems inevitable that any further reflection on the part of this believer will be based on the same systematically misleading pre-theoretical intuitions. Since these intuitions contain no incoherence that would alert the believer to his mistake, it is hard to see how further reflection based on these intuitions could lead the believer to correct his mistake. Something has caused the thinker to have the systematically misleading initial intuitions that he has—his upbringing, or his culture, or his character, or something like that. Whatever it was, I shall call it a “moral evil demon”—something that causes moral error in a way that makes that error undetectable by ordinary means to the one who is deceived.

There is a striking difference between these moral evil demons and their more famous cousins, the Cartesian evil demons, who deceive their victims by giving them systematically misleading sensory experiences. The Cartesian evil demons are creatures of philosophical fantasy; they are not to be found in the actual world. Of course, hallucinations and optical illusions do occur. But when they do occur, there is usually some sort of incoherence in the content of one's experiences so that it is possible to avoid being led into any mistaken beliefs. In real life, sensory hallucinations are never so systematic that they cannot be detected by the ordinary methods of empirical thinking. On the other hand, there is a good chance that the moral evil demons are *actual*. No doubt many moral disagreements are explained by procedurally irrational thinking on one side or the other, or by error or ignorance about relevant non-moral facts. But it seems that there are some disagreements that are more plausibly explained by people's pre-theoretical moral intuitions; and, in some of these cases, it is not clear that the intuitions of either of these disagreeing parties contains any sort of internal incoherence that would lead them to change their view on further reflection. In these cases, then, people's pre-theoretical moral intuitions are leading them astray, in a way that resists correction by ordinary moral thinking: that is, a moral evil demon has been at work. **(p.221)**

This suggests a different argument for scepticism from the argument that is based on the mere *possibility* of an evil demon. It seems that we all have reason to suspect that we live in a world in which moral evil demons are *actually* at work. So what entitles you to any confidence that *your* moral intuitions have not been led astray by such a moral evil demon? What reason do you have to think that you are immune to their malign influence? But then, if you think that there is a significant chance that your own moral beliefs have been distorted by the influence of such a moral evil demon, surely you should entertain some serious sceptical doubts about your moral beliefs?

2. Sidgwick's Principle

So far, this is still a very rough and impressionistic statement of this argument for scepticism about our moral beliefs. We need to lay out this argument in more explicit detail. It might seem that this argument for scepticism about moral belief has the same structure as the following. Suppose that you are in a prison where you have strong reason to suspect that prisoners are *actually* routinely anaesthetized in their sleep and then have their brains removed and placed in vats. Surely this should lead you to entertain very serious doubts about your ordinary perceptual beliefs. It may seem that, in just the same way, once we become aware that we have reasons to suspect that moral evil demons are actually at work, we should entertain serious doubts about our moral beliefs—indeed, perhaps we should even suspend judgment completely about the large parts of our moral thought that seem likely to be subject to disagreements of this sort.

In fact, however, it seems doubtful whether the argument from the probable actual existence of moral evil demons to a sceptical conclusion about our moral beliefs can be quite the same as the seemingly analogous argument in the case of those who are held in a prison where prisoners routinely have their brains placed in vats. In the latter case, you would have a compelling reason to think that there was a nearby possible world in which you used exactly the methods that you actually use to form the very same perceptual beliefs that you actually form (such as “There is a prison guard dressed in blue standing in front of me”), in which those perceptual beliefs are false. That is, in the prison case, you have compelling reason **(p.222)** to regard your perceptual beliefs as unsafe, and as formed by means of a method that is unreliable in the circumstances.⁴

On the other hand, it is not clear that, even if you have compelling reason to think that you live in a world in which moral evil demons are at work, it necessarily follows that you have compelling reason to think that your moral beliefs are unsafe, or formed by means of a method that is unreliable in the circumstances. Even if there are moral evil demons at work in the actual world, it does not follow that there is any nearby possible world in which *your* moral beliefs are false. Suppose that you believe the proposition “It is permissible for human beings to eat humanely killed chickens, purely for pleasure.” If this proposition is true, then it is presumably true at *all* worlds, except perhaps for some worlds that are very remote from the actual world indeed (such as worlds in which chickens are as intelligent as 4-year-old human children, perhaps). So, if it is true, there is no nearby world in which it is false—and a fortiori no nearby world in which it is false and you believe it.

According to a stronger conception of what it is for a belief to be “safe” (or formed through a “reliable” method), for one of your beliefs to be safe it is not enough that there is no nearby world in which you believe that very proposition and it is false; there must also be no nearby world in which you believe any sufficiently similar proposition as a result of a sufficiently similar method, and the proposition believed is false. But it is still not clear that, even if you do have a compelling reason to think that you live in a world in which moral evil demons are actually at work, it follows that any of *your* moral beliefs are unsafe. Perhaps your upbringing, your brain chemistry, and the cultural influences to which you have been subject have all been thoroughly salutary and benign. A world in which you were instead affected by a moral evil demon instead of these benign and salutary influences would be a world in which your whole life was significantly different from how it actually is; and such a world would presumably not be one of the relevant nearby worlds. So, even on this stronger definition of safety, we still do not have an argument for the conclusion that you have a compelling reason to doubt the safety of your moral beliefs. **(p.223)**

However, it still seems plausible that a sceptical argument of some kind could be developed out of the reasons that I have canvassed for thinking that we live in a world in which moral evil demons are actually at work. I propose that the reason why this seems so plausible is as follows. Once one recognizes that moral evil demons may be actually at work, this recognition awakens the suspicion that one's own moral intuitions may have been distorted by such a moral evil demon; and this leads us to think that this suspicion must be dismissed on some basis that is wholly *independent* of the moral intuitions in question. However, the very nature of the moral evil demons ensures that there can be no such fully independent basis for dismissing the suspicion that one's moral intuitions may have been distorted by a moral evil demon: it seems that any argument for the reliability of one's moral intuitions would itself have to depend on one's moral intuitions, and so would fail to count as an independent basis for dismissing this suspicion.

Why should the actual existence of moral evil demons give rise to any such suspicion? Of course, one way in which it might do so is because it makes salient the *possibility* that one is somehow deceived; but this way of arousing doubts gives no special role to the evidence that one has that these moral evil demons are not just possible but *actual*. So I propose that the way in which the actual existence of moral evil demons gives rise to sceptical doubts essentially involves a principle about *actual disagreement*—roughly, the principle that, whenever one believes a proposition *p*, and learns that some other thinker disbelieves *p*, then one should suspend judgment about *p* unless one has some independent grounds for regarding the other thinker as less likely to be right about *p* than one is oneself.

Versions of this principle have been defended by a number of philosophers. One prominent early example is Henry Sidgwick (1907: 342):

if I find any of my judgments, intuitive or inferential, in direct conflict with a judgment of some other mind, there must be error somewhere: and if I have no more reason to suspect error in the other mind than in my own, reflective comparison between the two judgments necessarily reduces me to state of neutrality.

There are also similar claims in other works of Sidgwick (2000: 168):

I suppose that the conflict in most cases [of philosophical controversy] concerns intuitions—what is self-evident to one mind is not so to another. It is obvious that in any such conflict there must be error on one side or the other, or on both. **(p.224)** The natural man will often decide unhesitatingly that the error is on the other side. But it is manifest that a philosophic mind cannot do this, unless it can prove independently that the conflicting intuitor has an inferior faculty of envisaging truth in general or this kind of truth; one who cannot do this must reasonably submit to a loss of confidence in any intuition of his own that thus is found to conflict with another's.

Many questions could be raised about how to interpret these passages. But I shall ignore these questions here. Instead, I shall just assume that the underlying principle behind all these claims is the following:

If you have a belief about a (first-order) question, and then acquire the (higher-order) information that another thinker disagrees with you about that question, you are rationally required to suspend judgment about that (first-order) question, unless you have *independent* grounds for thinking that the other thinker is less reliable about that question than you are yourself.

Just to have a label, I shall refer to this principle as “Sidgwick's principle.”

Many more recent philosophers have made claims that seem very similar to Sidgwick's principle. Thus, Adam Elga (2007) claims that, whenever you learn that another thinker attaches a different credence to a proposition p from the credence that you attach to p , you should adjust your credence to what Elga calls your *prior conditional credence* in p , conditional on the assumption that the other thinker disagreed with you in the way in which he actually does. As Elga explains, by referring to your “prior” conditional credence in this way, he means a conditional credence that is *prior to and independent of* any reasoning that led to your precise view about this particular proposition p . Broadly similar views have also been advocated by Richard Feldman (2006) and David Christensen (2007).

If Sidgwick's principle is correct, then, given that there is disagreement about a large number of moral propositions, you should suspend judgment about those moral propositions unless you have independent grounds for thinking that the dissenting thinkers' beliefs are less likely to be correct than yours. If the other party to the disagreement is less rational than you are, or less well informed about the issue than you are, then perhaps there will be such independent grounds for thinking that he or she is less likely to be correct than you are. However, if a moral evil demon has been at work, then the disagreement is explained simply by the fact that the two parties to the disagreement have different fundamental intuitions. In this **(p.225)** case, there will be no such independent grounds for thinking that the other thinker is less reliable than you are. So, according to Sidgwick's principle, you should suspend judgment about these moral propositions.

In this way, we can give a good interpretation of the problem of the moral evil demons if we view it as resting on something like Sidgwick's principle. For this reason, I shall respond to the problem of the moral evil demons by arguing that Sidgwick's principle is in fact incorrect, and that we need a different model to account for the epistemic significance of moral disagreement.

3. Philosophical Discussions of Disagreement

Sidgwick was particularly interested in moral disagreement. Some of the other philosophers whom I have cited, on the other hand, such as David Christensen and Adam Elga, are interested in disagreement more generally. In fact, the general question about disagreement has recently been discussed by quite a number of epistemologists.⁵

It seems intuitively clear that there is a considerable variety of cases in which one thinker learns that another thinker disagrees with him or her about some question. In some cases, for example, you should regard the other thinker as clearly more expert than you are about the question at issue, and you should unhesitatingly defer to him or her; you may also have many different sorts of reason for regarding the other thinker as more expert than you in this way. In other cases, it is rational for you to regard the other thinker as clearly mistaken; once again, there are many different reasons why you should think this. Then there are also many intermediate cases, where you should give some credence to the other thinker's belief, by weakening your own level of confidence in your own opinion, and shifting your opinion towards theirs, without completely deferring to the other thinker's view.

Many recent discussions of disagreement focus on the special case of disagreements among *epistemic peers*.⁶ There are various ways in which one (**p. 226**) may define the notion of an “epistemic peer.” One simple way would be by simply stipulating that your epistemic peers have *exactly the same evidence* as you have, and are *equally rational* (either in the sense that they are equally rational in the particular process of thinking that led them to their opinion about the particular question that is at issue, or perhaps just in the sense that they are generally speaking no less disposed to rational thinking than you are yourself).

There is a sense in which the problem of moral evil demons that I am focusing on here is similar, since this problem focuses on cases in which there is moral disagreement between thinkers who are *equally rational* (in the strong sense that they are equally rational in the thinking that led them to their dissenting opinion about the question at issue) and *equally well informed about the non-moral facts*. However, it is not obvious that this set of cases is exactly the same as the set of the moral disagreements between epistemic peers, since I have not described the cases that I am focusing on in terms of “evidence.” The reason for this is simple. There are at least two factors that influence what moral beliefs it is rational for one to hold. The first factor consists of one's non-moral beliefs, while the second factor consists of one's moral intuitions. It is not clear whether we should say that it is only the first of these two factors, or both of these two factors, that count as one's “evidence” for one's moral beliefs. Rather than getting into the question of what is the appropriate interpretation of the term “evidence,” I have avoided using this term in formulating the problem of the moral evil demons.

There are other ways of understanding what it means to call someone one of your “epistemic peers” with respect to a given question. For example, we might understand your epistemic peers to include everyone whom it was antecedently rational for you, prior to learning about any disagreement that you might have with them, to regard as equally likely to be correct about the question as you are. Alternatively, we might understand your epistemic peers to be everyone with respect to whom it was antecedently rational for you to regard it as just as likely, if you and they disagree about the question at issue, that they are correct and you are wrong as that you are correct and they are wrong—that is, it is rational for you to have the same conditional probability, given the supposition that you and they (**p.227**) disagree about this question, for the proposition that they are right about this question as for the proposition that you are right.⁷

These two further ways of understanding what it is for someone to be your “epistemic peer” are importantly different from each other, as we shall see later on. In part because there are all these different ways of understanding what it is to be someone's “epistemic peer,” I shall not make much use of this term here. Even without explicitly focusing on disagreements between epistemic peers, my arguments will be relevant to the debates that have explicitly focused on disagreements between epistemic peers. Many of the participants in those debates—including Elga, Christensen, and Feldman—have articulated principles (like Sidgwick's principle) that apply to *all* cases in which one learns that another thinker disagrees with one's belief about a given question, and my arguments will be immediately relevant to the evaluation of those principles.

4. Do the Moral Evil Demons Pose a *Special* Problem about Moral Disagreement?

Is the problem that I have identified a special problem for moral belief? This is a crucial question for our discussion. Several philosophers think that we are unusually resistant to adjusting our moral beliefs in response to learning that other moral thinkers disagree with us; as these philosophers put it, it seems to them that we are much more “intransigent” in the face of such moral disagreement than we are in the face of other kinds of disagreement.⁸ If this is right, then either we must concede that most people are irrationally overconfident in their moral beliefs, or else we must argue that there is something highly special and unusual about the epistemology of moral belief.

It seems doubtful to me whether the case of moral disagreement is a special case in this way. There are many other areas of thought in **(p.228)** which there are disagreements that are just as profoundly entrenched as moral disagreements. For example, there are some *theological* disagreements that do not obviously seem to involve any irrationality on either side (for example, consider a disagreement between an atheist who rejects all arguments for the existence of god, and a deist who accepts a version of the cosmological argument). Similarly, it is far from obvious that all *philosophical* disagreements must involve any irrationality on either side (for example, consider the disagreements between the various rival theories of the semantics of vague expressions in natural language, or the disagreements between different theories of how to understand the possibility that there might have been “additional” objects, which do not exist in the actual world). There also seem to be disagreements about some of the hard questions of history and social theory that do not obviously involve any irrationality on either side, or any error or ignorance about the relevant uncontroversial facts.

In at least some of these cases, the proponents of the various rival views are just as prone to stick to their guns and to refuse to adjust their opinions when they learn that another thinker disagrees with them as they are in cases of moral disagreements. Moreover, the subject matter of each of these disagreements is not in any obvious way a moral question; in many cases, it is not even a normative question of any other kind either. So it does not seem that there is anything unique or special about our tendency to be intransigent in the face of moral disagreement: there are many other non-moral questions on which we refuse to adjust our opinions, even when we learn that other thinkers dissent from our opinion. It seems *prima facie* more plausible that what we have here is a general phenomenon, not something that is peculiar to moral disagreements. For this reason, I shall assume that the solution to the problem of the moral evil demons will depend not on any special feature of the epistemology of moral belief, but on some much more general considerations about the epistemology of disagreement instead.

What makes it possible for different thinkers to reach different conclusions about a question, even if neither of the two thinkers is being in any way irrational, or ignorant, or misinformed about the uncontroversial facts that are relevant to the question? The answer seems to be that it is the same phenomenon that is sometimes identified by means of the slogan **(p.229)** that “theory is underdetermined by the data.”⁹ If we identify the “data” with the uncontroversial facts that are relevant to answering a given question—that is, the facts that both sides of the dispute rationally take for granted—then there are many theoretical questions that are not decided by the data alone. One's beliefs about these theoretical questions will also be influenced by some other aspects of one's overall state of mind—either by one's pre-existing beliefs, or by one's dispositions to have intuitions or impressions about various questions (such as which views are more or less plausible than others), or the like. If it is rationally permissible for these other aspects of one's overall state of mind to influence one's attitude to such theoretical questions, then this may help to explain how two equally rational thinkers may arrive at different views of such questions.

This point brings out a basic feature of the concept of rationality. A sort of relativism is in a way obviously true of the concept of rationality. So long as we reject relativism about truth, then a proposition is either true or not true *simpliciter*, without relativization to anything else. But it is not true in the same way of every proposition *p* that it is either rational to believe *p* or not rational to believe *p simpliciter*. On the contrary, it may be rational for one thinker to believe *p* at one time without its being rational for the thinker to believe *p* at another time, and without its being rational for other thinkers to believe *p* at any time. In this way, the rationality of believing *p* is obviously relative to a thinker and a time.

Still, even if this point explains how it is possible for equally rational thinkers to arrive at different conclusions about such questions, this point does not yet explain how these two thinkers should respond once they *learn* about their disagreement. This is the topic that I shall be focusing on in the remainder of this chapter.

5. A General Epistemological Framework

To make progress with evaluating Sidgwick's principle, we will need to see what the theoretical alternatives to it might be, and what implications these rival principles will have, in the context of our general epistemological **(p.230)** framework. So we will need to make a number of assumptions about the general epistemological framework within which we are working.

One of the main assumptions that I shall make here is that a version of what epistemologists call "internalism" about rationality is true.¹⁰ Roughly, this is the view that rationality supervenes on the relevant thinker's internal mental states. (By speaking of a thinker's "internal" mental states, I mean to exclude the so-called factive mental states, such as *knowing that p*, which by their nature are mental states that can have only a fact or a true proposition as their object.) What it is rational for a given thinker to believe at a given time depends purely on the facts about what internal mental states the thinker has at that time, and what mental processes she is going through at that time. To put it another way, in evaluating a belief as rational or irrational, we are not evaluating the belief on the basis of its relation to the external world; instead, we are evaluating the belief purely on the basis of its relation to the thinker's other internal mental states.

A second assumption that I shall make is that there are two kinds of epistemic rule or principle, which I shall call "special" epistemic principles and "general" epistemic principles respectively. Special principles are principles that specify the way in which it is rational to respond to some quite specific type of mental state. General principles, on the other hand, apply quite generally to all beliefs whatsoever. For example, special principles may include the following: (i) the principle that it is rational to take one's sensory experiences at face value (at least in the absence of any special reasons for doubting that one is perceiving properly in the circumstances); (ii) the principle that it is rational to take one's apparent memories at face value (at least in the absence of any special reasons for doubting that one is remembering properly in the circumstances); and (iii) the principle that it is rational to take one's moral intuitions at face value (in the absence of any special reason for doubting that one's moral intuitions are reliable about the relevant question). General principles might include principles of logical consistency, deductive and inductive coherence, and the like.

One framework that makes this distinction between special and general principles especially clear is a Bayesian framework. Within a Bayesian framework, the *general* epistemic principles are those that require (i) that **(p. 231)** one's degrees of belief should be probabilistically coherent (that is, that it must be possible to represent those degrees of belief by means of a probability function), and (ii) that, when one acquires new evidence, one should update one's degrees of belief by means of Bayesian conditionalization. However, a Bayesian framework will also need to assume that there are some *special* principles as well, in order to explain what it is to acquire "evidence" at all. For example, perhaps one version of this Bayesian framework will suppose that all evidence is acquired directly through sensory observation—but, in that case, it will be committed to the existence of a special principle to the effect that it is rational to treat one's sensory observations as evidence in this way.

For my purposes, however, the classical Bayesian framework is less natural than the less standard variant of the framework in which rational believers update their beliefs by Jeffrey conditionalization (instead of classical Bayesian conditionalization). The classical Bayesian framework requires a definite notion of "evidence," and maintains that the beliefs that it is rational to hold are determined solely by one's prior probabilities and one's "evidence." Moreover, if one is fully rational, then for every proposition p that forms part of one's "evidence," one would have to be *maximally* confident in p —as confident in p as one is in the simplest logical truths. By contrast, the variant of the Bayesian framework that invokes Jeffrey conditionalization allows that certain events that may not themselves count as one's acquiring any "evidence" (such as the event of one's having a sensory experience, or a memory, or a moral intuition) can change the degree to which it is rational to believe a proposition p , without making either p or $\neg p$ as completely certain as a logical truth. Then this approach implies that, to maintain coherence through one's whole system of beliefs, one should revise one's degrees of belief in all the other relevant propositions in accordance with Jeffrey conditionalization.¹¹

In this way, the variant of the Bayesian framework according to which the rational way of updating one's beliefs is by means of Jeffrey conditionalization has no need to appeal to any notion of "evidence" at all. All that it requires is that some event that may not consist in one's acquiring evidence can change the degree to which it is rational for one to believe a given proposition; the knock-on effects of this change for the rest of one's beliefs **(p.232)** are then explained by Jeffrey conditionalization. (According to an *internalist* version of this approach, this event that changes the degree to which it is rational to believe this proposition will always be some *internal* mental event—such as an experience or a memory or an intuition or the like.) As I explained earlier in Section 3, I have avoided speaking of "evidence" here; so it would be more natural for me to opt for the version of the Bayesian approach that involves Jeffrey conditionalization than to go for the classical Bayesian approach.

6. The Epistemic Significance of Information about Others' Beliefs

Within the context of the general framework that I have just outlined, Sidgwick's principle clearly counts as a *special* epistemic principle. In effect, Sidgwick's principle gives a *special* significance to information about others' beliefs. That is, Sidgwick's principle tells one how one should respond to a specific sort of mental state—namely, to the mental state of (rationally) believing that another thinker believes p (where one had previously believed $\neg p$ oneself). It is not a general principle that applies quite generally to all of one's beliefs whatsoever.

Admittedly, Sidgwick's principle is not a special principle about when it is rational to *form* a new belief (like the principle that it is rational to take one's sensory experiences at face value, at least in the absence of special reasons for doubt). Sidgwick's principle is concerned with when we are rationally required to *abandon* a belief, or in other words a principle about when our past beliefs are *defeated*: according to Sidgwick's principle, the information that another person believes p invariably defeats your prior belief in any proposition that is incompatible with p , unless there is *independent* reason for you to believe that the other thinker is less likely to be right about the question than you are yourself. Still, this is a principle that gives a special epistemic significance to information about the beliefs of others.¹² **(p.233)**

What rationale could there be for this special principle? This is a particularly pressing problem for the proponents of Sidgwick's principle, since there seems to be a much simpler way to conceive of the epistemic significance of information about others' beliefs—specifically, one could conceive of such information as simply one more piece of empirical information like any other, the epistemic significance of which is explained purely by *general* epistemic principles instead. On the face of it, there seems to be significantly more to be said in favour of this rival approach than in favour of the approach that is based on Sidgwick's principle.

Consider, for example, how it seems rational to respond to information about the state of measuring instruments, such as pieces of litmus paper. It would surely be misguided to postulate any special epistemic principles that are concerned solely with information about the states of pieces of litmus paper. Instead, we can explain how it is rational for you to respond to this information by appealing purely to general principles instead. If the relevant general principle is a version of conditionalization (such as Jeffrey conditionalization), then the rational way for you to respond to information that the piece of litmus paper has turned red is determined by the *conditional beliefs* that it was antecedently rational for you to have—such as the conditional beliefs that it was antecedently rational for you to have in the various relevant propositions, on the supposition that the litmus paper would turn red, and so on. Given an internalist approach, the conditional beliefs that it is antecedently rational for you to hold themselves reflect all the internal mental factors that have had an influence on what it is rational for you to believe—your sensory experiences, your apparent memories, your background beliefs, and so on. For example, if your mental life has been anything like mine, these internal mental factors have made it rational for you to have a high conditional degree of belief that the liquid into which the litmus paper was inserted is an acid, given the supposition that the litmus paper turned red.

In this way, my past mental life has made it rational for me to have a large stock of conditional beliefs about the world and how it works (including the dispositions of litmus paper). These conditional beliefs presumably include conditional beliefs about other people and how their minds work, and in particular about the circumstances in which people's beliefs are reliable, and in which they are unreliable. For example, the conditional beliefs that it is rational for me to have tell me that people's sensory perceptions **(p.234)** are usually fairly reliable, whereas, unless they have a special expertise, their beliefs about abstruse theoretical matters (such as about the age of the universe, or the correctness of various philosophical interpretations of quantum mechanics) are usually much less reliable. Given that my past mental life has made it rational to have this large stock of prior conditional beliefs, it could be that the way in which it is rational for me to respond to information about the beliefs of others is completely determined by these conditional beliefs, in accordance with some completely general epistemic principles (such as Jeffrey conditionalization).

There are several advantages to this approach, according to which the rational response to information about the beliefs of others is determined purely by general epistemic principles in this way. First, this approach is in an obvious way more economical: if Sidgwick's principle is true, it would surely call for some explanation or rationale; since this approach dispenses with any such special principle, it postulates fewer phenomena that cry out for explanation. Some special principles can be fairly easily explained. In particular, there must be some special principles specifying the epistemic significance of *non-doxastic* states (such as sensory experiences, apparent memories, intuitions, and the like): the epistemic significance of non-doxastic states cannot be captured by general epistemic principles, and can be accounted for only by special principles. But clearly this explanation does not apply in the case of one's beliefs about the beliefs of others. The epistemic significance of one's beliefs about the beliefs of others can easily be captured by general principles; there seems no obvious need for any special principle here.

Secondly, this approach can give an illuminating explanation of the wide variety of ways in which we respond to learning about the beliefs of others. In some cases, it was antecedently rational for you to believe that, whatever belief the other thinker has on the question, it is far more likely that the other thinker is right about the question than that you are right. (Perhaps this is because the question is about the facts of that other thinker's personal life, or about some topic on which that other thinker has world-renowned expertise.) In these cases, you should simply defer immediately to the other thinker's view. In other cases, the situation is reversed: it was antecedently rational for you to believe that, whatever belief the other thinker has about the question, you are much more likely to be right than they are. (Perhaps this is because the disputed question **(p.235)** concerns the grammar of a language of which you are a native speaker, while the other thinker has been studying the language for only a few months.) Then, of course, there is a wide spectrum of intermediate cases, where your prior rational conditional beliefs make it rational for you to respond to the information that the other thinker disagrees with you by *weakening* your degree of belief on the disputed question, but without simply deferring to the other thinker's view. In general, which of these ways of responding to the information that the other thinker disagrees with you is rational in the circumstances is simply determined by the conditional beliefs that it was antecedently rational for you to have; and that, in turn, is determined by the totality of your mental states (including your background beliefs, your experiences, your memories, and so on). There is no limit in principle to the ways in which the totality of your mental states may have determined which conditional beliefs were antecedently rational for you; in particular, both general epistemic principles and special epistemic principles may be involved in explaining how these mental states determined which conditional beliefs it was rational for you to have.

It seems to me that there is precisely the same spectrum of cases when the disagreement concerns a *moral* question as when it concerns any other question. In some cases, your prior rational conditional beliefs will lead you to regard the other thinker's moral sensibility as vicious and corrupt. In other cases, one's prior rational conditional beliefs will lead one to regard the other thinker as more reliable about moral questions of the relevant kind than one is oneself. For example, suppose that you know that you and the other thinker agree about almost all moral questions of this kind, but that, in the few cases in which you have initially disagreed, you have always in the end been persuaded that she was right and you were in fact wrong. Then it will presumably be rational for you to treat her intuitions as more reliable than your own.¹³ Between these two extremes, there are many intermediate cases where rationality will require you to *weaken* your confidence about your moral opinion without requiring that you simply defer to the other thinker.

Thirdly, this approach can also explain why it is rational in some cases to be intransigent in the face of disagreement. This approach does *not* require that your reason for thinking the other thinker to be less reliable than you **(p.236)** are yourself must be *independent* of the reasoning that led you to your view on the disputed question. In some cases, even though you might initially have thought it highly likely that the other thinker would be right about the question at issue (perhaps because, in your experience so far, the other thinker has always seemed impressively intelligent and well informed), the very fact that the other thinker believes *p* may rationally convince you that the other thinker is less reliable than you had previously thought. In these cases, you attach a high *unconditional* probability to the hypothesis that the other thinker will be right about the question—but this is only because you are confident that the other thinker will believe $\neg p$, which you regard as most probably the right answer to the question. You do *not* attach a high *conditional* probability to the proposition that the other thinker is right, on the supposition that the other thinker believes *p* (a belief that you think most probably wrong). Even if initially—before you find out that I disagree with you—you rationally thought that I was just as likely as you to be right, the information that I believe *p* may *by itself* give you sufficient reason to think that I am probably less reliable than you are.

This third sort of case shows how important it is to be clear about the definition of what it is for someone to count as one of your “epistemic peers.” On the one hand, suppose that we say that for you to regard another thinker as your “epistemic peer” (with respect to a given question) is for you to attach an equally high *unconditional* probability to the hypothesis that that thinker will be right about that question as to the hypothesis that you will be right about that question. Then, even if you start out by rationally regarding me as your epistemic peer with respect to a given question, it may be quite rational for you to respond to the information that I believe *p* (which you believe clearly false) by ceasing to regard me as your epistemic peer with respect to that question.

On the other hand, suppose that we say that for you to regard me as your “epistemic peer” with respect to a given question is for you to assign exactly the same *conditional* probability, on the supposition that you and I disagree about the question, to the proposition that I am right as to the proposition that you are right. Suppose that we also assume that the rational way to respond to new information is by some form of conditionalization. Then, if you rationally regard me as your epistemic peer in this sense, it could *not* be rational for you to respond to the information that you and I disagree by concluding that you are more likely to be correct than I am. **(p.237)** This point is not in any way a qualification of this way of understanding the epistemic significance of information about other thinkers' beliefs. This interpretation of what it is for you to regard someone as your “epistemic peer” makes it very unlikely that you will regard many people as your epistemic peers. On any less demanding interpretation of what it is for me to be your epistemic peer, it may be quite rational in certain cases for you to downgrade your assessment of my epistemic standing in relation to your own precisely in response to the information that I disagree with you.

Indeed, in some cases, it may be rational for *both* sides in a disagreement to regard it as more likely that the *other* side is mistaken about the disputed question than that they are mistaken themselves. It may be that some of the cases of apparently irresolvable disagreement that we considered earlier, in Section 1, are cases of this kind. It may be part of the explanation of people's intransigence in the face of these disagreements that such intransigence is in fact a rational response for all the parties involved. Since many people do respond to disagreements with this sort of intransigence, without believing that there is anything irrational about their response, the approach that I am describing here seems significantly less revisionary than the rival approach that is based on Sidgwick's principle. The fact that this approach is so much less revisionary than its rival seems to me another attractive feature of this approach.

7. Epistemic Egoism?

It may seem that the proponent of Sidgwick's principle has an easy way of responding to the complaint that there must be some explanation or underlying rationale for this special epistemic principle. Indeed, the explanation has already been sketched, at least in outline, by Allan Gibbard (1990: 176–81). The rough idea behind Gibbard's account is straightforward. First, Gibbard argues that it is impossible to get by as a thinker without having a kind of “self-trust”—that is, trust in one's own intellectual capacities and dispositions. Then it is argued that this sort of “self-trust” will commit one to a general “fundamental trust” of all minds, including other minds as well as one's own. This might seem to support Sidgwick's principle, because, if we are committed to this sort of fundamental trust in all minds, then it seems that one must always attach some credence to the beliefs of other **(p.238)** thinkers, unless one has some independent reason for discounting those beliefs as unreliable.

According to Gibbard, we can distinguish between two kinds of sources of information. On the one hand, there are those sources of information in which it is rational to have a sort of “fundamental trust”; that is, it is rational for us simply to be disposed to believe the pieces of information that derive from those sources, even in the absence of any *independent* reasons for regarding those sources as reliable. On the other hand, there are those sources of information that we “trust” (that is, we are disposed to believe the pieces of information that derive from those sources), but only because we have such independent reasons for regarding these sources as reliable.

It is clear that we have a fundamental need for a sufficient stock of beliefs about the world in the light of which to live and act. But we could never have any beliefs at all unless we placed such “fundamental trust” in at least *some* sources of information. Gibbard (1990: 178–9) argues that we need to place such fundamental trust quite generally in all of one's “judgments”: that is, even in the absence of any independent reason in favour of regarding one's beliefs as reliable, one should continue to rely on one's beliefs as though they really were reliable—at least so long as one does not have any special defeating reasons for thinking that one's beliefs are not reliable in the circumstances.

Having argued for this sort of “self-trust”, Gibbard (1990: 179–81) goes on to argue that such “self-trust” commits one to a general fundamental trust in all minds as such. There are two main links that connect self-trust with a more general trust in minds as such. First, it may seem that fundamentally the only rational ground for trusting one's own beliefs is simply that they are beliefs: there is surely nothing special about the fact that these beliefs are one's own. Secondly, since one has acquired such an enormous number of one's beliefs from what one was told by other people while one was growing up, if one were to give up this fundamental trust in the beliefs of other people, one would be faced with the practically impossible task of reconstructing one's whole belief system without relying on any of the beliefs that one acquired through one's earlier trust in others. So it may seem that we have no real alternative to having a fundamental trust in all beliefs as such, including the beliefs of others.

One point where this argument could be resisted is by pointing out that it is not clear that one really needs to have this sort of fundamental **(p.239)** trust in *all* of one's own beliefs as such. Perhaps one needs to have this sort of fundamental trust only in a certain *subset* of one's beliefs (such as certain pivotal entrenched background beliefs, perhaps), and also in certain *non-doxastic* mental states, such as one's sensory experiences, one's apparent memories, one's moral intuitions, and so on. If one is not committed to any such fundamental trust in all of one's beliefs as such, then there will be no reason to think that one is also committed to any such fundamental trust in all the beliefs of others.

Still, even if Gibbard's argument is mistaken in this way, we could still argue for a variant of Sidgwick's principle. Let us assume that it is true that one is inevitably committed to having this sort of fundamental trust in one's own moral intuitions. Then it may seem that one will also be committed to having a comparable trust in *all* moral intuitions, including other people's moral intuitions.

It is true that I have no reason to think that the mere fact that some moral intuitions are *mine* makes those intuitions any more reliable than anyone else's. But it does not follow that, if it is rational for me to have this sort of fundamental trust in my own moral intuitions, it must also be rational for me to have the same sort of fundamental trust in everyone else's intuitions. If it is rational for me to have this sort of fundamental trust in my own current moral intuitions, there must indeed be some feature of these intuitions that explains why it is rational for me to trust them in this way. But the feature of these intuitions that explains this need not consist simply in their being moral intuitions that are had by someone at some time. Another part of what explains why my current intuitions have the special rational role for me that they have is that it is possible for me to base my current formation of a moral belief *directly* on these intuitions.

It does not seem possible for me currently to form a moral belief *directly* on the basis of *your* moral intuitions. At best, I can only directly base my current formation of a moral belief on my *beliefs* about your moral intuitions. On the other hand, it *is* possible for me currently to form a moral belief directly on the basis of *my own current* moral intuitions. Moreover, it seems that we are disposed to be guided by our moral intuitions towards forming the corresponding moral beliefs: if I currently have a moral intuition, that moral intuition will immediately incline me to accept the corresponding moral belief (unless I have some special reason for doubting that intuition). **(p.240)** On the other hand, there is no such immediate tendency for *your* moral intuitions to incline me to accept the corresponding moral beliefs; even my own *beliefs* about your moral intuitions do not seem immediately to incline me to accept the corresponding moral beliefs.

At least assuming what epistemologists call an “internalist” view of rationality, the facts that make it rational for one to revise one's beliefs in a certain way must be capable of *directly guiding* one towards revising one's beliefs in that way. But, as I have argued, the fact that someone else has a certain mental state cannot directly guide one in one's revisions of one's beliefs. It is only one's own mental states that can do this. So it is simply out of the question that other people's intuitions should play the same role in rationally guiding my reasoning as my own intuitions. At most, it might be that *my beliefs* about other people's intuitions should play the same role in guiding my reasoning as my own intuitions. But *my intuitions* seem to be such different mental states from *my beliefs about other people's intuitions* that it is implausible to claim that they should play exactly the same role in guiding my reasoning.

Indeed, it is striking that, in these respects, it is only my own *current* intuitions that can play this role in guiding my current reasoning. I cannot directly base the formation of a new moral belief on my *past* intuitions (at best I can directly base the formation of a new moral belief on my *memory* of those past intuitions). Similarly, I cannot base my beliefs directly on my *future* intuitions, but only on my *expectation* of future intuitions. Moreover, if I either remember having a past intuition or expect to have a future intuition, but do not have the intuition itself now, then this memory or expectation will not now immediately incline me to accept the corresponding belief.

It might seem strange to claim that it is not rational to have the same sort of “fundamental trust” in our past moral intuitions as in our present intuitions. But the metaphor of “trust” is misleading here. What this claim really amounts to is the claim that, whereas it can be rational to form moral beliefs directly on the basis of one's current intuitions—even without any additional independent reason for regarding those intuition as reliable—it cannot be rational to form a moral belief directly on the basis of one's memory of a past intuition unless one has some further reason for regarding that past intuition as reliable. Moreover, we can certainly admit that it is rational to be guided by one's *entrenched current background beliefs*; and no **(p.241)** doubt some of these entrenched background beliefs will themselves reflect one's past moral intuitions.

With regard to our future intuitions, it does not seem so strange to me that, even if one receives the information that one will have a certain intuition in the future, one's response to this information should be guided by one's rational assessment of whether one's intuitions can be expected to become more or less reliable in future. There is also no need for this assessment to be independent of one's current intuitions. Indeed in some cases, the very information that I will come to have a certain intuition in future gives me reason to think that my moral sensibilities will deteriorate. For example, suppose that I receive the information that, in the future, I will have the intuition that it is an admirable form of tough-mindedness for the police to have a policy of torturing those whom they suspect of serious crimes. This information would seem to me all by itself to make it rational for me to think that my moral intuitions will probably be less reliable in future than they now are.

Gibbard's second argument is closely akin to a well-known argument that is given by C. A. J. Coady (1992) in the course of his argument for a version of the principle of *credulity*—that is, roughly, the principle that it is rational for us to believe everything that we are told unless we have some special positive reason to suspect that our informant is unreliable.¹⁴ Coady points out that we came to master a language only by coming to believe most of what we were told while we were growing up, and that we acquired many of these childhood beliefs in a fundamentally uncritical way. It may seem then that, unless the principle of credulity is true, we will be rationally required to pursue the Cartesian project of reconstructing our whole belief system from the ground up, without any initial reliance on anything that we have learned from others. Since that project does not seem feasible, it seems that we cannot really be rationally required to pursue it, and so—according to this argument—the principle of credulity must in fact be sound.

As the literature on the epistemology of testimony has made clear, there is a straightforward response that is available to those who are sceptical of the principle of credulity. Perhaps it is true that young children have to **(p.242)** pass through a phase of uncritically believing everything that they are told if they are to learn a language and acquire a sufficiently rich understanding of the world. But it does not follow that, once these children have acquired the status of being fully rational agents, they should continue to be so uncritical. Perhaps we cannot help but continue to rely on many of the beliefs that we acquired uncritically while we were children. But it does not follow that we should continue to acquire beliefs in such an uncritical way.

What makes it possible for this response to Gibbard's second argument to be coherent is that there is at least some plausibility in the suggestion that epistemology should treat our *continuing to rely on our old background beliefs* differently from our *currently forming new beliefs*. Since we are not engaged in Descartes's project of "pure inquiry," we are not required to reconstruct all our background beliefs from the ground up. We may continue to rely on our entrenched background beliefs, even if we cannot reconstruct them all in this way. But it does not follow that all the ways in which we originally acquired those old background beliefs are ways in which it is rational for us now, as mature adult thinkers, to form new beliefs about the world. The fact that we absorbed so many beliefs from other people while we were children shows at most that it must be rational for us to continue relying on the specific beliefs that we absorbed from others, not that we are rationally required to have a fundamental trust in all minds as such. One may have had to pass through a period of childish credulity while growing up, but as an adult thinker one should be able to put this sort of credulity aside.

In earlier work (Wedgwood 2007: ch. 11), I described my position on the epistemic significance of disagreement as a view that it is rational to have a sort of *egocentric bias* in forming beliefs. In a sense, this is an accurate description of my position. It is rational for my background beliefs, experiences, memories, and intuitions to guide me *directly* in a way in which it is simply not possible for your beliefs, experiences, or intuitions to guide me directly. That is, my thinking can be guided by my mental states, even if I am not thinking *about* those mental states at all. Indeed, it is *inevitable* that, if I form any beliefs at all, my thinking will be at least sometimes directly guided by my own current beliefs, experiences, memories, or intuitions in this way.

Moreover, in my view, it is rational for me, in being directly guided by my own current beliefs, experiences, memories, and intuitions, to have a tendency to form new beliefs, endorsing the content of those states, **(p.243)** directly on the basis of those states, even in the absence of any independent reason for regarding those states as reliable guides to the truth. We could metaphorically describe this tendency as a kind of “fundamental trust” in those states (although it should be emphasized that this “trust” in one's mental states need not involve *thinking about* one's mental states at all—one's attention may be fixed on the world, not on one's own mind). On the other hand, in rejecting the idea of any “special” epistemic principles defining the epistemic significance of information about the beliefs of others, I deny that it is rational to have the same sort of “fundamental trust” in the mental states of other people. So my view does give a profoundly different significance to one's own current mental states compared to the mental states of other people.

Unfortunately, my description of my view may have encouraged some misunderstandings. It is *not* my view that your beliefs and my beliefs both function for you as reasons or evidence of fundamentally the same kind, but it is rational for you to give greater weight to your beliefs than to mine. On the contrary, my view is that the role of your current beliefs and intuitions in guiding your thinking is profoundly different from the role of my beliefs or intuitions (and indeed the role of your own past or future beliefs or intuitions as well). The image of different “weights” for two bodies of evidence of essentially the same kind is in fact a travesty of my view.

Thus, I am also not saying, absurdly, that, when you learn that you and I disagree about something, it is rational for you to think to yourself: “I'm me, and he's not; so I'm probably right, and he's probably wrong.” On my view, when you consciously entertain a belief or an intuition with the content p , your attention is not on the fact that you are entertaining this belief or intuition, but rather on p itself. You might express this intuition or belief by saying something like “At least probably, p .” So, in entertaining this intuition or belief, your attention is on the facts, as they appear to you probably to be. According to my picture, when you learn that you and I disagree about p , it may in some cases (although certainly not in all cases) be rational for to continue relying on your original belief. So, in these cases, it may be rational for you to think to yourself: “At least probably, p ; but he believes that it is *not* the case that p : so he's probably wrong.” That does not seem to me to be in any way an absurd or irrational response to this information. **(p.244)**

The conclusion of this chapter, then, is this. It seems plausible that there is a sort of rational asymmetry between one's own moral intuitions and the intuitions of other people: it is rational to have a special sort of “fundamental trust” in one's own intuitions, but it is not even possible to have the same sort of “trust” in the intuitions of others. In consequence, even though we know full well that there is widespread disagreement about fundamental moral issues, and that this shows that in all likelihood “moral evil demons” have been at work, this knowledge may not always require us to suspend judgment about these moral issues completely. It may indeed often require us to *weaken* our degree of confidence in our beliefs about those moral issues. But, at least sometimes, it may be rational for each of us to continue having more confidence in the propositions that we believe than in the incompatible propositions that are believed by those who disagree with us.

Of course, if relativism is false, then, whenever two thinkers disagree, at least one of them has, as a matter of fact, got things wrong, and believes something false. But it is perfectly possible for false beliefs to be rational. Indeed, when you and I disagree, we could even recognize that each of us is rational—while I think that my belief is rational and true and your belief is rational but false, and you think that your belief is rational and true and my belief is rational but false. Even if we do not respond to learning about disagreement by abandoning our original belief, we do not have to be so dogmatic as to conclude that the other thinker is irrational—although we are committed to thinking that the other thinker is mistaken. Since this account of the epistemic significance of moral disagreement seems to be reasonably coherent and plausible, we may conclude that it can solve the problem with which we began. Even given the deep and irresolvable nature of some moral disagreements, the metaethical views that reject relativism can avoid being committed to any sweeping form of moral scepticism.

References

Bibliography references:

Boyd, Richard (1988). “How to Be a Moral Realist,” in Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism*. Ithaca, NY: Cornell University Press, 181–228.

Christensen, David (2007). “Epistemology of Disagreement: the Good News,” *Philosophical Review*, 116: 187–217. **(p.245)**

Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford: Clarendon Press.

Elga, Adam (2007). “Reflection and Disagreement,” *Noûs*, 41/3: 478–502.

- Feldman, Richard (2006). "Epistemological Puzzles about Disagreement," in Stephen Hetherington (ed.), *Epistemology Futures*. Oxford: Clarendon Press, 216–36.
- Frances, Bryan (2008). "Live Skeptical Hypotheses," in John Greco (ed.), *The Oxford Handbook of Skepticism*. Oxford: Oxford University Press, 225–45.
- Fricker, Elizabeth (1995). "Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony," *Mind*, 104: 393–411.
- Gibbard, Allan (1990). *Wise Choices, Apt Feelings*. Oxford: Clarendon Press.
- Harman, Gilbert, and Thomson, Judith Jarvis (1995). *Moral Relativism and Moral Objectivity*. Oxford: Blackwell.
- Jeffrey, Richard (1983). *The Logic of Decision.*, 2nd edn. Chicago: University of Chicago Press.
- Jones, Karen (1999). "Second-Hand Moral Knowledge," *Journal of Philosophy*, 96/2: 55–78.
- Kalderon, Mark (2005). *Moral Fictionalism*. Oxford: Clarendon Press.
- Kelly, Thomas (2005). "The Epistemic Significance of Disagreement," in John Hawthorne and Tamar Gendler (eds.), *Oxford Studies in Epistemology*, i. Oxford: Oxford University Press, 167–96.
- Lehrer, Keith (1976). "When Rational Disagreement is Impossible," *Noûs*, 10: 327–32.
- Miller, Richard W. (1992). *Moral Differences*. Princeton: Princeton University Press.
- Pettit, Philip (2006). "When to Defer to Majority Testimony—and When Not," *Analysis*, 66: 179–87.
- Plantinga, Alvin (2000). "Pluralism: A Defense of Religious Exclusivism," in Philip L. Quinn and Kevin Meeker (eds.), *The Philosophical Challenge of Religious Diversity*. Oxford: Oxford University Press, 172–92.
- Reid, Thomas (1764). *An Inquiry into the Human Mind on the Principles of Common Sense*. Edinburgh: Kincaid and Bell.
- Rosen, Gideon (2001). "Nominalism, Naturalism, Epistemic Relativism," *Philosophical Perspectives*, 15: 69–91.
- Sidgwick, Henry (1907). *The Methods of Ethics*, 7th edn. London: Macmillan.

— (2000). *Essays on Ethics and Method*, ed. Marcus G. Singer. Oxford: Clarendon Press.

Van Inwagen, Peter (1996). "It is Wrong, Everywhere, Always, and for Anyone, to Believe Anything upon Insufficient Evidence," in Jeff Jordan and Daniel Howard-Snyder (eds.), *Faith, Freedom and Rationality: Philosophy of Religion Today*. London: Rowman and Littlefield, 137–53. **(p.246)**

Weatherson, Brian (2007). "Disagreeing about Disagreement." TS, January. <<http://brian.weatherson.org/DaD.pdf>>.

Wedgwood, Ralph (2002a). "The Aim of Belief," *Philosophical Perspectives*, 16: 267–97.

— (2002b). "Internalism Explained," *Philosophy and Phenomenological Research*, 65: 349–69.

— (2006). "The Internal and External Components of Cognition," in Robert Stainton (ed.), *Contemporary Debates in Cognitive Science*. Oxford: Blackwell, 307–25.

— (2007). *The Nature of Normativity*. Oxford: Clarendon Press.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Clarendon Press.

Wright, Crispin (1992). *Truth and Objectivity*. Cambridge, MA: Harvard University Press.

Notes:

* Earlier versions of this chapter were presented to a Metaethics Workshop at the University of St Andrews and to the Fourth Annual Metaethics Workshop at the University of Wisconsin in Madison. I am grateful to the members of those audiences, and also to my colleagues David Charles, Bill Child, Antony Eagle, Lizzie Fricker, John Hawthorne, and Timothy Williamson, for helpful comments.

(1) For an example of a moral epistemology of this kind, see Wedgwood (2007: ch. 10)—although, for the purposes of this discussion, we do not need to presuppose the exact details of that account of what moral intuitions are, or where they come from.

(2) Richard Boyd (1988: 221) advocates a view that he calls the “rational supervenience” of moral beliefs on non-moral beliefs, which denies the possibility of any moral disagreements that do not involve either irrationality or disagreement about the non-moral facts—except in a “few cases” where either bivalence fails (so that in fact neither side in the disagreement is determinately correct or incorrect), or else “explanations in terms of nonculpable inadequacies in methodology or theoretical understanding are readily available” (p. 222). However, it is not clear what reason Boyd has for accepting these claims.

(3) For a longer and more careful argument for this point, see Richard W. Miller (1992: ch. 1).

(4) For this notion of safety, see Williamson (2000: 123–8); for the idea of a method that is unreliable in the circumstances at hand, see Wedgwood (2002a: 276–8).

(5) For some important recent contributions to this debate, see Keith Lehrer (1976), Peter van Inwagen (1996), Alvin Plantinga (2000), Gideon Rosen (2001), Thomas Kelly (2005), Richard Feldman (2006), Philip Pettit (2006), Brian Weatherson (2007), and Bryan Frances (2008).

(6) This term features particularly prominently in the work of Kelly (2005).

(7) Adam Elga (2007) works with this latter understanding of what it is for someone to be your “epistemic peer.”

(8) For this point, see especially Kalderon (2005: 8–36). My colleague Alison Hills is also developing an account of moral epistemology that is designed to explain this allegedly special sort of “intransigence” on the basis of some allegedly special features of moral thought.

(9) Cf. Wright (1992: 157–68).


(10) For arguments in favour of this sort of “internalism,” see Wedgwood (2002b, 2006).

(11) For an account of Jeffrey conditionalization, see Jeffrey (1983: ch. 11).

(12) This issue, about whether there is a “special” principle defining the epistemic significance of information about other people's beliefs, or whether the significance of this information can be completely explained by general epistemic principles, clearly mirrors the debate between “reductionists” and “anti-reductionists” in the epistemology of *testimony*; for this debate, see especially Fricker (1995).

(13) Thus, I need not disagree with the central claims of Karen Jones (1999).

(14) The inspiration for this principle of credulity derives at least in part from Thomas Reid (1764: ch. 6, esp. sect. 24).



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

Disputing about Taste

Andy Egan

DOI:10.1093/acprof:oso/9780199226078.003.0011

Abstract and Keywords

This chapter discusses disagreement about matters of taste. It contends that these disputes are best explained by saying that what is at stake in them is the self-attribution of certain properties and that resolution of the disputes amounts to the parties making the same self-attributions. It argues that this account explains a special sort of defectiveness to which disputes about taste are subject.

Keywords: disagreement, matters of taste, self-attribution, disputes

1. Introduction

“There’s no disputing about taste.” That’s got a nice ring to it, but it’s not quite the ring of truth. While there is definitely something right about the aphorism—there is a reason why it *is*, after all, an aphorism, and why its utterance tends to produce so much nodding of heads and muttering of “just so” and “yes, quite”—it is surprisingly difficult to put one’s finger on just what the truth in the neighborhood is, exactly. One thing that is pretty clear is that what is right about the aphorism, that there’s no disputing about taste, is not *that there’s no disputing about taste*. There’s *heaps* of disputing about taste. People engage in disputes about which movies, music, paintings, literature, meals, furniture, architectural styles, and so on are good, beautiful, tasty, fun, elegant, ugly, disgusting, and so forth all the time. This is obvious to anyone who has watched dueling-movie-critics shows, read theater reviews, or negotiated with a group or partner about which movie or restaurant to go to, or which sofa or painting to put in the living room. It takes great care and good aim to fling a brick without hitting somebody who is engaged in a dispute about taste.

We might suggest instead that what is right about the aphorism is that *there is no sensible, worthwhile disputing about taste*—that disputes about taste are, across the board, defective in some way that makes them nonsensical, irresolvable, bereft of a genuine subject matter, or otherwise second class (and therefore generally not worth pursuing). This can be a tempting thought. As Hume (1757/1965) notes, it is attractive to think that “a **(p.248)** thousand different sentiments, excited by the same object, are all right: Because no sentiment represents what is really in the object . . . To seek in the real beauty, or real deformity, is as fruitless an enquiry, as to pretend to ascertain the real sweet or real bitter.” Something has clearly gone wrong when I say that broccoli tastes better than Brussels sprouts, my friend Mira disagrees, and we launch into a deeply committed dispute aimed at getting to the bottom of this question once and for all, and uncovering the real facts of the matter. This sort of disagreement is defective in a way that makes it a mistake to invest it with any great significance, pursue it too deeply, think that anyone who takes a view contrary to mine must thereby be getting something wrong, or devote a lot of resources to resolving our dispute and arriving at a collective view on the question. The same seems to go for disputes about, for example, whether *Annie Hall* is a more entertaining movie than *Indiana Jones and the Raiders of the Lost Ark*, or whether (to steal an example from Peter Lasersohn (2005)) a particular roller coaster is or is not fun.

We might think, based on looking at these sorts of cases, that there is something defective about *all* disputes about taste. There is a project of engaging in a particular sort of argument and discussion, aimed at arriving at a common view about the matter in question, by both parties becoming convinced that one of the two views is the uniquely correct view to hold, that we engage in about ordinary matters of fact. That project is (we might think), across the board, a bad project to go in for in this domain, because the relevant facts in this neighborhood are either absent, observer-relative, or otherwise incapable of sustaining this sort of collective investigation and debate.

We might think this, but, as Hume points out not long after the passage quoted above, we would be wrong: “whoever would assert an equality of genius and elegance between OGILBY and MILTON, or BUNYAN and ADDISON, would be thought to defend no less an extravagance, than if he had maintained a mole-hill to be as high as TENERIFFE, or a pond as extensive as the ocean.” There are some disputes about taste that clearly *are* in order, and are not defective in the way that Mira's and my dispute about broccoli and Brussels sprouts was defective. If I maintain that some tune I have just idly pinged out on a piano is a more beautiful piece of music than Mozart's *Requiem*, or that the limerick about the guy from Nantucket is better poetry than Shakespeare's sonnet about love being not love which **(p.249)** alters when it alteration finds,¹ I have just got it wrong. There are a lot of cases in which the parties to the dispute should arrive at a common view—in which one of the candidate positions is clearly the one that both parties to the dispute ought to endorse. And disputes about taste often *do* wind up producing agreement. People do frequently manage to convince each other to revise their views about matters of aesthetics or personal taste, and in such cases the parties to the dispute wind up, at the end, with a common view about the matter, and in agreement about how and whether to apply the contested vocabulary to the items in question.

Think, for example, about the case in which Smith judges that Zingerman's pastrami is not tasty because she ate it with too much mustard, or when she was already over-full. Or the case where Jones judges that the symphony is not beautiful because the performance he attended was badly executed, he was angry at the conductor, he had an ear infection, his tinnitus was acting up, he was not paying attention, and so on. In these sorts of disputes, Smith's and Jones's interlocutors are likely to be able to convince them to revise their views by drawing their attention either to features of their situation that were liable to generate interference, or to features of the pastrami or the symphony that they had overlooked when forming their initial judgments.

Other cases are, of course, different. If it emerges that Smith's judgment was not the result of some sort of outside interference—if, for example, trying the pastrami with less mustard and on an empty stomach does not change Smith's mind—her interlocutor should probably conclude that they are just differently gustatorily constituted, and give up on the dispute. Continuing the dispute after this emerges would be a mistake. Similarly, if it emerges that Jones's judgment was not the result of interfering circumstances, or failure to attend to all of the relevant features of the symphony, it may well be that his interlocutor should conclude that they, too, are simply differently constituted, and give up on their dispute.

We are not using our time well if we spend a lot of it in deeply committed disputes about whether broccoli tastes better than Brussels sprouts or vice versa. And not just in the way that we are not using our time well if we spend it in deeply committed disputes about whether Caesar had more or less than 1,000 fleas on his person when he crossed the Rubicon. The **(p.250)** second dispute has got all sorts of things wrong with it—for example, which answer is correct is not very interesting and is not very important, and we are in an extremely bad position to figure it out. But there seems to be something more profoundly wrong with the first dispute—it is not just that it would be a mistake to care a lot about the answer, or that we are badly positioned relative to the evidence we would need in order to establish which answer is right. It is that the idea that there even is a unique correct answer to be discovered, however boring, trivial, and epistemically inaccessible, seems suspect.

Here are two phenomena to notice. (1) The status of a dispute as defective or not does not depend just on subject matter. In the pastrami and symphony examples above, we had pairs of disputes that were alike in their subject matter (the tastiness of the pastrami, the beauty of the symphony), of which one was defective and the other sensible, depending on the circumstances of dispute. (2) Disputes can start off sensible and *become* defective. Even the defective disputes above were in order (or at least, seemed to be in order) until it emerged that the parties to them were differently constituted with respect to their relevant bits of sensory apparatus (or whatever—I do not want, at this point, to commit myself to very much about just what sorts of differences between the disputants undermine the sensibility of these sorts of disputes).

A third phenomenon that is worth attending to is that there are variations in how *robust* different disputes about taste are. Disputes about *gustatory* taste—about what is tasty, for example—are comparatively fragile. It does not take that much to convince us that such a dispute is not worth engaging in. Disputes about the aesthetic value of literary works are much more robust. Others fall somewhere in between. (The robustness of a dispute, it should be noted, seems to vary both with the subject matter and with the circumstances of the dispute. We will—quite reasonably—stick with the project of arguing with one another, attempting to reach agreement, for longer in some cases than in others. I will stick with the project of convincing my friends of the musical virtues of the Ramones for longer than I will stick with the project of convincing my grandparents. And I will stick with the project of convincing my grandparents for longer than I will stick with the project of trying to convince a Martian, or an Earthling who comes from a completely different musical culture.) **(p.251)**

So here are the central phenomena to be explained. First, that there are some disputes about taste that seem to be perfectly in order, perfectly worth pursuing, and perhaps even liable (if all goes well) to wind up producing agreement at the end of the day, and other disputes about taste that seem to be somehow defective, shallow, not worth pursuing, and/or not subject to any sort of resolution. Second, that disputes which start off in the first category can move into the second when the parties to them get a certain kind of evidence—at a first pass, evidence that they are differently constituted with respect to a certain class of capacities or sensory faculties.

What is right about the aphorism that we started with seems to be not that there is *no* disputing about taste, or even that there is no *sensible* disputing about taste, but rather that *some disputes about taste are defective*, and that they are defective in a distinctive sort of way. Aesthetic disputes are in danger of falling into a certain distinctive sort of defectiveness that many other sorts of disputes are not subject to. My goal in what follows will be to offer a theory that explains the difference between the sensible and the defective disputes about taste, in a way that draws the line in what seems like the right place, and gives us some insight into the sort of defectiveness that is at issue. I will do this by giving an account of what is *at stake* in aesthetic disputes—of the upshot of these disputes' being resolved in favor of one of the competing views, and of what the parties to an aesthetic dispute are (typically and centrally) aiming to achieve by engaging in it.

Before we move on, it is worth noting a distinction between two different sorts of uses of aesthetic vocabulary, and pointing out that the phenomena I have just been drawing attention to arise for only one of them. One sort of use, which does not behave in the ways I have just been describing, is what we might call the “baby-food” use of aesthetic vocabulary: when feeding the baby, I might say, “those puréed green beans sure are tasty,” even though I know that they would taste absolutely revolting to me. Similarly, I might, after noting that Fido rolled enthusiastically in the three-day-old dead fish while ignoring the lilacs, say, “I guess the dead fish smells better than the lilacs.”² Disputes over *these* sorts of uses of the vocabulary of taste (**p.252**) do not seem to be subject to the special sort of defectiveness that will be the topic of this chapter. Call these (to have a more dignified and less misleading name than “baby-food uses” or “dead-fish uses”) *sympathetic* uses of aesthetic vocabulary. A natural thing to say about these cases is that I am sympathetically using the aesthetic vocabulary as if from someone else's perspective.

We can contrast these with *first-personally committed* (henceforth just *committed*) uses of aesthetic vocabulary, on which one is subject to accusations of dishonesty or hypocrisy if one first asserts “Vegemite is tasty,” and then balks (without special pleading) at eating it, or asserts “this opera is captivating,” and then yawns (without special explanation) through the performance.

When we are using aesthetic vocabulary sympathetically, our willingness to assert and assent to assertions of “the dead fish smells better than the lilacs” hinges not on our views about our own reactions or dispositions to react to the smells of dead fish and lilacs, but on our opinions about Fido's reactions or dispositions to react to the smells of dead fish and lilacs. I am not subject to any charges of insincerity if I assert, sympathetically, that “the dead fish smells better than the lilacs,” and then steer clear of the dead-fish-smelling areas and seek out the lilac-scented ones. Nor will you be subject to charges of insincerity if you accept my assertion, and then join me in my fish-avoiding, lilac-seeking behavior.

When we are using aesthetic vocabulary *committedly*, our willingness to assert, and to assent to assertions of, “the dead fish smells better than the lilacs” *does* hinge on (our views about) our own reactions, or dispositions to react to, the objects in question. I *will* typically be subject to charges of insincerity if I follow up a *committed* assertion of “the dead fish smells better than the lilacs” with behavior that reveals a clear, stable preference for the smell of lilacs over that of dead fish. And you will typically be subject to charges of insincerity if you *accept* my committed assertion when you in fact have a robust preference for lilacs over dead fish. The point about acceptance is important. It is not the case that we should accept committed assertions of aesthetic claims whenever we think *the speaker* has got the relevant sorts of preferences, dispositions, or whatever. The person who accepts an aesthetic assertion, no less than the one who makes it, is liable to charges of insincerity if they later betray that they lack the relevant **(p. 253)** preferences, attitudes, and so on that go along with that sort of aesthetic evaluation.³

At this stage, I want to avoid signing up for any specific proposals about the semantics of either committed or sympathetic uses of aesthetic vocabulary. What I want to do is just note the contrast, and set the sympathetic uses aside. It is only disputes involving *committed* uses of aesthetic vocabulary that are subject to the distinctive sort of defectiveness that we are investigating here, and that make it attractive to say things like “there is no disputing about taste.” Disputes about sympathetic uses are in general pretty straightforward—once we have figured out how Junior responds to the beans, or how Fido responds to the fish, that is the end of it. Any defectiveness in *these* sorts of disputes is just of the same kind as we find in the Caesar's fleas-type cases.

2. Ground Clearing: Contrasts and Desiderata

So far, I have been a bit sloppy about setting up the problem. Let us now be more careful. First, let us get clearer about what a “dispute about taste” is, exactly. In order not to close off any options prematurely, I am going to characterize disputes about taste as a certain sort of conversational exchange involving a certain sort of vocabulary.

I will call the sort of vocabulary in question *aesthetic vocabulary*, or *vocabulary of taste*, by which I mean to include predicates such as “fun,” “tasty,” “disgusting,” “beautiful,” “elegant,” and the like. I will not say much more about what it takes to be part of “and the like”—I intend the account that follows to apply broadly to aesthetic predicates and predicates of personal taste, but I want to proceed by identifying some core cases, offering an account that seems to work for them, and then letting the chips fall where they may as far as the exact boundaries of the class of expressions for which such an account is appropriate. (I have in mind something like Sibley's notion (1959, 1965) of an aesthetic expression.) **(p.254)**

The sort of exchange in question is the sort that happens when the following two things happen:

1. One party to a conversation assertively utters some sentence *S* and the other assertively utters a sentence (call it $\neg S$) that looks, as far as its surface structure goes, like *S*'s negation; call this a case of *superficial denial*. (I will be concerned, almost exclusively, with pairs of sentences of the form *x is F* and *x is not F*.)
2. The parties take their utterances to be in conflict, and go on to engage in a process of argument, negotiation, bullying, and so on with the aim of arriving at a common position, in which both parties to the dispute are prepared to assert, and to accept assertions of, one of the contested sentences, and both are prepared to deny the other. I will call this sort of process, aimed at this sort of outcome, a *dispute*, and, if and when the aim is achieved, I will say that the dispute has been *resolved in favor of* whichever sentence the parties to the conversation wind up collectively adopting as the one to assert and accept, and that the parties are *aligned* with respect to which of the contested sentences they are inclined to assert.⁴

Many cases of superficial denial are *merely* superficial. For example, crucial expressions are sometimes indexical or otherwise context dependent. (Philosopher Phil asserts “philosophy is my field,” and Linguist Larry responds with “philosophy is not my field.”) In these cases, there is typically no occasion for dispute, since the pairs of assertions are not in any tension with each other—one can perfectly well accept both, and each party to the conversation can perfectly well continue to stand by his own assertion while accepting the other's. There is no need to arrive at a common position about which of the sentences is *the* one for them, collectively, to assert, and indeed the project of getting **(p. 255)** aligned on which sentence to assert would often be deeply misguided and counterproductive. In Phil's and Larry's case, for example, getting both of them into the sort of doxastic state in which they could sincerely assert the same sentence would require one of them to believe something false.

It would be natural to contrast merely superficial denial with *semantic denial*, in which the proposition asserted by the second party to the exchange is the negation of the proposition expressed by the first. But, since there are other ways for a pair of assertions to be in conflict than by one expressing P and the other expressing $\neg P$, it will be useful to use a more general, though somewhat harder to pin down, notion of *genuine conflict*. There is a genuine conflict between an assertion of S by A and an assertion of $\neg S$ by B iff neither party can consistently accept the other's assertion without withdrawing, ceasing to stand by, and ceasing to be prepared to repeat their own. That is, A cannot consistently both stand by her original assertion and remain willing to assert S, while simultaneously accepting B's assertion of $\neg S$, and B cannot consistently stand by her original assertion and remain willing to assert $\neg S$, while simultaneously accepting A's assertion of S. When two assertions are in genuine conflict, each party's making, and continuing to endorse, its own assertion commits it to rejecting the other's.

We can think of this, if we like, in terms of the incompatibility of the *conversational demands* imposed by the two assertions: speech acts (including assertions) have, if the parties to the conversation all cooperate, distinctive sorts of effects on the conversational context in which they occur. We can model this by saying that speech acts impose certain sorts of *demands* on the context—they demand that the context be changed or updated in a certain way. Since the state of the context depends (at least to a large extent) on the actions and attitudes of the parties to the conversation, speech acts also impose, in the same sense, demands on the other parties to the conversation: demands that they do what they have to do, and adopt the attitudes that they need to adopt, in order for the context to undergo the relevant changes. There are, that is, certain things that the audience members must do, and/or attitudes that they must adopt, in order to *go along with* a given assertion (or other sort of speech act). If everyone goes along, their going along brings it about that the context has been changed in some distinctive way. We have a genuine conflict **(p.256)** between two assertions when they impose incompatible conversational demands.⁵

Disputes in the absence of genuine conflict are pretty clearly defective, and are bad uses of one's time and energy. If there is no genuine conflict (as in the case of Phil and Larry), the thing to do is typically for each party happily to accept the other's assertion, and move on to other topics. (Of course, one of the parties might have some independent reason to reject the other's assertion—Phil might suspect that Larry is lying about what he does for a living, for example. But the mere fact that Phil has gone on record with his own assertion of “philosophy is my field” does not present any obstacle to his accepting Larry's assertion of “philosophy is not my field.”)

The presence of genuine conflict is not all that is required in order to have a dispute that it makes sense to pursue, though. Return to the dispute about Caesar's fleas. Martha says, "Caesar had more than 1,000 fleas when he crossed the Rubicon," and Elizabeth responds, "Caesar did not have more than 1,000 fleas when he crossed the Rubicon." It would be a mistake for Martha and Elizabeth to spend a lot of time and energy fighting about this. What is wrong with their dispute is not that they have not succeeded in making conflicting assertions. Of course they have. Here, what is wrong with the dispute is that the project of resolving it just is not going to be a good use of their time. It will not be a good use of their time because (a) there is not much hope of success, and (b) success would not bring much of a payoff, anyway. It is going to be impossible (or nearly so) to get the evidence that they would need in order to resolve the dispute in favor of one claim over the other, and the question is (unless their circumstances are extremely odd) pretty uninteresting and unimportant.

Some genuine conflicts are sufficiently unimportant, given the purposes at hand, that the best thing to do is to pass over them in silence, rather than expending the effort that it would take to resolve them. Some would be extremely difficult to resolve, owing to the inaccessibility of the evidence that would need to be deployed in order to convince both parties to get **(p.257)** aligned on which sentences to assert and accept. Some would be extremely difficult to resolve, owing to the stubbornness with which one or the other party is likely to cling to the view with which they began. (These considerations interact with one another in all of the obvious ways. If the question is important enough, it is worth fighting against a lot of intractability, or going to a lot of trouble to obtain the relevant evidence, in order to resolve the dispute. The more intractable the disagreement, the more important it has to be in order for the project of resolving it to be worth engaging in, and so on.)

There are, then, two requirements on a dispute's being a sensible one to engage in: first, there has to be a genuine conflict. And, second, the conflict has to be one that it is worth resolving—the project of getting both parties aligned on which of the competing sentences is the one to assert and accept has to be one that it is worthwhile to engage in. Call these the CONFLICT and the WORTHWHILENESS requirements on sensible disputes.

Whatever has gone wrong in defective disputes about taste, it does not seem to be the same thing that has gone wrong in Phil's and Larry's dispute. That is, it does not seem to be a failure to satisfy the CONFLICT requirement. When I say "broccoli tastes better than Brussels sprouts" and Mira says "broccoli does not taste better than Brussels sprouts," it would (probably) be a mistake for us to invest a lot of time and energy in a dispute. But the reason for this certainly does not seem to be that we have failed to express conflicting views about broccoli and Brussels sprouts, so that we could just as easily simply accept each other's assertions with an interested nod and a polite, "yes, I see." Given what I have asserted, I cannot accept Mira's assertion without withdrawing my own. And, given what Mira has asserted, she cannot accept my assertion without withdrawing hers. (Similarly, a third party could not happily accept both of our assertions—an onlooker could not sincerely accept my assertion of "broccoli tastes better than Brussels sprouts" and then go on sincerely to accept Mira's assertion of "broccoli does not taste better than Brussels sprouts" without changing his views about the vegetables in question.)⁶

So it must be that the defective disputes about taste are defective because they fail to satisfy the WORTHWHILENESS requirement. But the failure of **(p.258)** WORTHWHILENESS in my dispute with Mira seems to be something deeper than the failure of WORTHWHILENESS in Martha's and Elizabeth's dispute about Caesar's degree of flea-riddenness. It is not just that the evidence that would point us toward the real, objective facts about whether broccoli tastes better than Brussels sprouts is difficult to come by, or that the facts are not very interesting or important. The problem is that the idea that there *is* any crucial evidence to be found, that there *are* any objective facts in this domain to be discovered, seems deeply suspect. It is not that it would be hard to figure out which of us is making the error, or that whichever of us has made the error is unlikely to admit it—it is that there is something suspicious about the idea that there ever needs to have been an error here at all. There is a pretty strong inclination to say that the project of resolving this dispute one way or the other, and establishing a common view about which of our assertions is the one to accept, is a bad one to go in for, not just because the facts about who has *really* got the right view about broccoli and Brussels sprouts are hard to figure out or are not very important, but because both parties have already *got* the right view about the relative tastiness of those vegetables for them to have.⁷

So here are some desiderata for our account of the special sort of defectiveness to which disputes about taste are subject:

The defect should be a failure of WORTHWHILENESS, not of CONFLICT.

The defective disputes about taste are ones in which there is a conflict, but the project of resolving it is a bad one to go in for, rather than ones in which there is no conflict to be resolved.

The failure of worthwhileness should be something deeper than what is happening in the dispute about Caesar's fleas, and should be something that explains why we are drawn to the talk of faultless disagreement and mutual correctness.

3. A Solution

Let us consider a sensible dispute about taste for a moment. Though Brett has never tried peanut butter and chocolate together, he has great (though **(p.259)** misplaced) confidence in his ability to imagine how it would taste, and imagines that it would taste terrible. He sees someone walking by with a peanut butter truffle. Being a philosopher,⁸ he first names the truffle "Alfred," and then asserts "Alfred is not tasty." Yuri, who has extensive experience with peanut butter truffles, and trusts that Alfred is a peanut butter truffle much like any other, responds with "Alfred is tasty."

One of the things that will happen if the dispute is resolved in Yuri's favor (so that, at the end of the day, Brett accepts Yuri's assertion, and Brett and Yuri are both prepared to sincerely assert "Alfred is tasty"), is that both Brett and Yuri will expect that, were they to eat Alfred (in typical circumstances), pleasant gustatory sensations would result. That is, there is a certain property, of being disposed to derive a certain sort of sensation from certain sorts of interactions with Alfred, that each of them will take themselves to have—that each of them will *self-attribute*. Let us call the property *being disposed to enjoy Alfred*.

Of course, something similar is true about the consequences of the dispute's being resolved in Brett's favor—if this is how things pan out, then, at the end of the dispute, both Brett and Yuri will take themselves to have the property, *not being disposed to enjoy Alfred*.

The same sort of thing is a quite common feature of aesthetic disputes, and of aesthetic discourse generally. It is a quite general feature of committed uses of the vocabulary of taste that, when sincerely asserted and sincerely accepted, the end result is a pair (or group) of people, each of whom takes himself to have a disposition to respond to some object or type of object in a particular way (and takes that disposition to be shared by the other parties to the conversation). Quite generally, in any (committed) dispute about taste, there will be some property P (a property of being disposed to have a certain sort of response to certain items) such that, if the dispute is resolved one way, both parties to the dispute will wind up self-attributing P, and, if it is resolved the other way, both will wind up self-attributing the complement of P (henceforth written '¬P').

This is not just an upshot of *disputes*—aesthetic assertions that are simply accepted without argument or discussion have the same sort of effect. One cannot sincerely accept a committed assertion of a sentence of the form *x is tasty* unless one takes oneself to be disposed to have a certain distinctive **(p. 260)** sort of positive response to certain sorts of sensory encounters with *x*. One cannot sincerely accept a committed assertion of *x is fun* unless one takes oneself to be disposed to enjoy certain sorts of interactions with *x*. And so on. Neither will one be inclined sincerely to *make* a committed assertion of the form *x is tasty*, *x is fun*, and so on, unless one takes oneself to have the relevant sorts of dispositions. It seems to be a condition on both the sincere assertion and the sincere acceptance of committed assertions about taste that one self-attribute the relevant disposition to respond. It looks as if, for each simple taste sentence—something of the form *x is F* where *F* is a predicate of taste—there is some property *P* (a property of being disposed to have certain sorts of responses to certain sorts of objects) such that it is a condition on one's either committedly asserting, or accepting a committed assertion of, the sentence that one self-attribute *P*.⁹

This fact helps to make sense of some important features of our aesthetic practice. One very major role that aesthetic discourse plays is a sort of connection-building role, in which people discover commonalities in the sorts of things that they enjoy, appreciate, or despise. This can be a substantial part of the process of building and maintaining interpersonal relationships, and in establishing and maintaining ties to communities and groups. Very many groups and subcultures are defined, at least in part, by the common aesthetic sensibilities of their members (and the contrast between their shared aesthetic sensibilities and those of outsiders). Think of, for example, such subcultures as goths, punk rockers, ravers, trekkies, bikers, and so on.

I propose that we should think of this effect of successful aesthetic assertions, and successful resolutions of aesthetic disputes, of inducing mutual self-attribution of certain dispositions to have a particular sort of response to a particular (kind of) object, as the central business of assertions and disputes about taste, and not as a mere side effect.¹⁰ **(p.261)**

Note that I am not yet trying to offer a theory about *why* mutual self-attribution of certain properties is at stake in disputes about taste—I am not yet offering an account of the semantics of aesthetic vocabulary that explains why that vocabulary is well suited for performing this purpose. I will take that task up later, but for now all I want to do is note that this *does* seem to be what is at stake (or among the things at stake) in these disputes, and use that fact to explain what is defective about the defective disputes about taste. Regardless of what turns out to be the best explanation of *why* successful resolutions of disputes about taste produce alignment in the self-attribution of certain sorts of dispositional properties, the fact remains that the successful resolution of such disputes does *in fact* produce such alignment.

Later in the chapter I will go on to say some things about (a) just what kinds of dispositional properties are at stake, and (b) what sort of semantic theory we ought to provide for predicates of taste, such that we can *predict* that this is what will be at stake. But one need not accept these further claims in order to accept the explanation of the special defectiveness of certain disputes about taste that I will be offering presently.

The project of disputing about taste is a project of bringing it about that we are alike with respect to self-attribution of certain properties. That is only a good project to engage in if we are, in fact, alike with respect to the properties whose self-attribution is at stake in the dispute. If Juan has P and Jim has $\neg P$, it is a bad idea for them to bring it about that they are alike with respect to their self-attribution of P or $\neg P$ —either way they align themselves, one of them is going to have to be getting it wrong and self-attributing a property that he does not have.

Clearly, this project of coordinating our self-attributions of certain sorts of dispositions is an unproductive project for people who are not alike with respect to those dispositions to engage in. It is also a bad project for people who have sufficiently good *evidence* that they are not alike in the relevant respects to engage in. This gives us a way to mark off the distinction between the sensible and the defective disputes about taste: the sensible disputes are the ones where the parties are, and reasonably take themselves to be, alike with respect to the dispositional properties that are at stake in the dispute. The defective disputes are the ones where the parties either are not, or do not, reasonably take themselves to be, alike with respect to the dispositional properties that are at stake. **(p.262)**

Let us distinguish three species of this genus of defectiveness. First, it can happen that the parties to the dispute in which property P is at stake really are different with respect to P, and so bringing it about that either both parties self-attribute P or both parties self-attribute $\neg P$ would require one of the parties to the dispute to self-attribute a property that she lacks. There is clearly something wrong with such a dispute. (Though the parties to the dispute might be blameless in pursuing it, if they do not realize that they are different in the relevant respect.) Call such a dispute *factually defective* (since what is wrong with it is that the parties to the dispute are *in fact* different with respect to the property at stake in the dispute).

Secondly, it can happen that, regardless of whether the parties to the dispute *are* alike with respect to the property P, they are *presupposing* that they are different, and so they are presupposing that the dispute is factually defective. This is also pretty clearly not an occasion in which it is a good idea for them to go in for a project aimed at making them alike in self-attributing P or in self-attributing $\neg P$, and so this sort of dispute is also defective. Only slightly better is the case where the parties to the dispute do not presuppose that they are *different*, but also do not presuppose that they are *alike*. In this case, it is an open possibility, as far as they are collectively concerned, that the dispute is factually defective. This is also not the sort of occasion in which it is a good idea to coordinate on self-attributing either P or $\neg P$. Call a dispute that is defective in either of these ways *presuppositionally defective*, since what is wrong with it is that the required presupposition of similarity is absent.

A final sort of defectiveness, in which we will be particularly interested in what follows, is one in which, whether the parties to the conversation are alike with respect to P or not, and whether they presuppose that they are alike with respect to P or not, it is not *reasonable* for them to presuppose that they are alike with respect to P. This is, again, going to be a sort of case in which it is not a good idea to coordinate on self-attributing P or on self-attributing $\neg P$, since it is an open possibility, as far as what they are *reasonably entitled to presuppose*, that the dispute is factually defective. Call such a dispute *justificationally defective*, since what is wrong with it is that the required presupposition of similarity is not justified.

Consider a dispute in which Smith asserts some aesthetic sentence S, and Jones responds with $\neg S$. There will be some dispositional property P (**p.263**) such that accepting Smith's assertion requires one to self-attribute P and accepting Jones's assertion requires one to self-attribute $\neg P$.

If Smith's and Jones's dispute is defective in any of the ways described above, it will have the following features:

INCOMPATIBILITY: Smith's assertion of S and Jones's assertion of \neg S cannot be simultaneously accepted—they really are in conflict.

POSSIBLE MUTUAL CORRECTNESS: It is either true, or compatible with everything Smith and Jones presuppose, or compatible with everything they are entitled to presuppose, that Smith is correct to self-attribute the property whose self-attribution motivates his assertion (since he really does have P), and Jones is correct to self-attribute the property whose self-attribution motivates her assertion (since she really does have not-P).

NON-SUPERFICIAL DEFECTIVENESS: Smith's and Jones's dispute is defective, and not just because they are talking past each other, or because the facts are hard to figure out, or the question is of no consequence. Their dispute is defective because the project of trying to resolve their dispute, so that Smith and Jones either both come to self-attribute P or both come to self-attribute not-P, is a bad one for them to engage in (since either way of resolving the dispute would in fact either require one of them to self-attribute a property that he lacks, or might, compatibly with everything they presuppose, or compatibly with everything they are entitled to presuppose, require one of them to self-attribute a property that he lacks).

This combination of features makes, I think, for a pretty attractive story about what's defective about the defective disputes about taste. INCOMPATIBILITY lets us say that, even in the defective disputes, we have really got a disagreement, and really got a conflict between the two parties' assertions—the two parties to the dispute really are attempting to press incompatible views on each other. POSSIBLE MUTUAL CORRECTNESS lets us respect the intuition that, in the case of defective disputes, there is something suspicious about the project of trying to figure out who *really* is getting it right, and who has made the mistake. There is a pretty widespread inclination to say that these could be cases of disagreement without error, in which both parties to the dispute are getting it right, and that part of what is wrong with the dispute is that the parties to it are not warranted in making the assumption that one or the other of them must be getting **(p.264)** it wrong, an assumption that they need to make in order to sustain the dispute. This account of what is at stake in these disputes allows us to see what is right about these thoughts. NON-SUPERFICIAL DEFECTIVENESS lets us say that what has gone wrong here is not just a shallow sort of talking-past, or a routine sort of unimportance or epistemic inaccessibility, but something more specific to discourse about taste.

Another nice feature of this account of the distinctive sort of defectiveness that disputes about taste are subject to is that it predicts, as seems correct, that we will have many more sensible aesthetic disputes with our friends and neighbors, and with others who share our perceptual apparatus and cultural background, than we will with those who are very different from us, either in the configuration of their sensory organs or in the sort of culture they come from. We will not have a lot of sensible aesthetic disputes with the Martians, but we will be able to have lots of sensible disputes with our fellow humans who were raised just down the block from us. Sensible disputes will become fewer and farther between, the farther our interlocutors get from our close cultural and biological neighbors, and the closer they get to the Martians.

We can also say something about how a dispute in which there is a danger of these sorts of defectiveness might be expected to proceed. Let us start by looking at an example.

Alan and Clare start off with a default presupposition that they are, in general, alike with respect to their dispositions to enjoy various foods. And so they presuppose that they are alike with respect to their disposition to enjoy Vegemite or not. Alan has good reason, based on his experiences with Vegemite, to believe that *he* is disposed to enjoy Vegemite. Presupposing that he and Clare are alike, and hoping to get her to recognize this similarity, he asserts “Vegemite is tasty.”

In order to accept Alan's assertion, Clare has to self-attribute *being disposed to enjoy Vegemite*. (Call this property V from now on.) Clare has good reason, based on her experiences with Vegemite, to believe that she lacks this property. One of two things could happen at this point:

- (i) She could take Alan's assertion of “Vegemite is tasty” to be a sign that he has never tried Vegemite, or that his experiences with it have always been in unusual circumstances, and he has mistakenly concluded, based on a few fluke pleasant interactions **(p.265)** with Vegemite, that he has got a general disposition to enjoy the stuff. In this case, she is likely to respond with “Vegemite is not tasty,” and attempt to get Alan to accept her assertion and self-attribute *not being disposed to enjoy Vegemite* ($\neg V$).
- (ii) She could take Alan's assertion of “Vegemite is tasty” to be good evidence that they are not, after all, alike with respect to their dispositions to enjoy Vegemite, and stop presupposing that they are. In that case, she will say something like “maybe it's tasty to you—I find it disgusting,” which will, probably, get Alan to accept that they are not alike in the relevant respect, and end the dispute.

Suppose Clare opts for (i). Alan now has a similar pair of options. He can take Clare's resistance to indicate that they are different in the relevant respect, or he can take it to indicate that Clare has had no interactions, or has had flukishly unpleasant interactions, with Vegemite. If he does the second thing, Clare and Alan will be engaged in a dispute.

It is quite likely that they will each, for a while, make efforts to convince the other to self-attribute the relevant property. Clare might encourage Alan to expose himself to Vegemite in the sorts of situations in which his true dispositions are likely to manifest themselves. Alan might attempt to draw Clare's attention to enjoyable-making features of her experiences of Vegemite to which he expects she has not attended.

At some point, one of two things will happen. Perhaps one or the other will be convinced by the other's efforts. Clare, for example, might finally try Vegemite spread to the correct thickness on the right sort of cheese, and become convinced that she really *is* disposed to enjoy Vegemite. She will then withdraw her original assertion and join Alan in asserting "Vegemite is tasty." Alternatively, they might eventually conclude, based on the persistent failure of their efforts to change each other's views, that they really are different with respect to their disposition to enjoy Vegemite. In that case, they will drop the presupposition of similarity.

They are likely, at that point, to stop insisting on their original assertions and retreat to explicitly relativized assertions. Alan is likely to stop asserting "Vegemite is tasty" (which Clare cannot sincerely accept without self-attributing V) and retreat to asserting "Vegemite is tasty to me" (which Clare *can* sincerely accept without self-attributing V—in order to accept **(p.266)** this assertion, all she has to accept is that *Alan* has V). Meanwhile, Clare stops asserting "Vegemite is not tasty" (which Alan cannot sincerely accept without self-attributing $\neg V$) and retreats to asserting "Vegemite is not tasty to me" (which Alan *can* sincerely accept without self-attributing $\neg V$ —in order to accept this assertion, all he has to accept is that *Clare* has $\neg V$). Since these new, explicitly relativized assertions are not in conflict—it is perfectly possible to accept both—this should be the end of the dispute.

In general, the predicted pattern is:

The dispute starts with a presupposition of similarity in place. The parties to the dispute make conflicting assertions. They try for a while to bring each other around. Eventually either one of them succeeds, or they get enough evidence for difference that the presupposition stops being reasonable. If they continue to engage in the dispute at that point, the dispute turns defective. (And, of course, the dispute will be defective right from the start if the presupposition was not reasonable to begin with.)

Some disputes will be more robust than others, because the robustness of the presupposition of similarity will be different in different cases. There will be some disputes about taste such that it is a bad idea to engage in them *at all*, since the presupposition is just not reasonable from the start. There will be others that we ought to give it up at the first sign of conflict. There will be other disputes such that it is just a bad idea to invest a lot of time in them, since the presupposition will very quickly (though not quite immediately) become implausible. Other disputes will, owing to the robustness of the presupposition of similarity, support a great deal of debate despite persistent disagreement.

When the presupposition of similarity *does* become implausible, the thing to do is typically to stop insisting on the original assertion (and thereby stop imposing the demand on one's interlocutor to self-attribute the property) and retreat to talking in explicitly relativized terms. That is, to stop trying to get your interlocutor to accept that *he* has the property in question, and just aim for the weaker goal of getting him to accept that *you* have it. (This is what happens when Alan stops asserting "Vegemite is tasty" and starts asserting "Vegemite is tasty to me.")

So the story is: disputes about taste are aimed at getting all of the parties to the dispute on the same page with respect to the self-attribution of some property—either everybody taking themselves to have it or everybody **(p.267)** taking themselves to lack it.¹¹ This explains why the disputes turn defective when the parties to them get enough evidence that they are different from each other in the right kind of way. It explains why there are some disputes about taste that are quite robust, and others that stop being sensible at more or less the first sign of disagreement. It explains why we get the sense that it is possible for both parties to such a dispute to be right—each really does have the property whose self-attribution they are attempting to push on their interlocutor.

There are still two big questions in need of answers. First, given that self-attribution of some dispositional properties or other is at stake in disputes about taste, what should we say about *which* properties, exactly? Secondly, what do we need to say about the semantics of predicates of taste (and perhaps about the pragmatics of the surrounding discourse) in order to make it *turn out* that disputes about taste aim at getting all of the parties to the conversation to self-attribute those properties? I will address the first of these in the course of answering a potential objection, the full response to which will also require a bit of a revision to the account of what is required for these sorts of disputes to be sensible. The second I will address, rather sketchily, at the end of the chapter.

4. Which Properties?

A natural worry to have about this sort of proposal is that it will predict too much defectiveness and fragility. The defective disputes about taste are the noteworthy exception rather than the rule, and we might be concerned that sensible, robust aesthetic disputes will be too hard to come by on such an account. Here is a reason to be worried. If, when I try to get you to accept my assertion of “broccoli is tasty,” what I am trying to do is get you to self-attribute a disposition to enjoy broccoli, why is it that these disputes do not just fizzle at the first sign of resistance? After all, you are in a better position to know what kinds of responses you get from broccoli than I am—should I not just take your word for it when you, by initially refusing to accept my assertion, signal to me that **(p.268)** you *do not* take yourself to be disposed to enjoy broccoli? Why think that the presupposition of similarity ever survives even the first sign of conflict? This is a potentially very serious worry—there is a danger that the account according to which what is at stake in disputes about taste is the participants' self-attribution of certain properties (having to do with one's dispositions to respond in particular ways to the objects in question) will predict that disputes about taste should be a lot less robust than they actually are.

There are two kinds of response to this sort of concern. First, we can look for properties such that the presupposition that the parties to the dispute are alike with respect to them will be reasonable in a lot of cases, and will remain reasonable even in the face of a fair bit of conflict. Secondly, we can expand the range of disputes that are predicted to be worthwhile to engage in by finding some circumstances in which disputes that aim at the participants' self-attributions of some property P are reasonable to engage in, even in the absence of a reasonable presupposition that the parties to the dispute are alike with respect to P. I will pursue the first strategy in this section, and the second in the next.

In figuring out which properties are at stake in which aesthetic disputes, we are not just out to maximize robustness. What we want is some properties such that the presupposition that we are alike with respect to them has the *right* degree of robustness. The right degree will be different for different disputes—what we want is, for each dispute, to find a property to be at stake in that dispute such that the presupposition that we are alike with respect to that property has the *same* degree of robustness as the dispute. We want the plausibility of the presupposition to run out at the same time as the sensibility of the dispute. Following are some ways to make the disputes increasingly robust, by making the presuppositions of similarity with respect to the properties at stake more robust. Which moves to make (or not) in a given case will depend on how much robustness we want in that case. (And, in fact, robustness will be a multidimensional affair—it matters, not just how *much* evidence it takes to undermine the presupposition of similarity, but also what *kind* of evidence.)

The first move toward greater robustness is one that we have already made: make sure the properties are dispositional. We can be wrong about our dispositions. For example, as mentioned in several examples above, it **(p.269)** could be that our past experiences with the items in question have been in deviant circumstances, which have prevented our actual, robust dispositions from manifesting themselves.

The second way to increase the robustness of a dispute about taste is to make the dispositions that are at stake ones that are liable to be widely shared. Dispositions that are products of widely shared features of our sensory apparatus, for example, would be good candidates. The more local the dispositions are, the more rooted in culture rather than biology, or rooted in individually variable biological features rather than species-wide features, the less robust the presupposition of similarity is likely to be.

This maps pretty nicely onto the sort of “fit with natural capacities” account of taste and aesthetics that we find in, for example, Hume (1757/1965) and Railton (2003, n.d.). Hume finds the foundation for sensible disputes about taste in “a certain conformity or relation between the object and the organs or faculties of the mind.” Railton, following Hume, finds it in the facts about “what matches best and most durably the potentials of our underlying structures” (Railton 2003: 96) or “a particular sort of robust and general match between objects or performances and widespread human sensory capacities and sentiments . . . that permits those objects and events to bring about intrinsically sought, perceptually based experiences in those who become acquainted with them” (Railton 2003: 102). On a Humean account of the sort Railton favors, what it is for something to be *beautiful* is, approximately, for it to be the sort of thing that is robustly disposed to interact with widespread human sensory capacities and sentiments in a way that produces experiences of an intrinsically desirable sort. The natural extension from *beauty* to *tastiness* will say, approximately, that what is *tasty* is what is robustly disposed to interact with widespread human gustatory capacities and sentiments in a way that produces gustatory experiences of an intrinsically desirable sort. This sort of account of aesthetic qualities is, I think, extremely appealing.

On the Railtonian account, then, “Vegemite is tasty” will express something like the proposition *that Vegemite is robustly disposed to interact with widespread human gustatory capacities and sentiments in a way that produces gustatory experiences of an intrinsically desirable sort*.¹² We can capture much of **(p.270)** what is appealing about this sort of proposal by modifying it slightly, so that what is at stake is not (or is not merely) acceptance of possible-worlds propositions about widespread human capacities being thus-and-so, but rather (or also) self-attribution of properties of *having capacities that are thus-and-so*. The property we find in the vicinity of the Railtonian proposal about “Vegemite is tasty,” then, is: *having gustatory capacities and sentiments that are robustly disposed to interact with Vegemite in a way that produces gustatory experiences of an intrinsically desirable sort*.

The general schema:

The property whose self-predication is at stake in a dispute about taste will be of the type: *having F-capacities that are robustly disposed to interact with x in a way that produces G experiences*. How we substitute for F and G will depend on which predicate is being deployed in the dispute, and how we substitute for x will vary depending on what the predicate is being applied to.

Besides inheriting many of the attractive-making features of Railton's Humean account, this will also get us quite a bit of robustness for aesthetic disputes. These sorts of dispositional properties are clearly the sorts of things that we can mistakenly take ourselves to have, or mistakenly take ourselves to lack. And so, it will not always be a good idea to defer to our interlocutors about whether or not they have the dispositions in question.

Another way to increase the plausibility and robustness of the presupposition of similarity is to make the properties in question ones having to do not with our current responses or disposition to respond but with the ones we would have in the ideal. For example, maybe what is at stake in a dispute over ‘x is beautiful’ is something like, *being someone whose suitably idealized self would be disposed to have experiences of an intrinsically desirable sort as a result of looking at/hearing/contemplating/etc. x*.

This is likely to make the properties in question even less luminous than the sorts of dispositional properties we have just been discussing, and allows more room for the parties to the dispute to be mistaken about whether they have the property or not. I could be disposed, *after idealization*, to enjoy Vegemite without being disposed, *as I am now*, to enjoy Vegemite. It also makes it more plausible to think that we are alike, even if we do not think that convergence in the ideal is *guaranteed*. This sort of account also helps us account for the apparent fact that the **(p.271)** proper appreciation of certain sorts of aesthetic qualities requires a sort of training, experience, and so on, and captures an attractive bunch of intuitions about taste—about the knowledgeable and experienced being more reliable trackers of the relevant features of things, and about the possibility of gustatory, culinary, and aesthetic self-improvement (rather than mere *change*).¹³

There is a species of the idealization move available here that will make the presupposition of similarity even more plausible. We might want to say that, in some cases anyway, congruence with one's neighbors is part of what constitutes being ideal in the relevant respect. Then just the fact that we are interacting with one another would give us reason to think that we will converge in the ideal, because part of what goes into determining where our ideals *are* is pressure toward convergence. This would be attractive on the sort of metaethical view advocated by Gilbert Harman (in, e.g., Harman and Thomson 1996), where morality is something like a negotiated system of norms whose purpose is to let us interact smoothly with each other. I suspect that this might have some appeal for at least some aesthetic qualities, too.

Relatedly, we might say that the properties in question are sometimes *group membership* properties—properties such as *being a member of a natural biological kind/cultural group/etc., typical members of which are disposed to have G experiences in response to x*. This would allow us another sort of flexibility—if I take myself to be an atypical member of my group, I can self-attribute the relevant group-membership property, even if I do not take myself to have the relevant dispositions.

There is a lot of room for variation, across different aesthetic predicates, in which sorts of properties one must self-attribute in order sincerely to apply them to something. There are, I think, a lot of extremely interesting questions here about just which aesthetic predicates to associate with which sorts of properties. But now is probably not the time to pursue the detailed questions about particular predicates. What I hope to have shown is just that the sort of account proposed here has a battery of resources at its disposal that seem well suited to capture the interesting differences between different sorts of aesthetic predicates, and to capture the variety of phenomena that we find in a satisfying way. **(p.272)**

5. Sensible Disputes without (or with Dubious) Presuppositions of Similarity

One way to expand the range of disputes about taste that your theory predicts will be sensible—the one we have just been exploring—is to be sure to pick the right properties to be at stake in those disputes. Another way—the one we are about to explore—is to draw attention to phenomena that either (a) lower the bar for how plausible the presupposition of similarity has to be in order for it to be reasonable to make, or (b) make room for sensible disputes about taste even in cases where the presupposition is absent.

One bar-lowering phenomenon that it is worth drawing attention to is the potential pragmatic importance of establishing a common view about some questions of taste and aesthetics. In some cases, a failure to get on the same page with respect to whether we have got the relevant properties is going to make trouble for our capacity to cooperate with one another, and to coordinate our actions. And in some of these cases the sort of coordination of action that failing to coordinate our self-attributions of the relevant properties would make trouble for will be quite important. Sometimes it is not a big deal whether we agree about the attractiveness of a certain sofa. But sometimes we are trying to decorate a house together, and it is important that we find a sofa that interacts nicely with both of our suites of sensory capacities and sentiments.

In these sorts of cases—where it is going to be a bit of a disaster if we turn out to be different in the relevant respect, and cannot coordinate on how we think that the object in question is liable to interact with our (possibly idealized) sensory capacities, and so on—it might be reasonable to cling to the presupposition that we are, after all, alike in the relevant respect, even after a fair bit of evidence has accumulated that we are not. This will be particularly likely in cases where there is not much to do about it if we are not alike except to give up on any kind of coordinated action in this domain—if diverging opinions on this sofa would scuttle the whole project of living in the same house, or diverging opinions on the tastiness of various foods would scuttle the project of cooking together, for example. (And, of course, the effect will be stronger if this is a domain in which it is particularly important to be able to coordinate our actions.) **(p.273)** This is, of course, going to be a degreed phenomenon: the more important establishing a common view is to the prospects of coordinating behavior, and the more important it is to coordinate in the relevant domain, the greater the incentive to stick with the presupposition of similarity in hopes that it will be borne out.

(A related way to lower the bar for how plausible the presupposition of similarity has to be in order for aesthetic assertions to be in order is to note the possibility of using assertions about taste as a sort of low-cost bid to establish similarity—and thereby to find potential partners for coordinated action. One can make the assertion and quickly withdraw it if it meets with resistance, but if somebody jumps up and says “yeah! that is absolutely right!,” you may have found a friend.¹⁴)

There are also two kinds of phenomena that make room for sensible disagreement even in the absence of a pre-existing presupposition of similarity.

The first is the familiar phenomenon of *accommodation*. In the absence of a pre-existing presupposition that we are alike with respect to some property F, I could still make an assertion that is felicitous only in the presence of such a presupposition, in hopes that the audience will, recognizing that I have just made an assertion that is felicitous only in the presence of a presupposition that we are alike with respect to F, accommodate my assertion by bringing such a presupposition into effect. This opens up quite a bit more room for sensible disputing about taste.

The sorts of pragmatic considerations canvassed above, about the potential importance of establishing common ground to coordinated action, also point to another kind of case where we could potentially find sensible disagreement without a presupposition of similarity. In these sorts of cases, we are liable to have reasons to try to *produce* similarity where it was previously absent. In these sorts of cases, we can understand the parties to the disagreement as trying to arrive at a common view not by one party recognizing that she has in fact had the disputed property all along, but by *acquiring* the property in response to pressure from the other disputant. These are cases in which the aim of the dispute is to *change* the other party's (p.274) taste, not to get him to correct his previously mistaken views about what his tastes were like at beginning of the dispute.

Suppose you want someone to self-attribute *being green*. One strategy for how to bring this about involves a mirror. Another involves a bucket of green paint. If your interlocutor is already green, and just does not know it yet, you can get him to self-attribute *being green* by holding up the mirror. If he is not already green, you can still get him to self-attribute *being green* by doing the right sorts of things with the paint. (In fact, there are two different paint-involving strategies. One is to paint him yourself, and ensure that he notices that he has been painted. The other is to convince him to paint himself. In normal circumstances, he will notice that he is doing this.)

There is a conversational precedent for this last, get-them-to-paint-themselves way of getting people to self-attribute properties. Think about the sorts of assertive orders discussed by Anscombe (1957) and later put to use by David Velleman (1989). The doctor asserts, in the presence of the orderly, “the orderly will take the patient to the operating room,” or “the orderly is taking the patient to the operating room now,” in order to *bring it about* that the orderly takes the patient to the operating room, not in order to get her audience (including the orderly) to accept the antecedently well-supported fact that the orderly was about to take the patient to the operating room. (Depending on one's view of future contingents, there may not have been such a fact prior to the doctor's assertion and the orderly's cooperation.)

Other examples of this are not hard to find—it is easy to construct the sorts of contexts in which “you will bring me a coffee,” “Brian will be at the club at midnight,” and “Juan and Carolina are taking down the front door, Sarah and Liz are covering the back,” could be uttered in order to get the audience to bring it about that things are as the assertion represents them to be, not to get the audience to recognize that things already were (going to be) that way.

We can understand these in terms of the usual sort of conversational demands being complied with in an unusual way. The assertion, “you will bring me a coffee,” even when Daniels uses it to give an order, in the first instance applies pressure on O'Leary to come to accept *that O'Leary will bring Daniels a coffee*. O'Leary cannot (in the usual case) accept that unless he takes it to be true, and he is in a position to make it true, or to make it false. So, in order to comply with the pressure to accept that he will bring Daniels **(p.275)** a coffee, he has got to make it true that he brings Daniels a coffee. And so the pressure to *accept* that the content of Daniels's assertion is true is leveraged into pressure to *bring it about* that the content of Daniels's assertion is true.

The same sort of thing could happen on the present account of disputes about taste, so long as the properties in dispute are ones such that we, or our interlocutors, have some control over whether we have them. It is, of course, an open question which, if any, of the relevant properties we have such control over. (It is easier to be sympathetic to this sort of picture if one bears in mind that we need not have *immediate, instantaneous* control—it need not be that we can just, as it were, flip some mental switch and acquire the properties in question. It could be that it requires some long-term Pascalian process of exposing ourselves to the right influences and so forth.)

There is also, at least in some cases, the prospect of using an analogue of the paint-them-yourself method. Consider an idealization property, and a situation in which there are multiple equally good ways for me to idealize from my present position—my present condition requires that I idealize either toward end-state A or toward end-state B, but does not determine which one. Once I have moved far enough toward either endpoint, though, I will be committed—further idealization from any point far enough along toward A can only move toward A, and further idealization from any point far enough along toward B can only move toward B. (Think of ideal endpoints as exerting a gravitational pull, and think about the positions in which one is just in between two equally strong attractors. Alternatively, think of movement toward the ideal as climbing a mountain, and think about the people who are presently living in valleys between two equidistant mountains.) In this sort of case, it could happen that the upshot of your arguing with me, emphasizing certain features of the object(s) we are talking about, getting me to undergo some experiences and do some thought experiments, and so on, is not that I come to realize that I was already going to idealize to be an *x*-enjoyer, but that you push me far enough down the road toward being an *x*-enjoyer that the previously indeterminate idealization-facts become determinate.¹⁵ (Also, on the sort of idealization story where convergence is part of what constitutes the relevant **(p.276)** sort of idealization, maybe you can push my endpoint of idealization around just by interacting with me, or by taking steps to ensure that you are in the group convergence with which matters.¹⁶)

There are also some disputes about taste that are sensible, not because of any prospects of arriving at a resolution, but for some other reason. Many cases of arguments between sports fans are like this—the participants engage in a sort of pretense of attempting to convince each other that their favorite team or player has the most electrifying offense, the most terrifying defense, the filthiest curveball, and so on, but what makes the dispute worth engaging in is not the prospects of a successful resolution. Sometimes the dispute is just enjoyable in itself. Sometimes it is valuable because engaging in the dispute helps one better to appreciate the aesthetic qualities of the items under discussion. Sometimes the value is in the extra appreciation of the merits of one's own view that one acquires in the process of defending it against attack. Sometimes the process of mutual discovery, in which the parties to the dispute come better to understand each other's aesthetic sensibilities, even without coming to *share* them, makes the dispute worthwhile.¹⁷

So while, to a large extent, the status of an aesthetic dispute as sensible or defective will track the presence or absence of a reasonable presupposition that the parties to the dispute are alike with respect to the property whose self-attribution is at stake in the dispute, these two will not walk *perfectly* in step.

6. Semantic Proposals

We have got a story about the upshot of disputes about taste—about what a successful resolution requires—that lets us explain the distinctive sort of defectiveness that they are subject to. What we are still missing is a story about the semantics that predicts that those disputes will have that sort of upshot. I will not attempt a detailed spelling-out of such a semantic theory here—instead, I will give a rough sketch of what seem to me to be the **(p.277)** two most promising proposals, and say a little bit about which I prefer and why.

The first proposal begins by adopting a Lewisian (1979) account of the propositional attitudes, on which the objects of belief, desire, and so on are *properties*, and a Stalnakerian (1978) account of assertion, on which accepting an assertion requires that one accept its content. Then one very straightforward way to explain the fact that, in order to resolve a dispute about S and $\neg S$, both parties need to come to self-attribute either P or $\neg P$, is to say that the content of S is P , the content of $\neg S$ is $\neg P$, and accepting an assertion requires believing (that is, self-attributing) its content.¹⁸ (Officially, we will want to allow for forms of acceptance that are weaker than belief, and so we will want to allow for something like for-purposes-of-this-conversation self-attribution. I will continue to abstract away from this.)

Here is some motivation for adopting the Lewisian view about the objects of propositional attitudes: there is a certain doxastic similarity between all of the well-informed people with burning pants, and a certain conative similarity between all of the kids who want to grow up to be firefighters. One way to capture these similarities is to say that there is some potential object of propositional attitudes that all of the well-informed people with burning pants believe, and some potential object of propositional attitudes that all of the kids who want to grow up to be firefighters desire.¹⁹ We cannot say this if we think that, necessarily, the objects of the propositional attitudes are always possible-worlds propositions. (The only candidate possible-worlds propositions in the neighborhood, when both Jane and Carlos want to be firefighters, seem to be the singular proposition about Jane and the one about Carlos, and the existentially quantified proposition. And it is not desiring any of these that marks the relevant conative similarity between Jane and Carlos. Jane and Carlos could both want to be firefighters without sharing a desire *that Jane be a firefighter*. And everybody, not just the aspiring firefighters, desires *that somebody be a firefighter*.) We can say it, however, if we think that properties are (or can be) the objects of propositional attitudes. What Jane and Carlos have in common is a desire directed toward the property, *being a firefighter*, and what all the **(p.278)** well-informed people with burning pants have in common is that they all self-attribute (that is, believe) the property, *having burning pants*.

Once we have made room for properties as potential objects of *belief*, we have gone a long way toward making room for properties as the objects of *assertion*. And if we think that the upshot of a successful assertion is the addition of the assertion's content to the conversation's presuppositions (to the stock of potential objects of belief which all of the parties to the conversation believe, believe that the others believe, etc.), then we have a straightforward explanation why successful resolution of a dispute over some aesthetic sentence S results in the mutual self-attribution of some property (of one of the sorts discussed above). It's because successful resolution of the dispute requires both parties to accept an assertion of a sentence (either S or its negation), the content of which is a property of the relevant type. Acceptance of an assertion requires believing its content, and to stand in the *belief* relation to a property is to self-attribute it.

Given this Stalnakerian view of the relation between assertion and content, and the view that sentences like "Vegemite is tasty" or "the symphony is beautiful" have self-locating content, we can explain all of the phenomena.

There is genuine conflict between assertions of Alan's assertion of "Vegemite is tasty" and Clare's assertion of "Vegemite is not tasty", because accepting Alan's assertion would require us to self-attribute (something like) *being disposed to enjoy Vegemite*, and accepting Clare's would require us to self-attribute *not being disposed to enjoy Vegemite*. No one can simultaneously accept both assertions, and so they are in conflict—they impose incompatible conversational demands on the parties to the conversation.

It is absolutely crucial to making this sort of story work that we take the relation between content and assertion to be the one described above, according to which the essential effect of an assertion with content P is that cooperative and credulous audience members *come to accept P*. (Which means, in the case of assertions whose content is some property P, that cooperative and credulous audience members come to self-attribute—that is, take *themselves* to have—P.) We do not get any sort of conflict if our view of the relation between assertion and content is that, in the case of self-locating assertions whose content is some property P, cooperative and credulous audience members come to accept *that the speaker has P*. **(p.279)**

In the case in which Alan and Clare are just differently gustatorily constituted, their dispute is defective because it is a bad idea for them to add either the property *being disposed to enjoy Vegemite* or *not being disposed to enjoy Vegemite* to their conversation's presuppositions. That is, it is a bad idea for them to get aligned on their self-attributions of *being disposed to enjoy Vegemite* and *not being disposed to enjoy Vegemite*. It is a bad idea, because both parties to the dispute are absolutely correct to have the view that they do, and getting aligned on their self-attributions of these properties would require that one or the other self-attribute a property that they in fact lack.

This, incidentally, shows why the very first place in which one might be inclined to look for self-locating content in natural languages—sentences involving first-person indexicals—is not in fact a very good place to look. (At least, it is not a good place to look so long as you are assuming that the role of an assertion of S is to add the proposition or property that S expresses to the conversation's presuppositions. And that is the only view about assertion on which attributing self-locating contents to sentences in context makes predictions that are interestingly different from those on which the sentences just express ordinary possible-worlds propositions about the speaker.)

Suppose that the sentence “I am John Malkovich” expressed the property *being John Malkovich*. Then introductions would be disastrous. The effect of Mr Malkovich's assertion of “I am John Malkovich” (if it were accepted) would be to add *being John Malkovich* to the conversation's presuppositions. Part of what would be involved in this would be all of the parties to the conversation self-attributing *being John Malkovich*. This is not what happens when people introduce themselves. Sometimes people do mistakenly come to self-attribute such properties as *being John Malkovich*, *being Hume*, or *being Napoleon*. They *do not* do this, however, simply by being credulous when Malkovich, Hume, or Napoleon introduces himself. So “I am John Malkovich” does not express the property *being John Malkovich*. Other attributions of self-locating content to sentences involving first-person indexicals are similarly disastrous. So we ought *not* to believe that sentences involving first-person indexicals have self-locating content. We ought instead to believe the usual sort of Kaplanian theory, according to which first-person indexicals are, well, *indexicals*—they refer to different individuals on different occasions of use. **(p.280)**

(How, then, do we learn anything other than the necessary truth *that John Malkovich is John Malkovich* from Mr Malkovich's introduction? By exploiting the sort of pragmatic mechanism Stalnaker (1978) sets out: we, as competent English speakers, know that, depending on who is speaking, “I am John Malkovich” expresses either a necessary truth or a necessary falsehood. If we take the speaker to be sincere and well informed, we trust that he is among the individuals who is in a position to say something true, rather than something false, with an utterance of “I am John Malkovich.” The only person in such a position is Mr Malkovich. And so we come to accept that the person before us who just uttered “I am John Malkovich” is John Malkovich.)

(A few clarifications for those who are concerned about how the semantic details are going to go. (1) This kind of account of aesthetic sentences does not force us to adopt a weirdly bifurcated theory of content that deals sometimes in possible-worlds propositions and sometimes in properties or self-locating propositions. Everything that we can do with a possible-worlds proposition, we can do with a self-locating proposition that does not distinguish between different positions within the same world.²⁰ (2) It also does not force us to say that sentences and predicates have the same kinds of semantic values. Whatever the semantic values of sentences are, the semantic values of predicates are functions from objects to sentence-type semantic values. So, if the semantic values of sentences are properties, the semantic values of predicates are functions from objects to properties. (See Egan 2006a, b.) (3) We also are not forced to add any additional formal apparatus to our formal semantic theories. What this sort of proposal requires is that we say that these sentences have contents that take truth-values relative to possible situations, positions, or predicaments rather than relative to possible *worlds*. We can do this by leaving the formal apparatus of our theory of types just the same (we still say that sentences in context have semantic values of type $\langle s, t \rangle$, for example), while revising our interpretation of what things of type s are like—we interpret them not as *worlds*, but as *positions* or *predicaments*.)

Another option is to adopt a straightforward contextualist account of the semantics of aesthetic vocabulary, and say that the connection between **(p.281)** accepting S ($\neg S$) and self-attributing P ($\neg P$) is pragmatic, rather than semantic.

We can do this by exploiting the possibility of fighting not over the truth or falsity of the propositions that are in fact the contents of our assertions, but over which propositions we are asserting—that is, over the nature of the context that we are in. For example, we can use assertions of, for example, “Joe is rich” and “Joe is not rich,” or “Michael is tall” and “Michael is not tall,” to fight not about truth or falsity of the propositions expressed, but about what the contextually salient standards of wealth or height are or ought to be. If we understand disputes about taste on this model, we can predict that what is at stake in such disputes is the self-attribution of the relevant sorts of properties.²¹

A little more detail about the sort of dispute in question. There are two kinds of disputes we might have about “Joe is rich.” In one kind, we are in agreement about how much Joe has to have in order to be rich, and we are fighting about whether he has that much or not. In another kind, we know how much Joe has, and we are fighting about whether that is enough to be rich. The second kind of dispute is naturally thought of as a fight about what sort of context we are in. (Of course there are also disputes—probably most of the ones we actually engage in—in which both questions are open.)

There are three things to notice. First, that sort of fighting-about-the-context use of “Joe is rich” and “Joe is not rich” makes sense only so long as we are presupposing that “Joe is rich” has the same truth value in each of our mouths. (The easiest way for this to be so is for our utterances to express the same proposition, which will happen only if our uses of “rich” have the same semantic values.) In order to have this sort of context-shifting fight, your utterance of “Joe is rich” has to be in *competition* with my utterance of “Joe is not rich.” If we thought that our two contexts might be relevantly different, such that we were (or might be) just talking past each other, we could just accept each other's assertions without revising our views about what *our own* contexts were like at all.

Secondly, notice that what is at stake in these sorts of context-shifting disputes just is the self-attribution of a certain sort of property. For any **(p.282)** context-sensitive sentence, there is a property that one must have in order to be in a position to use it to utter a truth, and that one must take oneself to have in order sincerely to assert it. In general, when S is context-sensitive, sincere assertors have to self-attribute *being in a context in which an utterance of S would express a truth*, and, to the extent that they are semantically competent with S, they have to self-attribute whichever properties one must have in order to be in such a context.

For “I am hungry,” the relevant property is *being hungry*. For “France is hexagonal,” the relevant property is *being involved in a conversation whose standards of precision are such that France counts as hexagonal*. It is plausible that we sometimes care about whether the context is a certain way, not for its own sake, but because we care about whether we, and the other parties to the conversation, are in the right condition to *make* the context that way. Whether I am in a context in which “Joe is rich” expresses a truth depends not just on how much wealth Joe has, but also on what sorts of attitudes I, and the other parties to my conversation, have toward wealth. How much do we think that one has to have in order to pursue the sorts of projects and life plans that we take to be important? Which sorts of things are we taking to be necessities, and which luxuries? The answers to these, and surely other, questions about our attitudes will be relevant to determining which property “rich” picks out in a given context. We can apply pressure on each other to change those semantic-value-affecting attitudes by applying pressure on each other to accept that, for example, “Joe is rich” expresses a truth in our context. (Note that this actually divides into two different kinds of fight: one in which we are attempting to change each other's views about the antecedent facts about the context, and one in which we are attempting to *shift* the context in one direction or another.)

In cases where we are presupposing that S has the same truth-value in your mouth as in mine, I will not be able to accept your assertion of S without taking myself to be in a position to truly assert it myself. And so my sincere acceptance of your assertion of S will require that I self-attribute the same property that your sincere assertion of S requires you to self-attribute. So, the upshot of a resolution of a dispute about some context-sensitive sentence will be the mutual self-attribution of whatever property (or properties) one must have in order to be in a position to use it to assert a truth. And, on plausible contextualist accounts of aesthetic vocabulary, these are liable to be precisely the sorts of properties discussed in Section 3. **(p.283)**

Finally, notice that this sort of dispute will be subject to the right sort of defectiveness. If aesthetic vocabulary is context-sensitive, then these sorts of disputes over aesthetic sentences, in which the aim of the disputants is to change the other's view about the context they inhabit (and therefore about which proposition is expressed), rather than to change their view about the truth or falsity of the particular proposition expressed, will be sensible to engage in only so long as the parties to the dispute take themselves to be alike with respect to the property (or properties) that one must have in order to be in a position to assert something true by use of the sentences in question.

So, at least at first glance, this sort of contextualist account looks pretty good. Let me say briefly before closing why I think that the previous, self-locating proposal is preferable.

My main concern is that it makes the wrong predictions about what we should do when the presupposition of similarity fails. On the sort of contextualist proposal we are now considering, once the presupposition of similarity fails, it should be clear that we are in a situation where the parties to the dispute are simply talking past one another, and their assertions are not really in conflict. Alan's assertion of "Vegemite is tasty" means something like *Vegemite is tasty to Alan, Vegemite is tasty to Australians, or Vegemite is tasty to Alan-type subjects*. Clare's assertion of "Vegemite is not tasty" means something like *Vegemite is not tasty to Clare, Vegemite is not tasty to North Americans, or Vegemite is not tasty to Clare-type subjects*. Once it is clear to the parties to the dispute that this is the case, the thing for them to do should be just to accept each other's assertions, and the thing for a third party to do should be just to accept both assertions. (In the same way that, once it becomes clear that we are using "here" to talk about different places, I should happily accept your assertion of "the Rock and Roll hall of fame is here," and you should happily accept my assertion of "the Rock and Roll hall of fame is not here.")

But this seems wrong. The right reaction to the failure of the presupposition of similarity *is not* for each party just happily to accept the other's assertion. The right reaction is to *stop asserting those sentences*. It would be very strange if, once Alan and Clare stopped presupposing that they were gustatorily similar, Alan continued to assert "Vegemite is tasty," and Clare responded by nodding acceptance and saying "that's interesting." It would also continue to be weird and infelicitous for a third party to accept both **(p.284)** Alan's and Clare's assertions. In short: the contextualist theory predicts that, once the presupposition of similarity fails, the dispute should become defective because the CONFLICT condition is no longer satisfied (since there is no longer any difficulty in accepting both parties' assertions). But this seems mistaken. Even after the presupposition has given way, the assertions are still in conflict. It is just that the goals one would be pursuing by continuing the dispute are not good ones to pursue—what gives out is WORTHWHILENESS, not CONFLICT. The self-locating account predicts this; the contextualist account does not.

7. Conclusion

We can explain the puzzling phenomena about disputes about taste with which we began with the chapter by saying that what is at stake in such disputes is the self-attribution of certain properties. The project of seeking to resolve such a dispute is the project of trying to bring it about that the parties to the dispute all self-attribute either the relevant property or its complement. If we say this, and take note of the conditions under which it makes sense to engage in that sort of project, we can explain the special sort of defectiveness to which disputes about taste are subject, and we can make the right kinds of predictions about the sorts of circumstances under which that defectiveness will manifest itself.

That is what I take to be the main point of this chapter. I have also argued more tentatively for some proposals about just which kinds of properties might be at issue, and about what sort of semantic theory we ought to endorse such that we can predict that the self-attribution of those sorts of properties is what is at stake in disputes about taste. I think that these are likely to be right, or at least on the right track, but they are detachable from the central point.

References

Bibliography references:

Anscombe, G. E. M. (1957). *Intention*. Cambridge, MA: Harvard University Press.

Barker, Chris (2002). "The Dynamics of Vagueness," *Linguistics and Philosophy*, 25: 1-36.

DeRose, Keith (2004). "Single Scoreboard Semantics," *Philosophical Studies*, 119/1-2: 1-21. **(p.285)**

Egan, Andy (2006a). "Secondary Qualities and Self-Location," *Philosophy and Phenomenological Research*, 72/1: 97–119.

— (2006b). "Appearance Properties?," *Noûs*, 40/3: 495–521.

— (2007). "Epistemic Modals, Relativism, and Assertion," *Philosophical Studies*, 133/1: 1–22.

Harman, Gilbert, and Thomson, Judith Jarvis (1996). *Moral Relativism and Moral Objectivity*. Cambridge, MA: Blackwell.

Hume, David (1757/1965). "Of the Standard of Taste," in *Of the Standard of Taste and Other Essays*, ed. John W. Lenz. Indianapolis, Bobbs-Merrill.

Lasersohn, Peter (2005). "Context Dependence, Disagreement, and Predicates of Personal Taste," *Linguistics and Philosophy*, 28/6: 643–86.

Lewis, David (1979). "Attitudes *De Dicto* and *De Se*," *Philosophical Review*, 88: 513–43;

repr. with postscripts in David Lewis, *Philosophical Papers*, i. Oxford: Oxford University Press, 1983.

MacFarlane, John (2007). "Relativism and Disagreement," *Philosophical Studies*, 132: 17–31.

Mill, John Stuart (1859/1978). *On Liberty*. Indianapolis: Hackett Publishing Co.

Nolan, Daniel (2006). "Selfless Desires," *Philosophy and Phenomenological Research*, 73/3: 665–79.

Perry, John (1979). "The Problem of the Essential Indexical." *Noûs* 13: 3–21.

Railton, Peter (2003). "Aesthetic Value, Moral Value, and the Ambitions of Naturalism," in J. Levinson (ed.), *Aesthetics and Ethics: Essays at the Intersection*. Cambridge: Cambridge University Press, 1998; repr. in Peter Railton, *Facts, Values and Norms*. Cambridge: Cambridge University Press, 85–130. Page references are to this reprinting.

— (n.d.). "Objectivity, Realism and Aesthetic Value." Manuscript.

Richard, Mark (2004). "Contextualism and Relativism," *Philosophical Studies*, 119: 215–42.

Sibley, Frank (1959). "Aesthetic Concepts," *Philosophical Review*, 68/4: 421–50.

— (1965). "Aesthetic and Non-Aesthetic," *Philosophical Review*, 74/2: 135–59.

Stalnaker, Robert (1978). "Assertion," *Syntax and Semantics*, 9: 315–32; repr. in Stalnaker, *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press, 1999, 78–95.

Stephenson, Tamina (2007a). "A Parallel Account of Epistemic Modals and Predicates of Personal Taste," in E. Puig-Waldmüller (ed.), *Proceedings of Sinn und Bedeutung*, 11: 583–97.

— (2007b). "Judge Dependence, Epistemic Modals, and Predicates of Personal Taste," *Linguistics and Philosophy*, 30: 487–525. **(p.286)**

Velleman, J. David (1989). "Epistemic Freedom," *Pacific Philosophical Quarterly*, 70: 73–97;

repr. in J. David Velleman, *The Possibility of Practical Reason*. Oxford: Clarendon Press, 2000, 32–55.

Williams, Bernard (1995). "Internal Reasons and the Obscurity of Blame," in Bernard Williams, *Making Sense of Humanity*. Cambridge: Cambridge University Press.

Wright, Crispin (1992). *Truth and Objectivity*. Cambridge, MA: Harvard University Press.

Notes:

(1) Sonnett 116.

(2) These are slight modifications of examples from Stephenson (2007b). She attributes them to Kai von Fintel and Danny Fox.

(3) This is one of the reasons why the first contextualist theory of taste that we might try out—the one according to which "tasty" always means *tasty to x*, where *x* is the speaker—does not actually look very attractive.

(4) We will be looking at only a subset of all of the very many kinds of disputes there are. There are all kinds of ways to have disputes—and perfectly sensible ones—where the sentences in question are not of this sentence and its negation form. (For starters, "This mug is hot"/"That mug is cold," "I think it is raining"/"No it is not," "None of the philosophers is in his office"/"Bob is in his office," "Everybody is at home"/"Bob is in his office," and so on). I will be attempting to draw the line between the sensible and the defective disputes only with this very specific sentence-and-negation form—the presentation would otherwise be a bit of a nightmare, and the extension to the rest of the cases should be pretty straightforward.

(5) This means that the presence or absence of a genuine conflict between a particular assertion of S and a particular assertion of \neg S will not always go along with genuine conflict between *later* assertions of the same sentences. In the case of, for example, sentences containing temporal indexicals, like “now” (or context-sensitive expressions designed to refer back to previously uttered sentences, like “your last sentence”), subsequent assertions of the same sentence are likely to impose different conversational demands than those imposed by the original assertion.

(6) These facts about our actual practices of disputing about taste are another reason why the first pass, ‘*tasty*’-means-‘*tasty-to-me*’ version of contextualism about aesthetic predicates does not look very attractive.

(7) This is a feature that has loomed large in discussions of aesthetic predicates and claims involving them. (See, e.g., Wright (1992) on *cognitive command*.) It has also loomed large in discussions of secondary qualities more generally.

(8) And wanting to avoid complications about generics.

(9) I will not, at the end of the day, quite want to endorse this. I want to allow that the properties I need to self-attribute in order to accept some aesthetic assertions will be *idealization* properties—properties of not having a certain disposition D just now, but of being such that a suitably idealized version of me would have D. I also want to allow that some of the relevant properties are *group-membership* properties—properties of being a member of some group typical members of which have D. This maps onto Railton's discussion (2003) of the *vertical* and *horizontal* features of aesthetic properties. More on this in Section 4.

(10) Tamina Stephenson (2007a, b) gives a similar account when she explains how we ought to understand the formal apparatus of her (independent, and somewhat differently motivated) semantics for the vocabulary of personal taste.

(11) More generally, everybody taking themselves to have F or everybody taking themselves to have some incompatible G.

(12) Or perhaps *of kind K*, where K is some intrinsically desirable sort of experience.

(13) We can also use this to draw a distinction between cases that *are* mere change and those that are genuine improvements.

(14) Both of these are akin to phenomena about the offering of reasons that Bernard Williams (1995) discusses.

(15) Here again, Williams (1995) says some similar things about reason-talk. It is also possible to read the later bits of MacFarlane (2007) as proposing that we understand most, or all, disputing about taste in this sort of way, though I am not completely confident that this is what is intended.

(16) Again, it is easier to be sympathetic to the idea that some disputes about taste work this way if one thinks about long-term, running arguments over the course of days, weeks, months, or years.


(17) Many of these positive functions of various sorts of disputes will be familiar from Mill (1859/1978).

(18) I argue for a similar account of epistemic modals in Egan (2007).

(19) Though not the only way. See Perry (1979) for an alternative. I think that Lewis's way (1979) is preferable, but that is a topic for another paper.

(20) See Lewis (1979). See Nolan (2006) for dissent.

(21) See, e.g., Stalnaker (1978), Lewis (1979), Barker (2002), DeRose (2004), and Richard (2004) for discussions of this sort of phenomenon.



Disagreement

Richard Feldman and Ted A. Warfield

Print publication date: 2010

Print ISBN-13: 9780199226078

Published to Oxford Scholarship Online: September 2010

DOI: 10.1093/acprof:oso/9780199226078.001.0001

(p.287) Index

- abortion 48-51
- acceptance 64-6, 252-3, 278, 282
- accommodation 273
- Adler, Jonathan 60
- aesthetic vocabulary 251-3, 271, 282-3
 - first-personally committed uses of 252-3, 259-60
 - sympathetic uses of 252-3
- agreement 144, *see also* consensus
- alethic exclusivism *see* exclusivism
- Alston, William 155n.
- Anscombe, G.E.M. 274
- Antony, Mark 92
- asymmetrical no independent weight view 136-7, 143, 159
- Barker, Chris 281n.
- Barnes, Barry 188
- Bayesianism 120-1, 200, 230-2
 - Bayesian conditionalization 231
- Beckstead, Nick 159n.
- belief 63-6
 - all-or-nothing beliefs 117, 196n.
 - beliefs about other's beliefs 232-5
 - degrees of belief 117-18, 195-6
 - genesis of vs epistemic status of 102-3
 - voluntariness of 60-1, 63-4, *see also* 'ought' implies 'can'
- Bergmann, Michael 102
- best explanation arguments 85-6, 216
- Bloor, David 188
- bootstrapping 81, 125-32, 144, 160, 163, 165-6
- Boghossian, Paul 189-91, 192n., 193n.
- Boyd, Richard 218n. 2
- Broad, C.D. 207n.

- Bronfman, Aaron 125n.
- Chomsky, Noam 100
- Christensen, David 32, 34-5, 43, 45, 47, 54, 56-7, 113n., 115n., 119n., 138n. 23, 149, 158n., 167n., 175, 183n. 9, 208-9, 224-5, 227
- Churchland, Paul 101-2
- Cicero, Marcus Tullius 146n. 27
- Clifford, W.K. 27-8
- Coady, C.A.J. 241
- cognitive character 63
- cognitive defects 99-107, 109
- Cohen, G.A. 147-8
- Cohen, L. Jonathan 63-4
- coherentism 82, 104
- common consent argument 146-7
- conciliatory views of disagreement 175-82, 185-6
 - partially conciliatory views of disagreement 175-6, 184-6
 - self-undermining problem for 178-82, 184-6
 - stubbornness problem for 177-8
- conditionalization 230-2, 233-6, *see also* Jeffery conditionalization
- Conee, Earl 127n. 15
- conjunction rule 200-1
- consensus 40-2, 44, 46, 51, 54, 137-8, 145-7
 - aesthetic consensus 41
 - mathematical consensus 42, 44-5
 - philosophical consensus 42, 44-5, 145
- Consumer Reports 180-2, 184-5
- context-sensitivity 253n., 254-7, 274, 279-84
- convergence of opinion, *see* consensus
- conversational demands 255
- credulity:
 - principle of credulity 241
- Darwin, Charles 10n. 1, 17-18, 56
- Dawkins, Richard 10n. 1
- defeat 75-81, 88-9, 91-3, 97-9, 106, 108-9, 232
 - defeat from doubt principle 79
 - self-defeat, *see* self-undermining
- degree of truth 194
- Dennett, Daniel 10n. 1
- DeRose, Keith 281n. **(p.288)**
- Descartes, Rene 91, 99, 107, 241-2
- descriptive pluralism, *see* relativism
- disagreement:
 - about disagreement 175-186
 - broad disagreement 48-51, 165-6
 - defective disputes 247-8, 250-1, 253, 256-7, 261-4, 266-7, 276, 283
 - factually defective disputes 262
 - justificationally defective disputes 262
 - presuppositionally defective disputes 262
 - epistemically benign disagreements 53, 93-7, 99

- extent of 1
- extreme disagreements 95-7, 149
- intrapersonal disagreement 125
- merely possible disagreement 34-5, 39-40, 56-7, 221-223
- moral disagreement 216-20, 225, 227-8, 244
- persistent disagreement 33, 53, 66-7, 168, 237, 266
- revealed peer disagreement 3-4
- sensible disputes 257-8, 264, 267, 273, 276
 - conflict requirement 257-8, 284
 - worthwhileness requirement 257-8, 284
- value of 66-8
- dispositions 252, 259-62, 264-71
- distribution of opinion 40-1, 44-7, 49, 129, 146, *see also* consensus
- downward epistemic push 158-9
- doxastic attitudes, *see also* belief
 - psychological limitations of 195-6
- Egan, Andy 277n. 18, 280
- Einstein, Albert 56
- Elga, Adam 33, 47-9, 65n., 113n., 116n., 127n. 17, 128n. 18, 134, 149, 153, 160-3, 165-6, 170, 175, 178n., 183n. 9, 224-5, 227
- E-norms, *see* epistemic norms
- epistemic culpability 205
- epistemic egoism 237, 242, *see also* self-trust
- epistemic non-absolutism 190
- epistemic norms 197-202, 205, *see also* epistemic principles
 - fundamental norms 199-200
 - derivative norms 199
 - probabilistic norms 200
- epistemic peers 3, 31, 37-8, 48, 51, 53, 55-7, 70-1, 166-7, 189, 225-7, 236-7
 - competent epistemic peers 55-6
 - possible epistemic peers 57
 - epistemic peers on the topic 71, 226, 236
 - epistemic superiors 58-9, *see also* experts
 - moral epistemic peers 216
- epistemic objectivism 188-9, 196
- epistemic pluralism 191
- epistemic principles 230
 - general epistemic principles 230, 232n., 233-5
 - special epistemic principles 230, 232n., 234-5, 243
- epistemic relationism 191
- epistemic relativism, *see* relativism
- epistemic systems 187-8, 191-4, 196-201, 203, 205-6, 210, 213
 - ranking epistemic systems 194-5
- equal weight view 112-35, 141-2, 144-5, 149-62, 165, 167, 169, 178n.
 - problems with all-or-nothing attitudes 117
- E-systems, *see* epistemic systems
- evidence:
 - dialectical conception of 170-2
 - doxastic evidence, *see* psychological evidence

- evidence of evidence principle 76-8, 84-8, 209-11
- first-order evidence 86, 123-4, 139, 141-2, 144, 146, 148-50, 153, 156, 159-60, 169, 208
- higher-order evidence 86, 123-4, 139-44, 148-50, 158-9, 208-9, *see also* evidence of evidence principle
- incommunicable evidence 25-6, 70-1
- intuitions as evidence 89-90
- material evidence 208-9, 213, *see also* first-order evidence
- mathematical evidence 168-9
- non-psychological evidence 128-34
- norm evidence 208, *see also* higher-order evidence
- private evidence 85, 89-90, *see also* incommunicable evidence **(p.289)**
- privileged access to 90, 91, 98, 100-1
- psychological evidence 128-34, 141, 143, 148, *see also* higher-order evidence
- public evidence 25-6
- role of evidence 132-3
- sharing evidence 152n.
- evidentialism 70
- exclusivism 10, 14
 - philosophical weak exclusivism 15, 21-22
 - philosophical strong exclusivism 17, 19-22
 - political strong exclusivism 18
 - scientific strong exclusivism 17
 - strongly exclusivist religions 13-14
 - very strongly exclusivist religion 13-14
 - weakly exclusivist religions 11-13, 15
 - weak exclusivism entailing strong exclusivism 21-4, 26-7
- experts 40-1, 43-6, 51-2, 56, 225
- extra weight view 153, 161-2
- extreme permissivism 188n.
- Feldman, Richard 31, 34, 54, 56-7, 59, 61, 65, 73n., 76n., 112-13, 117n., 118n., 119n., 127n.15, 17, 158n., 175, 183n. 9, 189-91, 210-11, 224, 225n. 5, 227
- Field, Hartry 149n., 185n.
- Foley, Richard 54
- Fox, Danny 251
- Frances, Brian 225n. 5
- Fricker, Elizabeth 232
- Fumerton, Richard 104n. 4
- Gassendi, Pierre 146n. 27
- geniuses, *see* experts
- genuine conflict 255-6
- Gibbard, Allan 237-9, 241-2
- Goldman, Alvin 95n., 146n. 26, 193-4, 211n. 17
- Goodwin, Brian 17-18
- Gödel, Kurt 25, 42
- Greco, Daniel 164n.
- groupthink 177, *see also* independence
- Grotius, Hugo 146n. 27
- Gutting, Gary 31n. 1

- Harman, Gilbert 216–17, 271
- Hills, Alison 227n. 8
- Hume, David 63, 247–8, 269–70
- Huemer, Michael 212
- incommensurable insights 70–1, *see also* evidence
- incompatibility, principle of 263
- independence 47, 146–8, 177–8, 185, 223, 225, 232, 236, 238, 243
- inductive method 179–82
- intellectual peers, *see* epistemic peers
- internalism 230, 232–3, 240
- Jeffrey conditionalization 231–4
- Jeffrey, Richard 231
- Jones, Karen 235n.
- justification:
 - doxastic justification 202–3
 - egocentric conception of 91–3, 106
 - higher-order justification 204–7, 213
 - propositional justification 202–3
- Kahneman D. 200
- Kalderon, Mark 227n. 8
- Kaplan, David 279
- Kelly, Thomas 34–41, 44, 47, 54–5, 58, 61, 111n., 115n., 116, 121n., 138n. 24, 145, 177n., 179n. 4, 183n. 10, 209n., 212, 225n. 6
- knowledge:
 - self-knowledge 103
- Kierkegaard, Søren 11
- Kitcher, Philip 67
- Klein, Peter 88
- Kornblith, Hilary 54, 65n.
- Lackey, Jennifer 199
- Laserson, Peter 247
- Latar, Apolonio 210
- Lehrer, Keith 225n. 5
- Lewis, David 23–4, 26–8, 46n. 25, 58–60, 180n. 7, 277, 280n., 281n.
- Locke, John 146n. 27
- MacFarlane, John 275n.
- Mill, John Stuart 146n. 27, 276n. 17
- Miller, Richard W. 218n. 3
- Monty Hall puzzle 94, 96, 105–6
- moral epistemology 217–18
- moral evil demon 216, 220–6, 228, 244 **(p.290)**
- moral intuitions 217, 219–21, 223–4, 226, 239–41, 243–4
 - misleading intuitions 219–21, *see also* moral evil demon
- moral skepticism 217, 221–3, 244
- neural quirk 27
- Neta, Ram 191n.
- Newcomb problem 39–42
- no independent weight view 115, 135–6, 141–2
- Nolan, Daniel 280n.

- non-superficial defectiveness, principle of 263-4
- Nozick, Robert 39
- 'ought' implies 'can', principle of 60, 71n.
- partially conciliatory views of disagreement, *see* conciliatory views of disagreement
- Pascal, Blaise 61, 275
- peer 2-3, *see also* epistemic peer
- Perry, John 277n. 19
- Pettit, P. 208n., 225n. 5
- phenomenal conservatism 212, *see also* credulity
- Plantinga, Alvin 115n., 225n. 5
- possible mutual correctness, principle of 263
- Pryor, James 79n., 127n. 16, 171n.
- quad, case of person on the 31-3, 36
- Railton, Peter 260n. 9, 269-70
- rationality:
 - doxastic rationality 61, *see also* justification
 - practical rationality 61
 - rational doxastic attitudes 70
 - rationality from rationality principle 74-6
 - reason from rationality principle 77-82, 84
- rational uniqueness, *see* uniqueness thesis
- recognition 156-7
- Reed, Baron 199
- reflective equilibrium 200-1
- Reid, Thomas 241n.
- relativism 187-92, 196, 201-2, 213, 229, 244
 - anti-relativism 216-7
 - descriptive pluralism 187
 - moral relativism 216-17, 219
 - nihilistic relativism 188-91, 202
 - non-nihilistic relativism 192
 - objectivity-based relativism 188-9, 201-3
- reliability 97-99, 107-9, 194, 222
 - reliability relative to a field 107-9
- restaurant check case 32-3, 36, 43, 45, 47-9, 167n.
- Richard, Mark 281n.
- right-reasons view of disagreement 183-4
- Rorty, Richard 27
- Rosen, Gideon 115 n., 120n., 191n., 225n. 5
- safety 222
- St Paul 13-14
- self-attribution 260-1, 263, 267-8, 270, 272, 276-9, 281-2, 284
- self-trust 237-40, 243-4
- self-undermining 178-82, 184-6
- Sellars, Wilfred 102
- semantic denial 255
- Seneca 146n. 27
- Shaked, M. 194
- Sibley, Frank 253

Sidgwick, Henry 221, 223-5, *see also* Sidgwick's principle
Sidgwick's principle 221, 223-5, 227, 229, 232-4, 237, 239
similarity 262, 264-9, 271-3, 277, 283-4
Smith, Holly 207n.
Sorensen, Roy 114n., 154n. 32
Sosa, Ernest 92, 141n.
spinelessness 57, 165
splitting the difference 112, 115-16, 122, 124, 126, 133, 135, 149, 151-2, 155, 166, 168, 203, 210, *see also* equal weight view
Stalnaker, Robert 277-8, 281n.
Stephenson, Tamina 251n., 260n. 9
Stich, Stephen 40 **(p.291)**
stubborn view of disagreement 176, 183-4
stubbornness 57-8, 176-8
superficial denial 254
symmetrical no independent weight view 116, 121, 135-6
Thompson, Judith Jarvis 216-17, 271
total evidence view 115, 135, 142, 144-5, 149-54, 157-63, 167-70, 172
 and externalism 169-70
Tverski, A. 200
underdetermination 229
uniqueness thesis 6, 21, 69, 71-2, 80-1, 83-90, 119-22, 136, 157, 164, 190-1, 202n. 10, 207-9
van Inwagen, Peter 10n. 2, 58-9, 70, 115n., 127n. 17, 225n. 5
Velleman, David 274
Vogel, Jonathan 81n.
von Fintel, Kai 251n.
Vranas, Peter 79n.
Weatherson, Brian 176n., 179n. 4, 225n. 5
Wedgwood, Ralph 217n., 222n., 230n., 242
White, Roger 33, 119n., 120n. 11, 121n., 188n., 196
Williams, Bernard 60, 273n., 275n.
Williamson, Timothy 69n., 89n., 106, 141n., 170, 171n., 184, 222n.
Wright, Crispin 229n., 258n. **(p.292)**