The Limitations of the Open Mind

Jeremy Fantl

For Norah and Toby

Table of Contents

Preface i

Part I: The Epistemology of Open-Mindedness

Chapter 1: Open-Mindedness

1.1 Open-Mindedness Toward Arguments 1
1.2 Conceptions of Open-Mindedness 4
1.3 Open-Mindedness and Outright Belief 21
1.4 Conclusion 32

Chapter 2: A Defense of (a different kind of) Dogmatism

2.1 Knowledge and Closed-Mindedness 34
2.2 Forward-Looking Dogmatism 37
2.3 The Case for Dogmatism^{FL} 44
2.4 Conclusion 64

Chapter 3: The Epistemic Efficacy of Amateurism

3.1 Coins and Counterarguments 65
3.2 Higher-Order Evidence 69
3.3 The Epistemic Efficacy of Amateurism 73
3.4 The Novice Knowledge Principle and the Case for the Other Side 95
3.5 Conclusion 99

Chapter 4: Psychic Phenomena and the Existence of God

4.1 Sources of Knowledge 103
4.2 Felt Obviousness 105
4.3 Prima Facie Goat and Atheistic Knowledge 114
4.4 Goat and Atheistic Knowledge, Undefeated 118
    4.4.1 The Fine Tuning Argument 118
    4.4.2 Statistical Arguments for Psi 130
4.5 Conclusion: the Position of the Expert 138

Part II: The Ethics of Participation in Argumentation

Chapter 5: The Obligation to Engage

5.1 Introduction 141
5.2 Is the Obligation a Role Ought? 144
5.3 The Duty to Maximize Good Epistemic Consequences 160
5.4 The Rights of Disputants 168
5.5 Conclusion 175

Chapter 6: Against Open-Minded Engagement (for Some People)

6.1 Introduction 177
6.2 The Case Against Open-Mindedness 178

6.3 From Knowledge to Closed-Mindedness 182
6.3.1 Worries about the linking premise 182
6.3.2 The argument for the linking premise 188
6.4 Dealing With the Arguments for Open-Minded Engagement 208

Chapter 7: Against Closed-Minded Engagement (in Some Situations)

7.1 The Pitfalls of Closed-Minded Engagement 215
7.2 The Pitfalls of Insincerity 222
7.3 The Pitfalls of Sincerity 238

Chapter 8: On Inviting Problematic Speakers to Campus

8.1 Introduction 247
8.2 Learning from Problematic Speakers 250
8.3 The Case Against Inviting (Some) Problematic Speakers 256
8.3.1 Problematic speakers and psychological harm 256
8.3.2 Problematic speakers and intrinsic harm 263
8.3.3 Problematic speakers and the respect for truth 277
8.4 Conclusion 281

References 284

Preface

> Mankind is so prone to prejudice, party bias, collective hysteria and unthinking acceptance of propaganda, that there is almost always and almost everywhere too little of the open mind, not too much. While, therefore, truth compels me to acknowledge the theoretical limitations of the open mind, in practice I should find it almost always wise to advocate it, since the utmost advocacy is not likely to produce enough of it, and since its absence is one of the chief reasons for the appalling dangers by which, in our age, the human race is beset. (Russell 1950, 39)

In his 2016 keynote address to the Society of Christian Philosophers (SCP), Richard Swinburne argued, among other things, that "homosexual sexual acts" are morally wrong because God has forbidden them; that homosexuality should be actively discouraged and treatment should be sought; and that the "husband is head of the family, and so wife and children have an obligation to obey him."

Reactions to the talk varied. But many of Swinburne's defenders made the same accusation: Swinburne's critics did not engage with Swinburne's arguments. One critic, it was claimed,

> like all other touchy-feely leftists, *produces no argument against [Swinburne's premises]*. Instead [sic.] all that is offered are claims that this contention is offensive, coupled with some nauseating pearl clutching. (Martel, Conservatrarian, et al. 2016, emphasis added)

Another critic

> *doesn't take on Swinburne's arguments at all*. This is entirely a 'who/whom' statement. (Dreher 2016, emphasis added)

A participant during the Q&A session following Swinburne's talk

> *never addressed [Swinburne's] actual arguments* but simply made this indignant accusation, to which Swinburne responded with admirable patience and grace. (Spiegel 2016, emphasis added)

And, finally,

> I, and many others in the room, had problems with Swinburne's paper. But it's Philosophy. Don't disrespect the opinions of someone else by getting hurt—*properly critique their argument and explain why they're wrong*. (Anonymous 2016, emphasis added)

The italicized demands are directed toward those who are objecting to Swinburne's presentation or publicly expressing views in opposition to Swinburne's conclusions. However, there may be some obligation to engage with arguments against what you believe, even if you aren't publicly condemning those arguments or expressing opposition to their conclusions. This seems to be the lesson Barack Obama tried to teach in a 2015 town hall speech in answer to a question from the audience:

> I've heard of some college campuses where they don't want to have a guest speaker who, you know, is too conservative. Or they don't want to read a book if it has language that is offensive to African-Americans, or somehow sends a demeaning signal towards women. I don't agree with that, either. I don't agree that you—when you become students at colleges—have to be coddled and protected from different points of view. . . .
>
> Anybody who comes to speak to you and you disagree with, you should have an argument with them. But you shouldn't silence them by saying, "You can't come because, you know, my — I'm too sensitive to hear what you have to say." That's not the way we learn, either.[1]

I take it that the ideal of argumentation, according to Obama, is something like this: if someone is expressing some offensive speech or making an argument for a problematic claim, you are either able to figure out where the speech goes wrong, or not. If you can figure out where it goes wrong, then you should argue with them so that you can correct the speech, change people's minds, and stop the spread of the falsehood. If you can't figure it out, then it just might not be wrong, and so you should engage with the speech in order to learn from the speaker. Presumably Obama would add that arguing with the person is the way that you figure out in the first place what's wrong with the speech.[2] Either way, you should engage.

---

[1] See (Office of the Press Secretary 2015).

[2] As Mill (1859/2003) says,

Obama’s argumentative ideal is not one we always live up to. In a 2014 study, Deidre Le Fevre, Viviane Robinson, and Claire Sinnema detail the difficulty people have participating in what they call “genuine inquiry” when dealing with cases of disagreement and difficult interpersonal relationships. Genuine inquiry is what they call personal interactions that are “motivated . . . by a desire to learn,” that come with a recognition that “one’s beliefs could be incomplete and misinformed,” and that involve “a willingness to think again despite having formulated a view” (2). Le Fevre et al. do not specify whether the obligation to participate in genuine inquiry is moral or epistemic, and I defer discussion of the status of any putative obligation until Chapter 5. But it is clear that, according to Le Fevre et al., it is an obligation; there is a normative requirement to participate in genuine inquiry. It’s what inquiry should be striving toward. According to Swinburne’s defenders, many of his critics did not approach the ideal of genuine inquiry.

Mere engagement with relevant counterarguments may not be sufficient. If you violate an obligation by failing to engage with some relevant counterargument, you seemingly also violate an obligation if you engage with the counterargument but fail to respond adequately. An adequate response might involve exposing some flaw in the argument—a false premise or invalid step. If you cannot do that, an adequate response might be the reduction of confidence in your position. Therefore, there are two separate argumentative obligations you might have: first, an obligation to engage with relevant counterarguments and, second, an obligation to reduce your confidence when, upon hearing out a counterargument, you find the steps individually

---

> There is the greatest difference between presuming an opinion to be true, because, with every opportunity for contesting it, it has not been refuted, and assuming its truth for the purpose of not permitting its refutation. Complete liberty of contradicting and disproving our opinion, is the very condition which justifies us in assuming its truth for the purposes of action, and on no other terms can a being with human faculties have any rational assurance of being right. (89)

compelling and you are unable to expose a flaw. If you have a public obligation to engage with relevant counterarguments, you often have a private obligation to do so *open-mindedly*.

The main conclusion of this book is that in some important and standard situations, Swinburne's keynote included, you shouldn't engage with relevant counterarguments open-mindedly. A secondary conclusion is that, in many of those situations, you shouldn't engage closed-mindedly, either. Hence, often you shouldn't engage with relevant counterarguments at all. This is not just because, in those situations, you aren't qualified to or aren't in a position to engage, or because you lack the time or resources. It's because you often can know that an argument is misleading, whether you fail to know the details of that argument or you know the details of the argument and are unable to expose a flaw. If you know that an argument is misleading, in many cases, you shouldn't hear out the argument if you haven't already, you shouldn't spend time publicly or privately discussing the argument with the one presenting it, and you shouldn't engage with the argument open-mindedly if you engage with it at all.

The argument requires three main premises:

1) There are important and standard situations in which you know that a relevant counterargument is misleading whether or not you have spent significant time with the argument, found each step compelling, and been unable to expose a flaw.

2) If you know, in a standard situation, that a relevant counterargument is misleading, you shouldn't engage with it open-mindedly.

3) If you know that a relevant counterargument is misleading there are important and standard situations in which you shouldn't engage with it closed-mindedly.[3]

The book is divided into two parts. The first part is epistemological and primarily devoted to the clarification and defense of the first premise and elucidation of the notions of closed-mindedness and open-mindedness. The second part of the book is ethical and devoted to defense of the second and third premises—the moral, social, and political consequences of the fact that you can know that relevant counterarguments are misleading even if you don't know where they go wrong.

---

[3] It is also required that a sufficient number of the situations that make premise 1 true are also situations that make premise 3 true.

In Chapter 1 I argue for what I call "Platonic" conceptions of open-mindedness and closed-mindedness. Open-mindedness is not simply a matter of being willing to change your mind in the face of relevant counterarguments. Willingness to change your mind is always conditional on a relevant counterargument having certain features. You have to be willing to change your mind conditional on spending significant time with the argument, finding each step compelling, and being unable to expose a flaw. If you are unwilling to reduce your confidence in your position conditional on spending significant time with the argument, finding each step compelling, and being unable to expose a flaw, then you are closed-minded toward the argument. If you are willing to do this, then you may be open-minded toward the argument provided you don't violate various procedural norms and aren't disposed to allow various affective factors to influence your beliefs.

On this Platonic conception of open-mindedness, we can explain how it's possible to hold an outright or "full" belief even while being open-minded toward arguments against that belief. It's possible to fully believe something while recognizing that you might not be able to figure out what's wrong with arguments to the contrary. If you also are willing to reduce your confidence if you can't figure out what's wrong with some counterargument, then you count as open-minded on the Platonic conception. Whether it's *rational* to maintain full belief while being open-minded in this way is another question, and the one to which much of the rest of the book is devoted.

In the second chapter I argue that Premise 1 is true: you can retain knowledge even if you manifest closed-mindedness toward a relevant counterargument—even if you spend significant time with an argument, find each step compelling, and are unable to expose a flaw, but nonetheless do not lose confidence that your position is true. When you are in this situation, to use the technical term I adopt, you retain confidence despite finding the counterargument *apparently flawless*.

There are uncontroversial cases about which everyone should grant that knowledge can survive finding a counterargument apparently flawless: trick arguments that 1=0, paradoxical arguments that nothing moves and that no one is bald, perceptual arguments that a rabbit has spontaneously appeared in a previously empty hat. But, I argue in Chapter 2, the premise is also true of counterarguments that concern controversial matters of ethics, politics, religion, science, and philosophy. This is because you often have what would otherwise count as knowledge-level positive support for your controversial beliefs. And this support is often not defeated by what turns out to be the rather weak evidence provided by the fact that you find a relevant counterargument apparently flawless.

Even if knowledge can survive exposure to apparently flawless counterarguments, it doesn’t always survive. Sometimes exposure to apparently flawless counterarguments destroys knowledge. What kinds of apparently flawless counterarguments does knowledge survive exposure to? I give a partial answer to this question in Chapter 3. The apparently flawless counterarguments with the least defeating power are those that it would be unsurprising to find apparently flawless even if they were misleading. For example, if a knower is unversed in Bayesian epistemology, it is unsurprising that the knower could be fooled into finding an argument that makes use of Bayesian techniques apparently flawless even if the argument is misleading. More generally, it’s unsurprising that knowers can be fooled by counterarguments that employ evidence-types or argumentative moves that the knower lacks sufficient expertise to reliably evaluate. Therefore, knowers should be inherently distrustful of their reactions to such arguments and can retain knowledge in the face of such arguments as long as their knowledge has a basis with which the knower does have sufficient facility. Amateurism, I argue, can be epistemically efficacious.

This is why it is so important in academic work to give a full treatment of the case for the opposition. It’s not simply because opponents are more easily convinced when they see you show respect for their position. It’s because if you do a good job presenting the case for the opposition, you train your reader in the skills and background knowledge required to reliably evaluate your positive argument. If you train your reader adequately, and they *still* find the steps in your argument compelling and are unable to locate a flaw, then it becomes harder for them to closed-mindedly dismiss your argument while retaining knowledge that you’re wrong.

In Chapter 4 I apply the conclusions of the previous chapters to two cases of controversial belief: atheistic belief and the denial of psychic phenomena (psi). In both cases, the felt obviousness—to the denier—of the non-existence of what they’re denying figures large in the production of their knowledge, if they have it. Felt obviousness can be knowledge-generating. This might be because it channels a lifetime of empirical support for what feels obvious, or because it provides *a priori* evidence, or is the guise in which our fundamental worldviews are presented to us, if such fundamental worldviews can be sources of knowledge.

Of course, felt obviousness does not always ground knowledge. Sometimes it can be defeated by contrary evidence or relevant counterarguments. In the case of atheistic belief, relevant counterarguments come in the form of the various theistic arguments (I focus on the fine-tuning argument). In the case of the denial of psychic phenomena, relevant counterarguments come in the form of meta-analyses of a number of controlled studies that purport to demonstrate the existence of psychic phenomena. Often, these arguments survive immediate objections to them, and the more refined objections may be beyond the atheist's or the psi-denier's abilities to evaluate. But, given the conclusion of Chapter 3, this need not destroy prior knowledge that there is no God or that there is no such thing as psychic phenomena. Prior knowledge can survive as long as atheists and psi-deniers have facility with the bases of their beliefs that they lack when it comes to the evidence-types and argumentative moves involved in the relevant counterarguments.

The second part of the book is devoted to the consequences of the epistemological conclusions of the first four chapters. Here's the primary thesis: because there are standard situations in which you know some counterargument is misleading, there are standard situations in which you shouldn't engage with that counterargument either open-mindedly or closed-mindedly. You shouldn't engage open-mindedly because to do so is to be willing to reduce your confidence in response to an argument you know has a false conclusion. You shouldn't engage closed-mindedly because in some standard situations the only way to do so requires you to be problematically insincere, disrespectful, or ineffective.

This conclusion runs afoul of principles common to much work in informal logic, political theory, and ethics: that you have at least *pro tanto* obligations to engage even with those arguments and opinions you find obviously false and abhorrent. In Chapter 5 I develop three general defenses of this standard position (which I go on to refute in the subsequent chapters). First, an obligation to engage with relevant counterarguments can be sourced in the nature of the practice of argumentation itself. One potential reason this is so is that argumentation essentially aims at rationally convincing others. But we can't rationally convince others unless we address their concerns, objections, and counterarguments. Therefore, insofar as you are participating in argumentation, you ought to address relevant counterarguments. My focus is on similarities between work in the informal logic literature on norms of argumentation and Timothy Williamson's seemingly unrelated work on norms of assertion.

Second, an obligation to engage with relevant counterarguments can be sourced in the good consequences that result from such engagement—in particular, the good epistemic consequences for you or the society at large. This source derives in part from Mill's argument for universal freedom of expression. Because the only way to ensure we are getting truth and avoiding falsehood is to allow unlimited free expression even of views we find abhorrent or false, and because this good is only achieved if those expressions are listened to and engaged with, there is just as much an obligation to engage with relevant counterarguments as there is an obligation to allow them to be heard.

Third, an obligation to engage with relevant counterarguments can be sourced in the rights or agency of the person offering the counterargument. Following Rawls and others, we might say that this is because it is only legitimate to impose legal restrictions on them if their voices had a chance to influence which restrictions were the outcome of the democratic process. Or we might say that people—in virtue of the fact that they are agents—have the right to be transmitters of their knowledge. If so, then if some agent is everywhere ignored, they seemingly have been wronged. This may ground, if not individual obligations to engage with relevant counterarguments, at least a group obligation to do so.

In Chapters 6 and 7 I reject these arguments and the conclusions supposed to follow from them. In Chapter 6 I argue that in standard situations you shouldn't engage with relevant counterarguments open-mindedly if you know they are misleading. Because, given the results of the first half of the book, there are standard situations in which you know relevant counterarguments are misleading, it follows that there are standard situations in which you shouldn't engage open-mindedly. This conclusion relies on a general principle: you should do whatever your knowledge a decisive reason to do; you should act in accordance with what you know. This principle doesn't say that you should do whatever there is a decisive reason to do. If there are decisive reasons to do something, but you are unaware of those reasons, there is an important sense of "should" according to which it's not the case that you should do that thing. But on this principle, if you know to be true the decisive reason to do something, then you should do it.

That a relevant counterargument is misleading is a decisive reason, in standard situations, to be unwilling to reduce your confidence in response to the counterargument. If you're willing to reduce your confidence in response to a misleading argument, you're willing to reduce your confidence in response to an argument whose conclusion is false. This is not, in standard situations, a good thing to do. Of course, you may be unaware that the argument is misleading; you may even rightly worry that it is sound. In that case, arguably, you should be willing to reduce your confidence in response to the counterargument. But, by the general principle just alluded to, if you *know* that the counterargument is misleading, you should do whatever that fact is a decisive reason to do: be unwilling to reduce your confidence in response to the counterargument, even if you will find it apparently flawless. Given the account of open-mindedness developed in Chapter 1, it follows that if you know that a counterargument is misleading, you shouldn't engage with it open-mindedly.

Nor should you, in many standard situations, engage with relevant counterarguments closed-mindedly. This case I develop in Chapter 7. There are situations in which, knowing that a counterargument is misleading, you have an obligation to engage with the counterargument in order to call out offensive or demeaning speech, correct easily correctable errors, or stand in solidarity with marginalized groups. But in many standard situations closed-minded engagement is problematically demeaning or deceptive. In other standard situations closed-minded engagement can give credibility to views you know to be false, demeaning, or dangerous. Far from stopping the spread of false or dangerous views, closed-minded engagement can often facilitate their spread.

There are a number of influential public engagement methods that land in either of the pitfalls of closed-minded engagement. There are two I especially focus on in Chapter 7. The *Listening Project* emphasizes non-judgmental question asking so that interlocutors can lead themselves into revising their problematic views. If effective, such techniques exploit misconceptions your interlocutors have about you in order to get them to change their attitudes. This is no more appropriate than misrepresenting yourself as a climate scientist in order to convince an audience of the perils of global warming. Of course, you could avoid deception by changing your attitudes—by becoming actually open-minded toward the views and arguments of those whose minds you're trying to change. But if you know that what they're saying is false, then—by the arguments of Chapter 6—you shouldn't be open-minded toward what they're trying to say.

Unlike the Listening Project, the methods of so-called *Nonviolent Communication* emphasize honest sharing of your needs and beliefs to make it easier for your interlocutors to share with you their needs and beliefs. Ideally, the process ends at a mutually agreeable meeting point. There is some evidence, however, that the kinds of honest but compassionate methods emphasized by Nonviolent Communication are ineffective, and those methods that are similar to

Nonviolent Communication but modified enough to make them demonstrably effective, slide back into potentially problematic deceptiveness.

I conclude that in many situations there is no way to closed-mindedly engage with interlocutors while being *non-manipulatively effective*. In many of those situations, this precludes the permissibility of closed-minded engagement. Given the conclusions of Chapter 6, in some of those situations you shouldn’t open-mindedly engage, either. Therefore, in some standard situations, you shouldn’t engage with relevant counterarguments at all, even if you’re qualified and well-positioned to do so, and you don’t lack for time and resources.

In the final chapter I apply these conclusions to a specific kind of public engagement—the decision to invite problematic speakers to campus. Speakers can be problematic in virtue of marginalizing or offensive content that you know will be part of their speech, or they can be problematic in virtue of their past verbal or non-verbal behavior. It is often, I argue in Chapter 8, morally impermissible to invite either kind of speaker if you know that their past or future behavior reflects accurately on their current attitudes or dispositions. This is because to do so can require you to stand in solidarity with the problematic speaker and thereby stand against those who have been or will be victimized by their speech or behavior. Inviting a speaker to campus comes with certain obligations to the speaker—obligations of politeness and respect. Because it is impermissible to satisfy those obligations to certain kinds of problematic speakers, it is impermissible to invite those speakers in the first place.

It is often argued by opponents of campus speech codes that students who object to various problematic speakers only do so to protect their fragile psyches, and so it is important to invite precisely those problematic speakers in order to teach students resilience and the best ways to deal with the inevitable conflicts they will experience upon leaving academia. This argument is wrong in all its steps. The primary problem with the argument is that students who are psychologically harmed by invitations to problematic speakers, along with the respectful behavior extended to those speakers, are in fact accurately perceiving something like a betrayal. Therefore, any attempt to build resilience in such a way that students no longer have this reaction results in students who do not accurately perceive when they are being intrinsically harmed. Students are right to feel betrayed in many of these contexts, because they are accurately responding to what is an actual betrayal.

Acknowledgements...

1

Open-Mindedness

> . . .we have set sail, and must go where the wind, or the argument, blows us. (Plato 2000, 394d)

1.1 OPEN-MINDEDNESS TOWARD ARGUMENTS

You can be open-minded toward new experiences, like the taste of unfamiliar foods, and new activities, like skydiving. You can be open-minded toward other cultures and other people’s behavior. In your cognitive life you can be open-minded toward opinions, ideas, concepts, and worldviews.[1] In this chapter I develop an account of what it is to be open-minded toward arguments.

Like courage, open-mindedness is often thought of as a standing character trait. Generally open-minded people are prone to be open-minded on specific occasions. But open-minded people can behave closed-mindedly and closed-minded people can behave open-mindedly. Because it’s possible to be open-minded on specific occasions while lacking the standing character trait, an account of the trait isn’t necessary for a proper account of isolated instances of open-mindedness. My focus is on the latter. What is it to be open-minded toward some specific argument?

An account of open-mindedness should start from paradigm cases.[2] Open-mindedness is what’s missing from some beliefs about vaccinations, many parents’ insistence that their child is innocent of a crime, some fans’ commitment to the innocence of their favorite athlete (say, Maria

---

[1] See (Baehr 2011, 145 ff.) for examples in which subjects are open-minded with respect to cognitive situations that don’t require evaluating arguments.

[2] Berislav Marušić makes a comparison to courage. Our concept of courage is bound up with specific cases: heading into battle, running into a burning building, etc. It’s not clear that any formal account that abstracts from specific cases is sufficient. My reaction to this comparison is the same as my reaction to the similar view about formal accounts of open-mindedness. Only certain features of running into battle or running into a burning building ground our judgments that those who do so are courageous.

Sharapova,[3] Tom Brady,[4] or Peyton Manning[5]), much support of Donald Trump, etc. But it's not every aspect of paradigm cases that contributes to those cases exemplifying closed-mindedness. We should look for those features of paradigm cases that ground the judgment that open-mindedness is absent. That's what I do in this chapter.

I share a majority view that to be open-minded is to have a disposition or a "willingness" to do something:

> Open-mindedness [is] *engagement*, that is, a willingness to make room for novel ideas in one's cognitive space and to give them serious consideration. (Kwong 2016b, 71)
>
> What [open-mindedness] does demand, is a readiness to re-examine these views in the face of alternative information and new evidence. Such re-examination involves a willingness to think again despite having formulated a view. (Le Fevre, Robinson, et al. 2014, 2)
>
> Open-mindedness appears to be a cognitive disposition: an open-minded person is disposed to gain, lose, and revise beliefs in a particular, reasonable way. (Arpaly 2011, 75)
>
> An open-minded person is characteristically (a) willing and (within limits) able (b) to transcend a default cognitive standpoint (c) in order to take up or take seriously the merits of (d) a distinct cognitive standpoint. (Baehr 2011, 152)
>
> To be open-minded is to be aware of one's fallibility as a believer, and to be willing to acknowledge the possibility that anytime one believes something, *it is possible that one is wrong*. (Riggs 2010, 180)
>
> Open-mindedness has to include the willingness . . . to re-examine one's own beliefs and, if called for, to let them go. (Cohen 2009, 56)
>
> Properly understood, open-mindedness is a fundamental intellectual virtue that involves a willingness to take relevant evidence and argument into account in forming or revising our beliefs and values . . . (Hare 2003, 76)

---

[3] Tennis pro Maria Sharapova was initially banned from International Tennis Federation events for two years following positive tests for the prohibited drug meldonium (Ubha 2016). The suspension was later reduced to 15 months (Wilson 2016).

[4] NFL quarterback Tom Brady was suspended for four games for his alleged involvement in the alleged deflation of footballs below minimum required levels prior to the New England Patriots' 2015 AFC Championship Game against the Indianapolis Colts (Manfred 2015).

[5] Al Jazeera America reported in 2015 evidence indicating that NFL quarterback Peyton Manning had taken the banned substance HGH in 2011 (Hackman 2016). Manning was cleared by the NFL in 2016 (Schefter 2016).

> Open-mindedness should be conceived . . . as requiring not neutrality or indecision about beliefs but a willingness to subject them to rational formation. (Hare and McLaughlin 1994, 241)

These views include all the following among the things you must be willing[6] to do in order to be open-minded:

1. Acknowledge the possibility that anytime you believe something, it is possible that you are wrong.
2. Transcend a default cognitive standpoint.
3. Give serious consideration to novel ideas.
4. Re-examine your views in the face of alternative information and new evidence.
5. Think again despite having formulated a view.
6. Take relevant evidence and argument into account in revising your beliefs and values.
7. Take relevant evidence and argument into account in forming your beliefs and values.
8. Gain beliefs in a particular, reasonable way.
9. Subject beliefs to rational formation.
10. Lose beliefs in a particular, reasonable way.
11. Revise beliefs in a particular, reasonable way.
12. If called for, let your beliefs go.

Some of these (nos. 8–12) have implications for what you believe. They require you to form certain beliefs or get rid of certain beliefs, depending on what is rational or called for. Others have no necessary implications for what you believe; they require only that you adopt certain cognitive habits (nos. 3–7) or have certain attitudes toward your beliefs (no. 1 and maybe no. 2).

---

[6] I discuss later in this chapter how willingness relates to dispositions and to what extent being open-minded requires various dispositions. Short answer: I remain neutral on whether open-mindedness requires a mere disposition or something more robust. From here forward I use the language of "willingness," while leaving it open whether willingness is anything more than a mere disposition to do what you're willing to do.

Are all these things required? Are any? In the first section of this chapter I proceed through some inadequate accounts of open-mindedness before arriving at the Platonic conception I endorse. In the second section I compare the Platonic conception to a competitor and argue that the Platonic account is preferable.

1.2 CONCEPTIONS OF OPEN-MINDEDNESS

It's not enough to be open-minded toward an argument that you are willing to hear the argument out; steadfast dogmatists can be willing to allow arguments to enter their ear canal. Nor is it enough that you are willing to believe the conclusion of the argument if you hear it out and find it convincing; steadfast dogmatists can be willing to do that, too. The problem with steadfast dogmatists is that they often should find arguments convincing but they don't.

This suggests that what it takes to be open-minded toward an argument is a disposition to adopt the rational response to the argument—to have whatever attitude the argument makes rational. As Hare says, open-mindedness "entails being prepared to take *appropriate* steps towards arriving at reasonable and justifiable conclusions" (76, emphasis added). Let's take this as a first stab at an account of open-mindedness toward an argument:

> (Open-Mindedness-as-Rational-Response) You are open-minded toward an argument iff you are disposed to respond rationally to the argument.

When I ask you to be open-minded toward an argument, I'm not asking you to automatically give up your belief in response to the argument.[7] When you promise to hear me out open-mindedly, you're not thereby promising to change your mind no matter what. What am I asking, and what are you promising, if not to change your mind? Open-Mindedness-as-Rational-

---

[7] In some contexts, my asking you to be open-minded implies that you are not responding rationally and that you should. In other contexts, my asking you to be open-minded lacks the implication that you are failing in any way. For example, I might ask you to keep an open mind right before I present what I know is a fantastical suggestion.

Response provides the answer; I'm asking you to respond to the argument as rationality requires, and you're promising to do so.

Still, Open-Mindedness-as-Rational-Response is not an adequate account of open-mindedness. This is why, when Wayne Riggs says of Hare's view that, "as Hare describes it, open-mindedness seems nothing short of rationality itself" (2010, 179), it's meant as a complaint or an objection. Equating open-mindedness with rational response, as Riggs puts it, "runs afoul of the first desideratum of an interesting view of open-mindedness: that it be a distinct and specific kind of excellence" (2010, 179).[8] When I tell you to be open-minded, I don't just tell you to avoid logical fallacies or to make your confidence match the evidence. I give you *advice*—advice that's more substantial than the advice to merely respond rationally.[9] The danger of blurring different ways you might fail to be rational in responding to an argument is that injunctions to be open-minded won't have the substance we take them to have.

Riggs' worry suggests a problem with the right-to-left direction of Open-Mindedness-as-Rational-Response. If Riggs is right, there is something more to open-mindedness than a disposition to respond rationally to an argument; open-mindedness is rational response *plus* something. There is a companion worry to Riggs' that counts against the left-to-right direction of Open-Mindedness-as-Rational Response. If open-mindedness is "a distinct and specific kind of excellence," then closed-mindedness is a distinct and specific kind of imperfection. There are lots of different kinds of irrationality and not all of them are failures to keep an open mind. You can respond irrationally to an argument by affirming the consequent, failing to see some obvious link between premise and conclusion, committing the base rate fallacy, or making other well-meaning but irrational mistakes that cause you to irrationally underrate the plausibility of the argument's conclusion. In these cases, you can still be open-minded toward the argument. Therefore, a disposition to respond rationally is not necessary for being open-minded toward the argument.

---

[8] A similar charge is leveled by Dennis Whitcomb and Heather Battaly et al. (2015, 4) against Peter Samuelson and Ian Church et al.'s (2013, 65) account of intellectual humility. I share Whitcomb and Battaly et al.'s concern. I share as well their commitment to distinguishing intellectual humility from the closely related virtue of open-mindedness, for reasons discussed below.

[9] Most virtue-theoretic notions admit of this advising possibility. You can modify my behavior by informing me about which virtues I am not exemplifying, and I can deliberate about my actions by reflecting on what the various virtues require of me. For example, I might wrongly take myself to be a modest person and wonder why you seem to be put off by my well-intentioned reports of my professional success. When you have a frank discussion with me about how my reports are coming off, it can inspire self-reflection: "*Am* I modest enough? I thought I had to work on other character flaws, but maybe it's *modesty* I'm lacking."

What is required for open-mindedness is to not refuse to, in some sense, let the argument in, or be vulnerable to the argument. Riggs says, "When an open-minded person encounters a challenge to one of his beliefs, he responds by (at least sometimes) *taking such a challenge seriously*" (2010, 180 emphasis added). Jason Baehr uses similar language: in being open-minded, you're supposed to transcend a default standpoint in order to take up or "take seriously" the merits of a distinct standpoint (2011, 152).[10] This idea of taking an argument seriously is neglected by Open-Mindedness-as-Rational-Response, but needs fleshing out. What is it to take an argument seriously?

The discussion of Open-Mindedness-as-Rational-Response points toward the answer. Whether you're being closed-minded in committing, say, the base rate fallacy seems to depend on why you're doing it. If you're committing the base rate fallacy because you were absent that day from your statistics class, that's consistent with open-mindedness. If you're committing the base rate fallacy because you have a self-serving need to protect your pre-existing beliefs, that seems closed-minded. You can open-mindedly commit fallacies, but not if you are committing those fallacies because you are prevented from recognizing genuine validity by various non-cognitive and vicious influences on beliefs—like prejudice, bias, desire, wish, or fear.

---

[10] See also (Kwong 2016a): "open-mindedness amounts to the idea that we should take a distinct novel standpoint seriously" (410).

For convenience, I'll call them "affective factors." This label is not ideal. Your cognitive life can rightly include emotional and other influences that might normally be classified as affective. "The heart has its reasons" (1999, 146), Catherine Elgin points out, arguing that "the understanding we achieve is not indifferent to emotion but that understanding is none the less objective for that" (147).[11] Still, the kinds of factors I'm calling "affective" do not have a proper role in your cognitive life, since they push you to unjustly protect yourself and unjustly discount others.

On this construal, you take an argument seriously only if no affective factors prevent you from being persuaded by the argument. This follows Nomy Arpaly's suggestion that the person morally praiseworthy for open-mindedness

> is the person whose moral concern insulates her from the pull of other concerns that would otherwise render her unresponsive to evidence, in contexts in which something morally significant might be at stake. (2011, 81)

On this account, the injunction to take an argument seriously is largely negative. It tells you what not to do—not to be influenced by affective factors in your response to an argument. Is this enough? Or when I ask you to take an argument seriously am I asking for something positive as well? For example, in addition to making honest reasoning mistakes and to allowing bias, desire, and prejudice to infect your responses to arguments, you can also "mismanage" (in Elizabeth Harman's (2011, 445) term) your beliefs by violating various procedural norms: not taking notes, not investigating carefully, listening only to news sources with a single political agenda, etc.[12] Harman attributes this conception of how we can go wrong in our reasoning to William FitzPatrick (2008), but strains of it can be found in William Clifford (1886), as well. Clifford condemns those who are “careless” (345) about their beliefs and condemns as well the believer who “purposely avoids the reading of books and the company of men that call into question or discuss” (346) some belief.

---

[11] Thanks to Terry Horgan for referring me to Elgin's book. Horgan suggests (in conversation) that one kind of factor that might naturally be called "affective," but still might have a justified role in our cognitive lives is the "fishy smell" that some arguments present to us. Whether this smell is properly called "affective" I don't include it under my technical use of the term, because it is not necessarily self-protecting or prejudicial in the right sort of way.

It seems possible to violate procedural norms without being led to do so by the interference of affective factors. But if you violate procedural norms by ignoring an argument—even if you are not led to do so by the interference of affective factors—we may want to say that you shouldn’t count as having taken the argument seriously or as having been open-minded. In that case, a complete account of open-mindedness can’t stop with the elimination of the influences of affective factors. Genuine open-mindedness toward an argument also must not involve the violation of any procedural norms.

There are ways of violating procedural norms that may seem consistent with open-mindedness. For example, suppose there can be reasonable disagreements about what the procedural norms are—about, say, what sample size is required to mount a legitimate statistical argument. If in addition, there is a fact of the matter about what the correct procedural norm is, then the incorrect side of a reasonable disagreement about that norm will often end up in violation of that norm. This shouldn’t prevent the losing side from being open-minded toward arguments when their response to those arguments involves the relevant violation.[13]

---

[12] Stipulate for now that included among the procedural norms are norms requiring the presence of various second-order attitudes about your beliefs of the sort endorsed by Riggs and Jonathan Adler: that you are fallible and, thus, any of your beliefs are possibly wrong. Even so, satisfaction of these norms is not enough for open-mindedness, as I argue below.

[13] Thanks to Al Casullo for suggesting this way of open-mindedly violating procedural norms.

There are also ways of unreasonably violating procedural norms while failing to be closed-minded. If you choose not to listen to an argument about the consequences of raising the minimum wage (so, you violate a procedural norm that requires you to attend to relevant arguments) because you are protective of your personal belief about the minimum wage, then you are closed-minded toward the argument. But if you choose not to listen because you are depressed, in a hurry, tired, concentrating on other things, or have social anxiety disorder, you are not closed-minded toward the argument even though you do not listen to it, and so violate a procedural norm.

In this case, though you are not closed-minded, you're not quite open-minded, either. You might be prevented from being open-minded, not by manifestation of the vice of closed-mindedness, but by other characteristics that prevent you from fully realizing an open-minded response. Closed-mindedness and open-mindedness, then, though mutually exclusive, are not exhaustive. You are open-minded toward an argument when both i) affective factors do not dispose you against being persuaded by an argument and ii) you are not disposed to unreasonably violate any procedural norms in your response to the argument. You are closed-minded toward an argument when either i) affective factors dispose you against being persuaded by the argument or ii) affective factors dispose you to violate some procedural norms in your response to the argument. These accounts are an improvement over Open-Mindedness-as-Rational Response:

> (Open-Mindedness-as-Unblocked-Belief) You are open-minded toward an argument iff both i) affective factors do not dispose[14] you against being persuaded by the argument and ii) you are not disposed to unreasonably violate any procedural norms in your response to the argument.

---

[14] The conditions are put in terms of dispositions because, prior to engagement with an argument, affective factors may not have had the opportunity to prevent you from being persuaded and you may not have had the opportunity to violate any procedural norms. But you still might fail to be open-minded toward the argument. Thanks to Jon Kvanvig for raising this issue.

> (Closed-Mindedness-as-Blocked-Belief) You are closed-minded toward an argument iff either i) affective factors dispose you against being persuaded by the argument or ii) affective factors dispose you to violate some procedural norms in response to the argument.

I think that Open-Mindedness-as-Unblocked-Belief (or something like it) expresses necessary conditions on open-mindedness and Closed-Mindedness-as-Blocked-Belief (or something like it) expresses sufficient conditions for closed-mindedness. But you can fail to be open-minded and you can be closed-minded even if unblocked. There are other ways to closed-mindedly fail to be moved by an argument.

Your failure to be moved by an argument can be the product of your fundamental worldview, epistemic standards, framework judgments, prior probability assignments, or the like. Your response may[15] be irrational. But whether you're rational or not, there's no guarantee that your resistance is the product either of affective factors or the violation of procedural norms.[16] Even so, in some such cases, you are closed-minded.[17] Suppose you and I are arguing about whether Oswald shot JFK. You offer this familiar argument:[18]

- According to the official story, and in fact, it takes 2.3 seconds to reload and shoot with Oswald's gun.
- Therefore, it takes 6.9 seconds to fire three shots with Oswald's gun: 3×2.3=6.9.
- Therefore, the official story that Oswald shot three times in 5.6 seconds is false.

---

[15] It also may not. Perhaps it's rational, as Jennifer Faust (2008, 74) claims, to maintain views that are products of framework judgments, fundamental worldviews, or deep epistemic standards. See also (Schoenfield 2014, 199).

[16] Clifford (1886) may think otherwise. In his cases of clear sin in the matter of belief, the believers err in "stifling . . . doubts" (340) and "listening to the voice of prejudice and passion" (341). The conclusion Clifford draws from these cases is that your belief must be proportioned to the evidence, as if the only way that belief can fail to be proportioned to the evidence is if it is moved by self-serving stifling of doubts. I am convinced by Harman's (2011, 455) criticisms of this sort of view (her targets are (FitzPatrick 2008) and (Moody-Adams 1994)).

[17] Riggs (2010, 186) is sympathetic to the view that you are open-minded in such cases.

[18] See, for example, (Jones 2015): "Three shots at 2.3 seconds each resulted in a minimum time frame of 6.9 seconds. And yet the Warren Report claimed Oswald fired . . . three shots total in 5.6 seconds" (24).

I offer one standard response to your argument:[19]

- It doesn't take 2.3 seconds to fire the first shot; the clock starts at the first shot.
- Therefore, it takes only 4.6 seconds to fire three shots with Oswald's gun.
- Therefore, Oswald could shoot three times within the 5.6-second limit.

This standard response is a decisive argument; it proves that your argument is a failure. Nonetheless, after you hear me out and understand the standard response you are not at all persuaded; you continue to maintain with full confidence, on the basis of your argument, that the official story that Oswald shot three times in 5.6 seconds is false. Supposing that you make this error in reasoning without violating any procedural norms or being motivated by affective factors, are you closed-minded toward the standard response?

It depends. As already discussed, while it's possible for you to remain unmoved by the standard response, and though it's arguably irrational for you to remain unmoved, your vice need not be closed-mindedness. If you think that the first premise of the response is false—you think that the sound of the gun starts the clock and then it takes 2.3 additional seconds for the first bullet to exit the gun before Oswald starts the 2.3 second reloading process—then you may be exhibiting some other vice. You're being irrational, if you understand the premise and fail to find it plausible. Or you're being dense if you fail to understand the premise. In neither case need you be closed-minded.

But what if you are unwilling to accept the standard response even on the condition that you grasp the premise, you find it plausible and compelling, you see how each subsequent step in the argument follows from the steps prior to it, and you admit that you can't identify a flaw in the argument? In this case, your vice does not seem to merely be stupidity or irrationality. It seems to be a particular sort of irrationality, one of a refusal to allow counterarguments—no matter how convincing they seem—to affect belief. It seems, that is, like closed-mindedness. Suppose you say,

---

[19] See, for example, (Myers 2007): "FBI firearm experts concluded that the minimum firing time between shots . . . was 2.3 seconds—that's a total minimum firing time for three shots of *4.6 seconds* (don't forget, the clock starts running with the firing of the first shot)!"

> Oh, I see what you mean. It does look really plausible that it's illegitimate to multiply the time to fire a shot by 3, because it sure looks like Oswald won't have to wait 2.3 seconds to fire that first shot. I don't have an answer to that. But I'm completely unconvinced. There's got to be something wrong with your argument, even though I have no idea what it is. Oswald simply can't get three shots off in 5.6 seconds when each shot takes 2.3 seconds. Rest assured, though. I am completely open-minded toward your argument.

Here, we can stipulate, you have followed all the procedural norms. You've thought the argument through. You recognize that you're generally fallible and that you could even be wrong about your views on the Kennedy assassination. You aren't motivated by a desire to have been right all along. You're genuinely modest; you aren't trying to protect yourself or your beliefs. You haven't allowed wishful thinking or prejudice to infect your reasoning. You are unblocked. But your conviction that something-you-know-not-what must be wrong with the standard response makes you unwilling[20] to be persuaded even though you find every step in the argument compelling and have no idea what's wrong with the argument.

This means that you are closed-minded toward the argument. Your final insistence that you are open-minded toward the argument is bizarre, given your immediately prior admission that, while you have no clue what's wrong with the argument and every step seems compelling, you're unwilling to be persuaded in the slightest. If I had asked you to stay open-minded and you sincerely promised that you would, then in responding to me as you do, you have broken your promise. I can at least expect from a sincere commitment to be open-minded that you will be willing to pull back on your confidence if I offer an argument each step of which you find compelling and in which you cannot, by your own admission, find a flaw.

---

[20] You are expressing an unwillingness to be convinced, not merely an observation that you are not able to be convinced. You're not confessing that, despite your best efforts at reducing your confidence, you can't help but fail. You're not condemning your disposition to remain unmoved. Even so, your unwillingness is compatible with your satisfaction of all procedural norms and your invulnerability to affective factors (just as your unwillingness to be persuaded by trick proofs that 1=0 need not be the result of the influence of affective factors or the violation of procedural norms).

This is even clearer when, prior to hearing me out, you are unwilling to be moved conditional on hearing me out and being unable to locate a flaw. If I tell you that I'm going to present an argument against your view, and you promise to be open-minded, I'll not expect you, in the next breath, to say,

> Yes, I promise I am, and will remain completely open-minded toward your argument. But even if every step in your argument is very compelling to me and I can't figure out where your argument goes wrong, I'm not going to be persuaded in the slightest. Trust me, this isn't because I'm going to violate any procedural norms; I'll be very conscientious and think carefully about your argument. Nor is it because I'm going to be influenced by affective factors. I have no particular stake in being right about this. It's just that, in this case, I believe that even arguments that look completely sound must have a tricky flaw, so I simply won't budge if your argument looks completely sound. Nonetheless, yes, I am completely open-minded toward your argument.

There is a tension between your claim to be open-minded toward my argument and your unwillingness to be even slightly persuaded by my argument conditional on finding each step compelling and being unable to expose a flaw.[21] What is needed in the accounts of open-mindedness and closed-mindedness is a clause that explains this tension. I call the accounts I endorse "Platonic" conceptions[22] because they make essential to open-mindedness the respect for pure argument that is expressed in the epigraph opening this chapter and throughout Plato's works.[23] If you can't refute an argument, you must accede to it:

---

[21] This is in keeping with Jack Kwong's (2016b, 82) contention that, upon failing to refute a counterargument, open-mindedness can demand revision of your position.

[22] These Platonic conceptions differ from what Hare calls "Socratic Open-Mindedness," since Socratic Open-Mindedness requires only "a willingness to give consideration to further evidence that may yet emerge, recognizing that such further evidence may require revising previous conclusions that had seemed settled" (Hare 2009, 6). It does not require a willingness to be significantly persuaded conditional on finding the steps compelling and being unable to locate a flaw.

> (The Platonic Conception of Open-Mindedness) You are open-minded toward an argument iff i) affective factors do not dispose you against being persuaded by the argument, ii) you are not disposed to unreasonably violate any procedural norms in your response to the argument, and iii) you are willing to be significantly persuaded conditional on spending significant time with the argument, finding the steps compelling, and being unable to locate a flaw.

> (The Platonic Conception of Closed-Mindedness) You are closed-minded toward an argument iff either i) affective factors dispose you against being persuaded by the argument, ii) affective factors dispose you to violate some procedural norms in response to the argument, or iii) you are unwilling to be significantly persuaded by the argument conditional on spending significant time with the argument, finding the steps compelling, and being unable to locate a flaw.

Though I put the Platonic conceptions as stating necessary and sufficient conditions, I am not averse to additional conditions; there may be other ways to be closed-minded than via the stated disjuncts, and there may be other conditions you have to satisfy to count as open-minded. What is important for this book is that it is sufficient for counting as closed-minded and necessary for counting as open-minded that you satisfy the respective condition (iii). To be open-minded is not just to be unblocked. It's to be willing to be persuaded when a certain sort of argument comes along. Closed-mindedness need not be blocked belief. You can also be closed-minded when nothing blocks belief, but you are nonetheless unwilling to be persuaded when a certain sort of argument comes along.

---

[23] Here are some representative passages, all from (Plato 1997):

> . . . if you don't mind, we ought to follow where the argument leads us. (*Laws* 667a)
>
> . . . let's either refute our argument and show that we were wrong, or as long as it remains unrefuted, let's never say that the soul is destroyed by a fever or any other disease. (*Republic* 610b)
>
> . . . if my questioner was one of those clever and disputatious debaters, I would say to him: 'I have given my answer; if it is wrong, it is your job to refute it.' (*Meno* 75c)
>
> . . . we either refute Parmenides' claims or else agree to accept them. (*Sophist* 241e)
>
> . . .let someone refute us and persuade us that we've made a mistake—or else, so long as he can't do that, he should say just what we say. (*Sophist* 259a)
>
> If some mistake was made then, anyone now has the opportunity to take it up again and correct it. (*Philebus* 60d)

Just to make it explicit, note that according to the Platonic conception, open-mindedness toward an argument requires only that you be conditionally willing to be *significantly* persuaded. Therefore, it doesn't require that that you be conditionally willing to form a full belief in the conclusion of the argument; it requires only that you be (*ceteris paribus*[24]) conditionally willing to adjust your confidence to some significant extent. Because my primary conclusion in the book is that often you shouldn't be willing to adjust your confidence at all, it follows that you should be closed-minded toward certain counterarguments no matter how minimally we spell out "significant." Therefore, I leave open the precise amount by which you must be willing to adjust your confidence.

What is it to be "willing" to be persuaded by an argument? I remain neutral between various conceptions of willingness and unwillingness. First, willingness could be a straightforward disposition to do the very thing you're willing to do. Unwillingness could be a disposition to refrain from doing the thing you're unwilling to do. More plausibly, willingness could be an endorsed disposition to do that thing, while unwillingness could be an endorsed disposition to refrain from doing it. If you're disposed to do something and you approve of your disposition, you're willing to do it. If you're disposed to refrain from doing something, and you approve of being disposed to refrain from doing it, then you're unwilling to do it.

---

[24] The degree to which you are willing to be persuaded has to do with your willingness to adjust your confidence. You are persuaded by an argument when the argument, *ceteris paribus*, moves you to be more confident in its conclusion than otherwise. The *ceteris paribus* clause is necessary because there are arguments for conclusions you already believe that should count as persuasive for you, even though they don't, when you hear them, change your confidence; you already believed the conclusion anyway.

On these construals, you can't be willing to be persuaded by arguments you can't be persuaded by, and you can't be unwilling to be persuaded by arguments you can't help but be persuaded by.[25] Therefore, you can't be open-minded toward arguments you can't be persuaded by, even if you sincerely try to be persuaded or endorse being persuaded. And you can't be closed-minded toward arguments that you can't help but be persuaded by, even if you sincerely try not to be persuaded or condemn being persuaded.

These may be the right results. I won't decide the matter. There are other construals of willingness and unwillingness that don't have these consequences. Willingness and unwillingness might amount to patterns of dispositions to try or strive by an act of will. For example, unwillingness to be persuaded could be a disposition to try not to be persuaded. Alternatively, willingness and unwillingness could be a pattern of endorsements and condemnations of various dispositions you may or may not have. For example, in being willing to be persuaded you might endorse the disposition to be persuaded.[26] This construal allows you to be open-minded toward arguments that you can't, but want to, be persuaded by.

These latter construals have the support of ordinary notions of willingness and unwillingness, which tell us that the mind can be willing though the flesh is weak. You can be willing to do what you are unable to do. You can be unwilling to do what you can't help but do. Any time you are forced to do something against your will—any time you do something unwillingly—you are unwilling to do something even though you can't help but do it. Any time you are unable to do something you try hard to do, you are willing to do something you can't do.

---

[25] This assumes that if you are disposed to do something in a set of circumstances, then you can do that thing when you're in those circumstances.

[26] To endorse or condemn a disposition, in the relevant sense, is to endorse or condemn your having the disposition (which, of course, you can do even if you do not have the disposition, the way you can endorse your winning the lottery even if you don't win it). You might endorse your being disposed to be persuaded by an argument and not endorse my being disposed to be persuaded (because, perhaps, you know that I have different evidence than you do).

These facts tell against construing willingness and unwillingness simply as (endorsed) dispositions to do (or refrain from doing) what you are willing (or unwilling) to do. Still, I remain neutral between these construals and the alternatives: willingness and unwillingness construed as dispositions to try or strive, or as patterns of endorsements and condemnations. Thus, I remain neutral on whether you can be open-minded toward arguments you can't be persuaded by.[27]

There are cases in which you are willing to reduce your confidence conditional on being unable to expose a flaw in a counterargument, but—miracle-of-miracles—you can *always* expose what, to your satisfaction, is a flaw. For example, suppose you believe that vaccinations cause autism and, dozens of arguments—both bad and good—to the contrary later, you haven't reduced your confidence at all. This is not because you are unwilling to reduce your confidence in response to those arguments conditional on being unable to expose a flaw in them. Rather, it's that no matter what arguments are presented, you are disposed to find flaws in them.[28] If so, you do not satisfy condition (iii) in the Platonic conception of closed-mindedness, and you do satisfy condition (iii) in the Platonic conception of open-mindedness. But, intuitively, you count as closed-minded and you do not count as open-minded.

There are a few different versions of this kind of case. In one version, you persistently find flaws that aren't there. In a second version, you persistently find flaws that are there, though you aren't nearly as adept at finding flaws in flawed arguments for your position. In a third version, you are perceptive enough to find flaws in all flawed arguments, whether they support your position or not. The final version is not one in which, intuitively, a proper account of closed-mindedness must count you as closed-minded. I consider only the first two versions.

---

[27] This neutrality is legitimate because, I argue in Chapter 6, you should be closed-minded toward some arguments regardless of how willingness is construed: you shouldn't have a disposition—endorsed or otherwise—to be persuaded by them (conditional on finding them apparently flawless). Nor should you try to be persuaded by them or endorse the disposition to be persuaded by them (conditional on finding them apparently flawless).

[28] Thanks to Marc Johansen for raising this concern.

One way to respond is to add a fourth condition to the accounts having to do with a persistent disposition to find flaws in counterarguments that have none. But adding this fourth condition is arguably unnecessary. The Platonic conceptions of open-mindedness and closed-mindedness are only committed to the incorrect result if you satisfy the other conditions on open-mindedness and fail to satisfy at least one of the other conditions on closed-mindedness. But, plausibly, in all the relevant cases, the right conditions are satisfied. If you are persistently able to expose flaws only in arguments in favor of vaccinations, then there must be some explanation for this. The explanation cannot simply be your essential perceptiveness, since somehow you aren't able to easily expose flaws in arguments for your side. The explanation has to make reference to some self-serving affective factors. Therefore, in all cases in which you count as closed-minded toward the relevant arguments, it looks like the Platonic conception yields the desired verdict.

If shrewd evaluation of the evidence is what moves you to expose flaws in—though only in—arguments against your beliefs, some won't have the intuition that you are closed-minded toward those arguments. If so, these cases aren't counterexamples to the Platonic account. Presumably, you only count as closed-minded toward those arguments if you are moved not by shrewd evaluation of the evidence, but by self-serving affective factors. But it's exactly in those cases that the Platonic account—rightly—counts you as closed-minded.[29]

Moving on: open-mindedness comes in degrees. If you are unwilling to hear an argument out, then you are less open-minded toward the argument than someone who is willing to hear the argument out, even if they are unwilling to think deeply about the argument. If you are willing to think deeply about the argument but unwilling to change your mind no matter what, then you are less open-minded than someone who is willing to change their mind if they can't figure out what's wrong with the argument. If you are willing to modify other kinds of attitudes and beliefs in response to the counterargument, but unwilling to change your mind about the conclusion of the counterargument, then you are more open-minded than if not, but less open-minded than if you're willing to change your mind about the conclusion of the counterargument, as well.[30] But unless you're willing to reduce your confidence upon spending time with the argument, finding each step compelling, and being unable to expose a flaw, you're not open-minded, full stop.[31]

---

[29] Thanks to Evangelian Collings for persuading me that these cases aren't as obviously putative counterexamples as I initially thought.

You sometimes should be open-minded in this sense; you should be willing to reduce your confidence if the steps seem individually compelling and you can't expose a flaw. But sometimes you shouldn't. When it comes to Sorites arguments, Zeno's arguments, and trick "proofs" that 1=0, you should retain full confidence in the denials of the arguments' conclusions. In other cases, there might be moral reasons to be unwilling to be persuaded by certain arguments.[32] Therefore, the Platonic conceptions allow that sometimes you should be closed-minded and shouldn't be open-minded. Open-mindedness can be bad and irrational. Closed-mindedness can be good and rational.

---

[30] Thanks again to Evangelian Collings. To use her example, in response to Swinburne's argument that homosexuality is a disability, you might be unwilling to increase your confidence in Swinburne's conclusion. But you might be willing to reduce your confidence in the claim that it is of utmost importance that one's children be biologically their own. If you are, then you count as more open-minded toward the argument than you would otherwise.

[31] You might also decrease in degree of closed-mindedness if you satisfy fewer of the disjuncts of the Platonic conception. Students in an introductory ethics course might start out closed-minded toward the arguments of those they disagree with because they are unwilling to be persuaded by arguments and are moved by affective factors to violate procedural norms. By the end of the course they might stop being moved to violate procedural norms, but still be unwilling to be persuaded by arguments. They still count as closed-minded, though not as closed-minded as they were before. Thanks to Asia Ferrin for this example.

[32] Thanks to Josh Stein for suggesting I mention this possibility at this point. I discuss this possibility briefly in Chapter 6.

Is this a problem with the Platonic conceptions? Is “open-mindedness” necessarily a term of praise while “closed-mindedness” is necessarily a term of condemnation? Perhaps the Platonic conceptions are better thought of as accounts of related attitudes, like flexibility or dismissiveness. When you are wrongly flexible, you are not open-minded; you are gullible. When you are rightly dismissive, you are not closed-minded; you are steadfast.[33]

I agree that when you are wrongly flexible, you are gullible. But the gullible are still open-minded. Open-mindedness and closed-mindedness are not like murder. Murder is wrongful killing. So if I find out that a killing I previously thought was wrongful was actually not—if it was performed in justified self-defense, for example—I retract my claim that the killing was murder. I don’t exhibit the same patterns of retraction with respect to claims of open-mindedness and closed-minded. If I know that you are willing to reduce your confidence in response to some argument I think you should be convinced by, then I’ll think you’re open-minded toward that argument. If I find out that the argument you’re willing to be convinced by is not the one I thought you should be convinced by, but rather an argument that there is no motion, I won’t say, “I guess you aren’t open-minded toward the argument after all.” Finding out that your willingness to be convinced is irrational doesn’t change my mind about whether you are open-minded. It changes my mind about whether you should be.

Similarly, when you are rightly dismissive, you are not open-minded; you are closed-minded. If I overhear you say, “I don’t care how good your argument seems or that I can’t figure out what’s wrong with it; I’m not changing my mind in the slightest,” I can render my verdict that you’re closed-minded toward whatever argument you’re responding to without knowing what argument it is. My verdict doesn’t depend on whether the argument you’re dismissing is, unbeknownst to me, an argument that human behavior is a primary cause of global warming or an argument that homosexuality is morally perverse. In the first case, your dismissal is irrational. In the second it's rational. In neither case are you open-minded toward the argument. You're closed-minded in both.[34]

---

[33] Thanks to Leo Iacono for suggesting “gullibility” and “steadfastness” as more natural labels for what I am calling irrational “open-mindedness” and rational “closed-mindedness,” respectively.

Open-mindedness is often thought of as essentially good because satisfying the conditions of the Platonic conception is often thought of—incorrectly—as essentially good. There seems to be something essentially good about being willing to reduce your confidence in response to arguments you can't expose flaws in and each step of which you find compelling. There seems to be something essentially bad about being unwilling to reduce your confidence in response to arguments you can't expose flaws in and each step of which you find compelling. These normative intuitions are false. But it takes argument to show that they're false. The natural tendency to take these normative intuitions for granted explains the natural tendency to call open-mindedness good and closed-mindedness bad. This explanation is only available if the Platonic conceptions are correct.

1.3 OPEN-MINDEDNESS AND OUTRIGHT BELIEF

I spend much of the rest of this chapter comparing the Platonic conception of open-mindedness to a competing account presented by Jonathan Adler and developed by Riggs. According to Adler and Riggs, at the core of open-mindedness is a recognition of "myself as generally fallible" (Adler 2004a, 130)—an "acknowledgement that, being human, he could always have got things wrong" (Riggs 2010, 181). If I present you with a counterargument and your response reveals an admission that you might be mistaken, that is intuitively an open-minded response. Their account explains why:

---

[34] Again, contrast with murder; if all I know is that you caused someone's death, I don't judge that it's murder until I know more details about the killing. Whether an act is murder depends on whether it's wrongful. Whether a response is closed-minded doesn't.

> (Adler/Riggs Open-Mindedness) You are open-minded toward a challenge to your belief only if you recognize that your beliefs—among which is the challenged belief—may be mistaken.

One appealing feature of this account, they argue, is that it explains how open-mindedness is compatible with full—or outright[35]—belief. Not only can we be more or less confident that some proposition is true, we can also believe outright that the proposition is true. On one view, outright belief is the attitude you have when you are willing not just to say that some proposition is very probable or almost sure, but that it is true full stop.[36] If that's right—if flat assertion is the expression of outright belief—then there is no easy mapping between outright belief and any confidence level less than absolute certainty, especially if confidence-level is a matter of probability judgment. That's because, for any confidence-level less than absolute certainty[37], it seems consistent to say, "I'm not saying it *is* true. I'm really confident it's true, yes. But I don't outright believe it."[38]

---

[35] Adler and Riggs most often talk of "full" belief, though Adler sometimes uses "all-out" or "all-or-nothing" instead of "full," while Riggs occasionally speaks of "full-blooded" beliefs. I prefer the term "outright belief," because "full" seems to connote that no further increases in confidence are possible, whereas it should initially be an open question whether full belief allows for some uncertainty. Still, to respect Adler's and Riggs' discussions, I use "full belief" interchangeably with "outright belief."

[36] This is a simplification of Mark Kaplan's (1996, 109 ff.) Assertion View of outright belief. I don't mean to commit to the Assertion View, but rather use it to highlight the attitude of outright belief I'm referring to. As Ram Neta puts the point:

> Sometimes, stating one's account of some property can help others to identify which particular property one is talking about, whether or not those other people end up agreeing with one's account of that very property. (2007, 198)

[37] The view that outright belief requires certainty is in the minority, though see (Clarke 2013) for a defense. Other views include "Lockean" views according to which to outright believe something is to have a degree of confidence above some threshold less than certainty (see (Foley 2009, 37), (Sturgeon 2008, esp. 147 ff.), and perhaps (Locke 1693/1995, Book IV, Ch. XV, §3)) and pragmatist views, according to which to outright believe something is to be prepared to act on it (see (Fantl and McGrath 2009, Ch. 5), (Ganson 2008, 444), (Weatherson 2005, 421 ff.), (James 1962, 242), and (Clifford 1886, 342)). The Platonic conception allows outright belief to be compatible with open-mindedness, on all of these views (see note 43).

[38] Thanks to Christina Stiso for encouraging me to say more about the notion of outright belief than I initially had.

Is outright belief compatible with open-mindedness? According to Adler and Riggs (and me), it had better be: the injunction to be open-minded had better not require that we give up our full beliefs. But Adler and Riggs want more. Riggs, for example, wants it to be the case that the open-minded person takes challenges to their full beliefs seriously:

> When an open-minded person encounters a challenge to one of his beliefs, he responds by (at least sometimes) taking such a challenge seriously. This can be so even if he believes the challenged belief quite strongly. (Riggs 2010, 181)

On the assumption that Riggs doesn't think that the open-minded person is irrational in taking challenges to what they fully believe seriously, then Riggs seems to want a proper account of open-mindedness to require that, even when you *rationally* maintain full belief, you (sometimes) *should* take challenges seriously.

Adler suggests a similar desideratum. He says:

> I hold both that your view should be tolerated, not suppressed or censored, despite its definite falsity, and *that I may learn from arguing with you* over it . . . I have reason to engage in critical discussion with you *as an expression of my open-mindedness* in regard to a belief of mine. (Adler 2004a, 133 emphasis added)

Adler and Riggs, then, both seem to want their account of open-mindedness to satisfy two desiderata:

1) It is possible to open-mindedly maintain full belief.
2) Open-mindedness sometimes requires, even when you rationally maintain full belief, that you take challenges seriously.

To clear up a possible ambiguity, according to the second desideratum, it's not merely a condition on being open-minded that you take challenges seriously, the way that selfishness, in some sense, requires that you put yourself before others. The second desideratum says that sometimes because you should be open-minded, you should take challenges to what you fully believe seriously.

Adler and Riggs think their account of open-mindedness satisfies both desiderata. I argue A) their account does not satisfy both desiderata because B) the desiderata cannot be jointly satisfied. Furthermore, C) the second desideratum is both illegitimate and false (of standard situations), and D) their account offers neither necessary nor sufficient conditions on open-mindedness in any case.

Adler offers an example that he hopes explains why acknowledging the general possibility of error can generate a requirement to seriously consider challenges even to what you fully believe. In his example, a widget company has a policy of randomly selecting one out of every batch of ten widgets for quality-testing. Because the incidence of imperfection is so low, a widget monitor can fully believe that the chosen widget—say, widget 30—has no defects or imperfections. But because of the general acknowledgement of fallibility, the company (rationally) implements the widget-checking policy anyway. From the monitor's point of view—prior to inspection—widget 30 is defect-free. But the monitor should check the widget anyway, because there is a policy requiring that he do so.

So far, so good. Arguably, the monitor should check the widget even if it is *certain* for the monitor that the widget is defect-free, if there is a general company policy that the widget be checked. This may be Adler's story about how the monitor's full belief in the widget's defect-free status can be rational even though the monitor should check the widget. As Adler tells it, the widget monitor's "actual basis for the examination"—what runs through the monitor's mind—is that

> [t]he process by which my company manufactures widgets is a reliable one, and so each widget is justifiably certified as error-free. Nevertheless, probably some few that we judge as lacking in defects or imperfections by our manufacturing process do (or will) suffer imperfections, which we care to correct. So we adopt a policy to randomly check one out of every ten widgets, and it so happens that widget 30 falls under this policy. (2004a, 132)

This is a little misleading. If the widget monitor is moved simply by the fact that the company has a policy mandating random widget inspection, then there is no difficulty reconciling the widget monitor's behavior with the full belief that widget 30 is defect-free. "It's my job," the widget monitor might say, "So I'll do my job." But the widget monitor in Adler's story also takes care to justify the company policy by noting the general possibility for any particular widget to be flawed. This creates a tension when the widget monitor is forced to attend to the possibility that widget 30 is flawed. For how should we continue the widget monitor's train of thought? If the widget monitor fully believes that widget 30 is defect-free, then the continuation has to go like this:

> . . . it so happens that widget 30 falls under this policy. That's too bad, because widget 30 is just fine. So now I have to take time out of my day to test a perfectly good widget. Of course, some widgets aren't just fine, but widget 30 is. So it's a bit silly I'm testing this one, but since the company has a policy, I guess that's what I have to do.

If this is how the monitor finishes the story, then it is not open-mindedness—the general acknowledgment of fallibility—that requires that the monitor check widget 30; it's the fact that the company has a policy and that, as an employee of the company, the monitor should obey the company's policy. Of course, the policy might be a good one because of the truth of the general acknowledgement. But the application of the policy to this case—from the monitor's point of view—is a waste of time, because, as the monitor notes, the policy in this case requires inspecting a widget that is defect-free. Generally speaking, unless there is some mitigating factor (which there is in this case—that the widget monitor is an employee), we shouldn't follow even generally good policies if in a particular case, following that policy is a waste of time. Therefore, if this is what's going through the monitor's mind, what makes it that the monitor should test widget 30 is not the general acknowledgement of fallibility, but simply the fact that the monitor's employer has a policy requiring that widget 30 be tested.[39]

Perhaps Adler thinks that this isn't what should be going through the mind of the widget monitor immediately prior to testing the widget. Maybe what Adler would expect is more like this:

> . . . it so happens that widget 30 falls under this policy. That's fine, because widget 30—though it is defect-free—is just a widget like any of the others. So it stands just as much a chance—namely, some—as any other widget of having a defect. So I'll go ahead and check it just in case it has a defect.

If this is what's going through the widget monitor's mind, then the general acknowledgement of fallibility is playing a pivotal role in making it the case that the widget monitor should check widget 30. Therefore, this depiction of the widget monitor's reasoning makes sense, on Adler and Riggs' interpretation of open-mindedness, of how open-mindedness can require that we take challenges seriously.

What's less clear is that the widget monitor, in this case, counts throughout the reasoning as having a full belief. In this depiction of the widget monitor's thinking, the widget monitor vacillates. The widget monitor at one point says—outright—that widget 30 is defect-free. In another breath the widget monitor expresses an intention to go ahead and check it just in case it has a defect: "Widget 30 is defect free, but I'd better check it just in case." If there is no vacillation, inserting clauses that make clear that there is no vacillation should result in a coherent assertion. But it doesn't. Once it's clear that the widget monitor isn't vacillating, the utterance becomes absurd:

> Widget 30 is defect free, but though it is defect-free, it might have a defect (that it doesn't have), and if it does (which it doesn't) and I don't check it, a defective widget will get through (which it won't). It's defect-free and I'd better check whether it's defect-free.

---

[39] This is further confirmed by the fact that the monitor would have run through the same reasoning even if it were certain for the monitor that widget 30 were defect-free.

To be clear: it's coherent to fully believe something while recognizing the chance that what you fully believe is false. What's incoherent is fully believing something while allowing the chance that what you believe is false to bear on your behavior.[40]

It does no good to object that the widget monitor shouldn't be committing in the first place to the claim that widget 30 has no defect. This may, in fact, be what the widget monitor should do—withhold commitment to the claim that widget 30 has no defect. But it won't help Adler or Riggs shore up their point—that open-mindedness can require you to take challenges to what you *fully believe* seriously. Their point is that recognizing your general fallibility is consistent with you being able to take a stand—even on members of the set of beliefs you generally acknowledge you're fallible about. That's why Adler posits that one of the widget monitor's judgments is simply, "This widget has no defects or imperfection" (2004a, 132).

The widget example is just an analogy, of course. But analogous difficulties beset the issue of whether open-mindedness, on their construal, makes it that you should take seriously challenges to your full beliefs. Adler says that,

> In obvious ways, the reconciliation proposed above carries over . . . Since the open-mindedness is justified by my view of myself as fallible, it does not follow from my judging it worthwhile to engage in the exchange that I do not know, or should not be certain of, my belief (and so too that yours is erroneous). Tolerance for your position is a political, not epistemic, judgment of equality, community and liberty. (2004a, 133)

Again, Adler wants his view of open-mindedness to show how it's possible that rational full believers should engage seriously with challenges. One way to show how this could be is by using the observation Adler makes at the end of the passage—that "tolerance for your position is a political, not epistemic, judgment of equality, community and liberty." But Adler needs to say more than this, because political justifications of serious engagement having to do with equality, community, and liberty don't essentially involve open-mindedness. Adler wants the full believer to engage in serious critical discussion, not out of independent political considerations but as noted earlier, "as an expression of my open-mindedness" (132). On Adler's account, this means that the full believer is to engage in serious critical discussion as an expression of their general fallibilistic acknowledgement that any of their beliefs could be wrong.

---

[40] For further details on how this sort of argument would go, see (Fantl and McGrath 2012, 32–36). I exploit this point again in Chapter 6 to argue against a general obligation of open-minded engagement with relevant counterarguments.

Mirroring the treatment of the widget monitor, if Adler wants your open-mindedness to require that you engage seriously with the challenges of others, you should be reasoning like this:

> . . . So we adopt a policy to test our beliefs, and it so happens that belief 30 falls under this policy. That's just fine, because belief 30—though true—is just a belief like any other belief. So it stands just as much a chance—namely, some—as any other belief of being false. So I'll go ahead and seriously consider counterarguments just in case it's false.

Here, your open-mindedness requires taking challenges seriously. But, as we saw with the widgets, what is expressed in this reasoning is not full belief; it is vacillation. That it's vacillation is indicated by the fact that, when an attempt is made to eliminate the possibility of vacillation, the resulting expression is absurd:

> Belief 30 is true, but though it is true, it might be false (which it isn't), and if it is (which it isn't) and I don't check it, I'll believe a falsehood (which I won't). It's true and I'd better check whether it's true.

So, while full belief is consistent with a general acknowledgement that your beliefs can be wrong, that general acknowledgement—and, hence, on Adler's and Riggs' account, open-mindedness—does not allow satisfaction of their second desideratum: that open-mindedness can require you to take challenges to your full beliefs seriously. Therefore, Adler and Riggs can't salvage that desideratum. This, in fact, is all to the good. For while the first desideratum—that it be possible to maintain open-minded full belief—may be legitimate, the second desideratum is not. In Ch. 6 I argue that when you rationally maintain full belief, you shouldn't take challenges seriously just because you may generally be mistaken; so, the second putative desideratum is false. But even if it were true, it is not a legitimate *desideratum*.

A desideratum is a prior constraint on an account—a principle that states that, if an account has a particular implication, then it's a bad account. It might be that open-mindedness in fact sometimes requires you to take challenges to your full beliefs seriously, but whether you are required to take challenges to your full beliefs seriously is a substantive ethical question and should not be decided by stipulation. What is a legitimate desideratum is that it not be analytic that you shouldn't take challenges to what you rationally, fully believe seriously. And though I argue in Chapter 6 that you shouldn't take challenges to what you rationally, fully believe seriously, it's analytic neither that you shouldn't nor that you should.

Therefore, though Adler and Riggs think it speaks well of their account that it satisfies the second desideratum, it's fine that it doesn't. But their account is not a good account of open-mindedness in any case. The condition that you generally acknowledge the possibility that your beliefs are wrong is neither necessary nor sufficient for open-mindedness toward an argument. First, it is not sufficient.[41] If you fully believe that you could be wrong, but you nonetheless are unwilling to be persuaded by my argument even after you find all the steps in my argument compelling and are unable to find a flaw, you are not open-minded toward my argument; if you have promised me that you will be open-minded toward my argument, but you are unwilling to bend under those conditions, then you have broken your promise.

---

[41] Riggs doesn't say that the condition is sufficient. He also requires for open-mindedness that the second-order attitude manifest itself in certain open-minded behaviors. These are 1) a habit of being willing to accept claims "about one's cognitive weaknesses and strengths" (2010, 182) and 2) a habit of self-monitoring, especially when presented with challenges, so that you guard against the weaknesses identified in (1) (where these include tones of voice, body language, etc.). Adding these conditions does not result in a set of sufficient conditions for open-mindedness, for the same reasons adduced in the body of the text.

Second, the proposal is not necessary for open-mindedness. If you don’t acknowledge the possibility that any of your beliefs could be false, but you are still willing to engage with my argument and revise your opinion if you can’t figure out where the flaw is, if you are not disposed to and don’t violate any other procedural norms, if you’re not prevented by affective factors from being persuaded, then you are open-minded toward my argument. You may be irrational for failing to acknowledge the general possibility that any of your beliefs could be false, given how many beliefs you have. But it is not the sort of irrationality that constitutes a failure to be open-minded toward my argument, provided the other conditions on open-mindedness are satisfied. Therefore, the Adler/Riggs condition is neither necessary nor sufficient for open-mindedness, though satisfying the condition might make you, *ceteris paribus*, more open-minded than you would be otherwise.

Something like the Adler/Riggs condition might be a good candidate for an account of intellectual humility, rather than open-mindedness. Allan Hazlett, for example, construes intellectual humility as “a disposition not to adopt epistemically improper higher order epistemic attitudes, and to adopt (in the right way, in the right situations) epistemically proper higher-order epistemic attitudes” (2012, 220). If a healthy appreciation that your beliefs may be false is generally an epistemically proper higher order epistemic attitude, then Hazlett’s account naturally dovetails with the Adler/Riggs condition, as a condition on intellectual humility. But because the Adler/Riggs condition is not an appropriate condition on open-mindedness, if Hazlett is right about intellectual humility, intellectual humility should be distinguished from open-mindedness.

As I see the relationship between intellectual humility and open-mindedness, your intellectual humility might often explain why you are open-minded. And if intellectual humility is generally a virtue, this might help explain why open-mindedness is generally a virtue, insofar as reasons to be intellectually humble provide reasons to be open-minded. But intellectual humility is an attitude toward yourself, while open-mindedness—of the sort that interests me in this chapter—is an attitude toward arguments and, especially, the arguments of others.[42]

The Platonic conception of open-mindedness allows open-mindedness to be consistent with full belief. It is consistent with full belief that you are willing to revise that belief if you can't figure out what's wrong with a counterargument. Because you have a full belief, you may be fully confident that you will be able to figure out what's wrong with the counterargument. But you still might be willing to reduce your confidence if you can't figure out what's wrong with it, even if you now have full belief that you're right. You might even have full belief in a proposition even if you do not have full belief that you could figure out what's wrong with some argument to the contrary; there are, after all, arguments for falsehoods that seem really convincing. Knowing this, it might sometimes be irrational to be willing to reduce your confidence if you can't figure out what's wrong with the argument (as I argue in later chapters). But A) it might sometimes be rational and B) whether it's irrational or not, it's possible to be willing to reduce your confidence conditional on being unable to locate a flaw in a counterargument even if you have a full belief prior to hearing out a counterargument.[43]

---

[42] A further attitude in the vicinity both of intellectual humility and open-mindedness is tolerance. Tolerance, like open-mindedness, is not limited to attitudes you have toward yourself; you can be tolerant toward the arguments of others in the same way you can be open-minded toward the arguments of others. But tolerance, unlike open-mindedness, does not require satisfaction of the third conjunct: you can be tolerant toward the arguments of others even if you are unwilling to be moved by those arguments conditional on being unable to figure out what's wrong with them. Tolerance more plausibly requires some degree of satisfaction of the other two conjuncts in the Platonic conception of open-mindedness. If you are moved by fear of some argument to violate procedural norms, for example, then it is natural to say you fail to live up to the ideal of tolerance. Thanks to Asia Ferrin for discussion.

[43] This is also true on pragmatist views of belief according to which you fully believe something just in case you are prepared to act on it and on views of belief according to which to believe something is to be certain of it (see note 36). If you're certain that *p* or fully willing to act on *p*, you still might be willing to reduce your confidence conditional on being unable to expose a flaw in a counterargument. You might be willing to do this if you take an inability to expose a flaw in a counterargument as a reason to reduce confidence.

This puts the open-minded full believer into a somewhat odd position. Open-minded dispute ends up with the following basis: “I’m right. But I should engage with you because you might present an argument whose errors I can’t expose, and if you do, I’ll change my mind.” Why would you now be willing to change your mind when confronted with an argument that you now believe would be misleading? Answer: often, you shouldn’t. This is the argument I present in Chapter 6. But it’s one question whether, when you have a full belief, you should be open-minded toward counterarguments. It’s another question what open-mindedness is and whether open-mindedness is consistent with full belief.

1.4 CONCLUSION

There are four stages at which we can evaluate a closed-minded attitude toward an argument—four stages on the path from there being no salient counterarguments to there being a relevant counterargument you’ve spent time with, in which you find each step compelling, and in which you are unable to expose a flaw. At the first stage you believe that *p* and there is no salient counterargument against *p*. At the second stage, you know about the existence of the counterargument, but you haven’t heard it out yet or become familiar with any of its details. At the third stage, you’ve heard out the counterargument, but haven’t spent time investigating it and trying to figure out what’s wrong with it. At the fourth and final stage, you’ve heard out the counterargument, spent time trying to figure out what’s wrong with it, find each step compelling, and are unable to expose a flaw.

In the first three stages, open-mindedness requires that you’re willing to reduce your confidence on the condition that you reach the fourth stage. At the fourth stage, though, to be open-minded toward the argument, you must be willing to reduce your confidence conditional on arriving at the very stage that you've arrived at. As noted, I take no stand on whether willingness to do something entails a disposition to do that very thing. But if it does, once you've arrived at the fourth stage, to be open-minded toward the argument, you must reduce your confidence. Only at the fourth stage have you necessarily *exercised* either open-mindedness or closed-mindedness, either by manifesting your disposition to reduce your confidence upon reaching that fourth stage, or by manifesting your disposition not to.[44]

At all of these stages, in many cases, you should be open-minded, because in many cases your full belief is not justified. But in other cases, your full belief is justified and even counts as knowledge. In many of those cases, open-minded engagement with counterarguments should be avoided. This latter point is the conclusion of the main argument of Chapter 6. The first step in that argument is that knowledge can survive even in those cases in which you can't figure out what's wrong with arguments to the contrary—even in those cases in which you've reached the fourth stage and exercised closed-mindedness. It's to this topic I turn in Chapter 2.

---

[44] You can at earlier stages manifest your disposition to, once reaching the fourth stage, reduce your confidence in your position. You can do this, for example, by preparing for the possibility that you'll reduce your confidence if you do reach that fourth stage. But there are no steps that you must take, at earlier stages, to exhibit open-mindedness toward the argument. You can have a disposition to reduce your confidence upon reaching the fourth stage that manifests itself in no way in the first three stages.

2

A Defense of (a different kind of) Dogmatism

> He considered justly, that it was not requisite, in order to reject a fact of this nature, to be able accurately to disprove the testimony, and to trace its falsehood, through all the circumstances of knavery and credulity which produced it. (David Hume (1748/2007, 108))

2.1 KNOWLEDGE AND CLOSED-MINDEDNESS

You know that things move even if you don't know how Zeno's arguments fail. You know that 1 doesn't equal 0 even when stymied by a "proof" that 1=0.[1] Does the same hold for more controversial domains? Can knowledge survive confrontation with challenging counterarguments for conclusions about the theory of evolution, psychic phenomena, the efficacy and dangers of vaccinations, convoluted conspiracy theories, repugnant moral positions, the existence of God, and whether the Holocaust occurred?

In all these domains, there is some significant disagreement about their central propositions. It might seem you cannot know whether these controversial propositions are true unless you figure out how arguments on the other side go wrong. For example, Mill says,

> He who knows only his own side of the case, knows little of that. His reasons may be good, and no one may have been able to refute them. But if he is equally unable to refute the reasons on the opposite side; if he does not so much as know what they are, he has no ground for preferring either opinion. The rational position for him would be suspension of judgment, and unless he contents himself with that, he is either led by authority, or adopts, like the generality of the world, the side to which he feels most inclination. (Mill 1859/2003, 104)

According to Mill, you must suspend judgment on—and, hence, cannot know—propositions even for which you have good and unrefuted reasons, if you cannot refute the reasons on the other side and even if you don't know what those reasons are. More recently, with respect to controversial moral matters, Sarah McGrath argues that in a battle between two non-fallacious arguments for opposite sides, neither side can know that the conclusion of their argument is true:

---

[1] For a defense of a kind of closed-mindedness regarding arguments like Zeno's, see (Kelly 2011).

> Of course, the premises of my argument seem more compelling to me than the premises of Alice's argument; but by the same token, the premises of Alice's argument seem more compelling to her than the premises of my argument. . . . Can I break the symmetry, then, by assuring myself that the reasons that I have are more compelling than hers? This seems no better than simply privileging my judgment about a given shade of green over Alice's contrary judgment. (2008, 103)

Supposing that knowledge requires justified belief, one principled basis for the view that knowledge requires you to expose flaws in your interlocutor's counterargument is provided by so-called "dialectical conceptions of justification"—conceptions of justification according to which being justified in believing something requires adequately defending your view to actual or possible opponents. For example, Jan Willem Wieland's "Dialectical Assumption" holds that

> If S wants to find out the truth about a belief P, then S is justified in believing P only if S can defend P to a possible opponent in a non-question begging way. (2013, 199)[2]

Proponents of dialectical conceptions normally conjoin the conception with a substantive thesis about what it takes to adequately defend your view: being able to expose where your opponent's argument goes wrong.[3] Hence, Scott Aikin's "dialectical requirement for justification":

> If S is justified in holding that *p*, then S can successfully address the standing cases for not-*p* and against *p*. (2011, 17)

The intuition that knowledge cannot survive closed-minded dismissal of arguments in controversial domains is strong.[4] But closed-mindedness seems least legitimate when it comes from the opposition. It is illegitimate for you to ignore my argument or, possibly worse, hear me out and then dismiss my argument without finding a flaw. Suppose that the best you can do with my argument is tell me that, while all the steps seem pretty plausible, something just seems a bit fishy. I press you further: “Can you help me out? What exactly seems fishy?” You make some suggestions—this premise, that use of a concept—and, suppose, I answer them, to your satisfaction. But, you continue, “I guess I can’t put my finger on what’s wrong. But I just don’t buy it.” This response is wholly unsatisfying, especially regarding the controversial issues about which I feel strongest. And if you can’t do it to me, why would I get to do it to you? As Stephen Carter says,

---

[2] According to Markus Lammenranta (2011), being justified requires that we can "defend our beliefs for each other or a third party in a way that is dialectically effective" (12) and that we "have reasons that would convince" others (14).

[3] Wieland, for example, maintains that the key planks in the Dialectical Assumption require you to "provide reasons against possible challenges" (2013, 15). Lammenranta says that those who satisfy the requirements of the dialectical conception of justification can "respond to our challenges" (2011, 12).

[4] Lammenranta is sympathetic to skepticism in those "many areas where controversy prevails, such as politics, religion, and philosophy itself" (2011, 15).

> if I expect you to listen respectfully to me and give me a genuine opportunity to convert you to my way of thinking, civility requires me to first listen respectfully to you and give you a genuine opportunity to convert me to your way of thinking. (1999, 139)

That’s in theory. In practice I feel much more comfortable with my closed-mindedness than with yours. While I usually feel that, with a little time, I’d be able to figure out how your arguments go wrong or, at least, have some idea where the weak spots are, I wouldn’t be too concerned if I didn’t. There is a certain amount of guilt involved—in this as in all things—when I am reflective about my practices. But that doesn’t change the fact that the practice is there. To the extent that others are like me, all of us act as if closed-mindedness is often legitimate.

The good news offered in this chapter is that at least some of us don’t have to feel so guilty about our closed-mindedness. Even when we lack special training, some of us, I argue, know controversial propositions. For those of us who do, knowledge can often survive the manifestation of closed-mindedness. Briefly, the reason is this: for any controversial proposition you know, there could easily be an apparently flawless argument that the proposition is false. If so, your knowledge will sometimes not be destroyed if you are exposed to an apparently flawless counterargument and respond closed-mindedly. Therefore, you can retain knowledge even if you are exposed to an apparently flawless counterargument and respond closed-mindedly.

2.2 FORWARD-LOOKING DOGMATISM

Sometimes knowledge is destroyed by counterarguments because those counterarguments raise doubts in your mind; knowledge can be destroyed by psychological defeaters. But sometimes you willingly retain full confidence despite the counterarguments. When you willingly retain full confidence, can knowledge survive exposure to counterarguments in which you cannot expose a flaw?

Boring stuff first: terminology. You *dismiss* a counterargument when you willingly retain full confidence that the argument's conclusion is false despite being aware of the counterargument. You *closed-mindedly* dismiss a counterargument when it is irrelevant to your dismissal whether, after extended engagement, you continue to find the steps in the counterargument compelling and are unable to expose a flaw. When does knowledge survive closed-minded dismissal of a counterargument and when doesn't it? One simple answer is this: knowledge survives only when hardly anyone would be convinced by the argument. Hardly anyone is convinced by magic shows or Zeno's arguments. So knowledge can survive closed-minded dismissal in those cases. The same can't be said for many arguments about controversial philosophical, scientific, religious, and moral propositions.

This simple answer is incorrect. Suppose that you are a murderer and that you discover that someone else's fingerprints are on the murder weapon. This surprises you. You have no idea how those fingerprints came to be there. In this case, you retain knowledge that you committed the murder without having any idea what's wrong with the argument to the contrary. Knowledge survives closed-minded dismissal of the argument even though lots of people would be convinced by the argument; they would think you're not guilty.[5]

The difficulty with this case is that those who would be convinced of your innocence do not have all the evidence you do that you are guilty. Were they aware of your evidence—your memory of committing the murder—they would not be convinced by the argument that you're innocent. This can only go so far, of course. Let enough evidence come in that you're innocent and even you could be convinced to disregard your apparent memories. Nonetheless, the mere fact that lots of people would be convinced by an argument isn't sufficient for your knowledge to be destroyed by that argument. You might have evidence they lack.

A more promising line is that closed-minded dismissal destroys knowledge if lots of people would be convinced by the argument, and would remain convinced if they had your evidence to the contrary.[6] Call such an argument a "relevant counterargument." How many people must a counterargument be convincing to in order to be a relevant counterargument; how many people is a "lot"? More than one and allowably fewer than all. Fortunately, being more precise isn't necessary. That's because the domains of interest are paradigmatically controversial: the number of people on each side isn't in the gray area of "lots." There are lots of people on each side who would be convinced by similar arguments, so I'll grant that, for every side on these issues, there are relevant counterarguments. Nonetheless, knowledge often survives closed-minded dismissal of these arguments.

---

[5] An example from the television show *Seinfeld*:

> Jerry: I don't get it. He assigns it to you. You don't do it. Somehow it gets done, and now he's telling you what a great job you did.
>
> George: Maybe somebody did it and didn't take credit for it. Maybe it was already done and didn't need doing in the first place. I have no idea who did it, what they did, or how they did it so well. And you know what? Jimmy crack corn and I don't care.

[6] Not included in the evidence that must be shared are any feelings of sheer obviousness regarding various premises and conclusions. It is plausible enough that if lots of people are convinced by an argument and would remain convinced if they were aware of all your arguments then your knowledge can't survive closed-minded dismissal of their argument even if they don't share your feelings of sheer obviousness. I argue in Chapter 4 that feelings of obviousness can generate knowledge and, when they do, the resultant knowledge can survive closed-minded dismissal of relevant counterarguments.

The view that knowledge can survive closed-minded dismissal is a kind of dogmatism,[7] though not the kind defended by Jim Pryor (2000), according to which you can know that a proposition is true even if your belief is not justified by any non-question-begging reasons. Pryor's dogmatism is *backward-looking* in the sense that it comments on the sources of belief. Backward-looking dogmatism is suggested in Moore's proof of an external world, because Moore takes himself to know the premises in his proof (e.g., that here is a hand) though he cannot prove those premises. But Pryor notes Moore's additional "idea":

> that the proposition that there is a hand, though only defeasibly justified, is more certain than any of the premises which might be used in a skeptical argument that he does not know that there is a hand. (2000, 518)

The dogmatism suggested by this second idea is *forward-looking* in the sense that it comments on what might require belief revision. Of the two, forward-looking dogmatism seems closer to the ordinary use of the term; a dogmatic person isn't sensitive to counterevidence.[8] Still, the term is already taken, so I'll refer to forward-looking dogmatism as "dogmatism^{FL}." Dogmatism^{FL} does not entail that knowledge can survive closed-minded dismissal of all relevant counterarguments. But knowledge can survive closed-minded dismissal of relevant counterarguments that you lack sufficient expertise to reliably evaluate. I make the case for this in Chapter 3.

---

[7] In the psychological literature and in ordinary use, dogmatism is a character trait or a disposition. A dogmatic person, it might be said, exhibits their profound dogmatism. I use the term, following philosophical tradition, to refer to a *view*—a view that entails that knowledge can survive closed-minded dismissal of relevant counterarguments.

[8] This is reflected in the psychological literature. Milton Rokeach claims that dogmatism, along with the similar construct of rigidity, "refer to forms of resistance to change, but dogmatism is conceived to represent a relatively more intellectualized and abstract form than rigidity" (1954, 196). See also (Altemeyer 2002), (Shearman and Levine 2006), and (Martin, Staggers, et al. 2011), all of whom characterize dogmatism in large part in terms of flexibility and resistance to external pressure.

Arguments can be closed-mindedly dismissed subject to varying levels of scrutiny. You can dismiss arguments without having heard them out, having heard them out but without understanding them, having understood them but without taking the time to evaluate them, and having taken the time to evaluate them but without exposing flaws in them. Perhaps it's easier for knowledge to survive closed-minded dismissal if you haven't taken the time to understand or critique them than it is for knowledge to survive closed-minded dismissal of arguments you have taken time with but still can't diagnose. If you haven't spent time with an argument, there's a much greater chance that engagement would reveal a flaw than if you have spent time with the argument and still can't figure out what's wrong with it.[9]

Often, when you're evaluating a counterargument, you are pretty sure that you've spotted the weak part. Or, if you're not pretty sure, you have hunches about which steps or premises are problematic. Even Moore, who didn't much care where Russell's argument for skepticism broke down, thought that he'd isolated at least one false premise: premise 4—that what is "based on an analogical or inductive argument is never certain knowledge" (1959/1983, 225 ff.).[10] As Moore says, he is "inclined to think" that this premise is false (1959/1983, 226). Does a hunch that some premise or step is problematic constitute "exposure" of the argument's flaw? Is it sufficient to be "inclined to think" that some premise is false? Must you be pretty sure that a step is fallacious? Must you have knowledge of what the flaw is in order for knowledge to survive closed-minded dismissal of the argument?

---

[9] There are a number of reasons why you might take time to engage with an argument that you have already closed-mindedly dismissed. You might have purely intellectual interest in how the argument fails, you might have rhetorical interest in convincing some interlocutor, you might have made a bet on whether you can find a flaw in the argument, etc. None of these reasons amount to a concession that the argument has a significant potential for being sound.

[10] My thanks to an anonymous referee for emphasizing this point.

We need not to take a stand on what precisely it means to "expose a flaw" in an argument. Instead, we can say that dogmatism^{FL} comes in varying degrees, corresponding to varying degrees to which a flaw can be exposed. The degree to which a flaw is exposed varies along two dimensions: the degree of exposure, and the specificity of the flaw. The degree of exposure is a function of the degree of certainty that the flaw is exposed and the number of people the exposure would be convincing to. The specificity of the flaw is a function of the degree to which you've narrowed down where and what the flaw is. You, your interlocutor, or those on the fence can be certain, pretty sure, suspicious, or clueless that X is the exact flaw in the argument. Or you, your interlocutor, or those on the fence can be certain, pretty sure, suspicious, or clueless that something involving something like feature F is what the argument does wrong.

I argue that knowledge can survive closed-minded dismissal of some counterarguments even when you're very close to the extreme end of the continuum—when you find each step compelling and you have at most a vague hunch where and what any potential flaw in the argument might be. Call such arguments "apparently flawless" (though this is not an ideal label[11]). From here forward I reserve the label "dogmatism^{FL}" for the view that knowledge can survive closed-minded dismissal of that kind of argument:

> (Dogmatism^{FL}) Knowledge can survive closed-minded dismissal of apparently flawless relevant counterarguments.

---

[11] It is not ideal because the arguments in question are often obviously misleading and, thus, obviously have a flaw. A trick proof that 1=0 can count as apparently flawless in the technical sense I intend because each step can seem compelling and you can fail to expose a flaw. But it doesn't seem quite accurate, in the ordinary sense, to say that the trick proof is apparently flawless, since it obviously has a flaw somewhere. The needed term is one that describes an argument as having the "trappings of soundness" even if it is quite clear that the argument is misleading. It also needs to have an adjectival form. I am unable to come up with anything better than "apparently flawless." I was happy to discover, after settling on the term, that Tom Kelly uses the expression "appears to be flawless" in a similar context (see his quoted passage, below).

Some elaboration is called for regarding the notion of an apparently flawless argument. Inductive arguments, after all, are not logically sound; it is possible for the premises to be true while the conclusion is false, even in good inductive arguments. But noticing that an argument is inductive does not count as exposing a flaw in the argument. Inductive arguments can be apparently flawless.

When you've exposed a flaw in an argument, you've explained why the argument fails to confer justification on its conclusion.[12] Because inductive arguments confer justification on their conclusions, noticing than an argument is inductive does not explain why the argument fails to confer justification on its conclusion. Knowing that the conclusion of an argument is false also does not explain why the argument fails to confer justification on its conclusion. When it comes to trick proofs that 1=0, you may know that the conclusion is false while failing to explain why why the argument fails to confer justification on its conclusion. But when you notice that one of the steps requires division by zero, then you've explained why the argument fails to confer justification on its conclusion.

You can find all the steps in an argument compelling even though you've also exposed a flaw in the argument. The possibility of cognitive illusions entails that, while finding a step compelling, you can recognize that it is fallacious in such a way that it prevents the argument from conferring justification on its conclusion. You can also fail to find the steps in an argument compelling while failing to explain why the argument fails to confer justification its conclusion. For example, some argument that abortion is morally impermissible might use the premise that personhood begins with conception. You might not find that premise compelling, but still not want to count on that as an explanation for why the argument fails to confer justification on the conclusion. But some arguments have both: you find each step compelling and you lack any more than a vague explanation for why the argument fails to confer justification on its conclusion. Those arguments are apparently flawless.

---

[12] There can be multiple explanations for why arguments fail to confer justication on their conclusions because there can be multiple flaws in arguments.

With respect to some of your most strongly held beliefs you might not have had the opportunity to closed-mindedly dismiss such counterarguments. Perhaps you usually can find a flaw. But this doesn't mean that you're not fully prepared to closed-mindedly dismiss such counterarguments if they ever come your way. For many of the issues you have strong positions on—the existence of God, the moral irrelevance of sexual orientation, whether the Holocaust occurred—you might not much care whether, when faced with a relevant counterargument, you end up flummoxed. You might happily maintain outright belief even if you have no clue where the flaw is. If knowledge—or even knowledge-level justification—is the norm of belief, then your disposition is proper only if dogmatism^{FL} is true.[13]

There is an even stronger grade of forward-looking dogmatism according to which, not only can knowledge survive closed-minded dismissal of many relevant counterarguments, but it is sometimes the case that there is no possible counterargument that knowledge wouldn't survive, even if those arguments were known about and understood. According to this view—call it "anti-Quinean dogmatism"—some bits of knowledge are immune from rational revision in the light of new evidence:

> (Anti-Quinean Dogmatism) Some knowledge can survive all possible apparently flawless counterarguments.[14]

---

[13] If dogmatism^{FL} is not true then, in having this disposition, you are disposed to believe improperly. And, to misquote Clifford, it is not possible so to sever the disposition from the belief it suggests as to condemn the one without condemning the other.

[14] This specification of anti-Quinean dogmatism should be tweaked to require that the positive epistemic support for the belief in question be held constant. It might be that some bit of knowledge can survive all possible counterarguments, but only given the kind and strength of evidence in favor of the bit of knowledge. Perhaps this kind of dogmatism is true of some beliefs about simple and obvious conscious episodes (for example, that you are in pain). The untweaked version wouldn't be.

I discuss anti-Quinean dogmatism no further, except to make explicit that dogmatism^{FL} does not entail anti-Quinean dogmatism.

2.3 THE CASE FOR DOGMATISM^{FL}

If you're like me, you sometimes dismiss relevant counterarguments for which you can't, even after reasonably lengthy attempts, produce more than some vague and unconvincing suggestion of where and how the argument breaks down. When it comes to those controversial beliefs for which you have substantial epistemic support—the controversial beliefs that count as knowledge—knowledge can survive closed-minded dismissal of these counterarguments nonetheless. After all, *what's the greater miracle*—that your well-supported belief is wrong or that some clever person has come up with a misleading argument in which you cannot find a hole? The latter happens—at least to me—all the time.

Tom Kelly marshals a similar argument:

> In deciding how to respond to any argument which appears to be flawless, one is in effect in the position of performing an inference to the best explanation . . . If . . . the better explanation of one's failure is one's own cognitive limitations, then one should remain unmoved in the face of the argument . . . Notice that, if this is dogmatism, there is a respect in which it is an unusually modest variety. For when one reasons in this way, one's refusal to change one's beliefs is due to the weight that one gives to one's own cognitive limitations. (2005, 183)

Intellectual modesty is normally thought to require uncertainty in your beliefs. Kelly's point is that intellectual modesty requires you to be just as leery about your ability to expose flaws in relevant counterarguments. Just because you can't expose flaws in some relevant counterargument doesn't mean that flaws aren't there. This is especially so when those counterarguments rely on evidence-types or methods with which you lack sufficient familiarity. If you are an expert on some subject, then this situation might not arise often; perhaps you can usually expose flaws in misleading arguments. But it can happen often with respect to arguments for or against "lay" propositions: propositions about whose subject you lack special training.

Apparently flawless arguments against controversial lay propositions shouldn't be too surprising. After all, there are apparently flawless arguments even for the uncontroversial and obvious falsehoods that nothing moves and that I am not bald. If there are apparently flawless arguments for propositions as obviously and uncontroversially false as those, then there could easily be apparently flawless counterarguments against any controversial lay proposition—even those controversial lay propositions you know.[15] This is the first step in the argument for dogmatism^{FL}:

> (The Principle of Modesty) For any controversial lay proposition you know, there could easily be an apparently flawless relevant counterargument.

Suppose you know some controversial lay proposition, *p*. Because *p* is controversial, there are smart people with an interest in devising convincing arguments for not-*p*. Many of the smart people who devise those arguments would be convinced by them even if they had all of your evidence for *p*. Therefore, those arguments are relevant counterarguments. Because *p* is a lay proposition, many of those arguments will employ methods or evidence-types you can't reliably evaluate. As a result, you won't be able to expose flaws in many of those arguments. And because the people who devised them are smart, one of the arguments could easily be an argument each step of which seems compelling. Therefore, for any controversial lay proposition you know, there could easily be an apparently flawless relevant counterargument. That's the Principle of Modesty.

---

[15] An argument is only apparently flawless when you've actually examined it, find the steps compelling, and are unable to expose a flaw. So, strictly speaking, when you haven't examined an argument, it couldn't easily be apparently flawless. When I say that there could easily be an apparently flawless relevant counterargument, I mean that there could easily be an argument such that, were you to examine it, it would be apparently flawless.

Whether there "could easily" be apparently flawless arguments for not-*p* is an epistemic matter. If your epistemic position excludes the possibility that there are apparently flawless arguments for not-*p*, then not only couldn't there easily be such arguments, there couldn't be such arguments, period. If your epistemic position is consistent with there being such arguments, but makes the existence of such arguments extremely unlikely, then though there could be such arguments, there couldn't *easily* be such arguments. But if your epistemic position allows some significant chance that there are such arguments, even if your epistemic position leans against the existence of such arguments, then there could easily be such arguments.

On this epistemic construal of "could easily," whether there could easily be an apparently flawless relevant counterargument might depend on how versed you are in the literature. If you have made the Holocaust or evolution your life's work and have seen pretty much all the contrary arguments that are likely to be on offer and feel justly confident you have exposed all their flaws, then it might not be the case that there could easily be apparently flawless relevant counterarguments. But this is not so for most people's controversial beliefs, because most people are not extremely well versed in the literature about most of the controversial propositions they believe. *p*, though, is a controversial *lay* proposition. So, there could easily be an apparently flawless relevant argument that not-*p*.

Does the fact that there could easily be an apparently flawless argument that not-*p* mean that you don't know that *p*? It had better not. Because if it does, then many people fail to know that things move simply because they can't figure out how Zeno's arguments go wrong. For many people, there *is* an apparently flawless argument that nothing moves. If knowledge can survive an apparently flawless counterargument, then knowledge can survive it being the case that there could easily be an apparently flawless counterargument. So the mere fact that there could easily be an apparently flawless argument that some controversial proposition is true does not by itself mean that you don't know that the proposition is false: your knowing *p* is consistent with it being the case that there could easily be an apparently flawless argument that not-*p*.

Nothing in the nature of beliefs about controversial propositions entails that—controversiality aside—it is impossible to have very strong support for those propositions: support that in many other cases would suffice for knowledge.[16] Controversial propositions—again, controversiality aside—can enjoy very high levels of evidential support, be true, be unGettiered, etc. Some of the propositions in question are theoretical or abstract, but that alone doesn't preclude having very strong support for them. We all know a range of theoretical and abstract propositions. In any case, many controversial propositions are not theoretical. They are about the approximate age of the earth, whether a certain person was guilty of sexual harassment, and the number of people killed in the Holocaust.

Nor does anything in the nature of beliefs about controversial propositions entail that—controversiality aside—you are prevented from having extremely strong support just because you are a layperson. Laypeople have extremely strong support for many propositions about which they lack the relevant expertise. You have enough support, for instance, to know that the earth revolves around the sun, that the green color in plants is the result of chlorophyll, and that dead animals in rivers increase the likelihood of disease transmission downstream.[17]

Controversiality aside, then, you can have very strong support for the controversial lay proposition, *p*. But, of course, we can't just put controversiality aside. The controversiality of *p* provides two kinds of evidence that might outweigh[18] your strong support for *p*.[19] First there is the mere fact that there is disagreement about whether *p*. Second, there is the fact that there could easily be an apparently flawless argument that not-*p*. I consider the first kind of evidence below. With regard to the second, is the fact that there could easily be an apparently flawless argument that not-*p* guaranteed to be sufficient to outweigh otherwise very strong support for *p*?

---

[16] Here and throughout the book I do not assume that there is any fixed level of evidential support such that when a proposition enjoys that level of support for a believer, the proposition is true and unGettiered, and the belief is based in the right way, then the proposition is known. But I do assume that in every particular case there is some minimal (perhaps vague, but non-zero) degree of evidential support that is required for knowledge, so that falling at or below this level can be the reason why a proposition fails to be known.

[17] Thanks to Elizabeth Brake for suggesting these examples.

Nathan Ballantyne (2015b) might be taken as arguing that the fact that there could easily be an apparently flawless argument that not-*p* makes it extremely difficulty to maintain knowledge that *p*. Ballantyne's concern is what he calls "the problem of unpossessed evidence" (316). The worry, with respect to many of our otherwise rational beliefs, is that when it becomes clear to us that there is a large body of evidence we aren't aware of, and thereby becomes clear that there could easily be undefeated defeaters to our positions and arguments, otherwise rational belief becomes irrational. Ballantyne asks us to consider a case he calls WWW:

> Fifteen years ago, you thought carefully about economic ideas and arguments. Then life changed. With a family and busy job, you haven't kept pace with recent discussion. Now you wonder what has happened in the intervening years, so you search on Google and JSTOR with some relevant keywords ('government spending economic growth') and your searches return thousands of results. As you quickly recognize, there are hundreds of articles and books, all potentially relevant to figuring out what to think about this one economic issue, about which you once had carefully considered views. You knew the arguments, replies, and counter-replies, and you had a good rationale for your favoured positions. It's evident some of the recent discussion challenges your thinking. But you

---

[18] Perhaps disagreement about *p* doesn't *outweigh* your otherwise knowledge-level evidence for *p*. Perhaps disagreement *implies* that you don't have knowledge-level positive evidence for *p*; if knowledge-level positive evidence existed, everyone would already believe that *p* and there would be no disagreement. This objection is unsuccessful because it assumes, falsely, that people are always convinced by knowledge-level evidence.

[19] There is another way that the salience of an apparently flawless counterargument might destroy knowledge. David Lewis, in developing his contextualist account of knowledge, argues for a "Rule of Attention" (1996, 559), according to which mere attention to uneliminated possibilities of error can raise the standards for knowledge high enough that ascriptions of knowledge all go false. We might think that attention to the possibility or actuality of apparently flawless counterarguments is automatically knowledge-destroying in this same sense.

I disagree, for the same reasons that most subsequent contextualists differ from Lewis in requiring not mere attention, but serious attention to the possibility of error in order for the standards to be driven to knowledge-destroying heights. Knowledge survives mere non-serious attention to salience of error, and knowledge survives mere attention to apparently flawless counterarguments, as it does with Zeno's paradox and trick arguments that 1=0. I see no reason why the situation is different when it comes to controversial propositions.

> have not studied any of it. Rehearsing your earlier rationale, it still seems perfectly right to you, but you know of new evidence you do not have. (2015b, 315 ff.)

According to Ballantyne, in many cases like WWW, though you may have started with rational belief, upon learning of the unpossessed evidence, your belief is no longer rational. If you lose rational belief, I'll suppose, then you lose knowledge too.

One way that Ballantyne considers in which your rational belief could survive confrontation with unpossessed evidence is if you have certain kinds of defeaters for that evidence. For example, you might "learn that Big Smoke, a cigarette manufacturer and proud sponsor of stock car racing, has funded scientists to investigage the longer-term health effects of smoking" and that "Big Smoke's research team concludes that smoking is not a health hazard" (2015, 322). Though you don't bother to investigate the details and so don't possess the evidence that Big Smoke is presenting, you have a defeater for that evidence: Big Smoke is unreliable with respect to the evidence in question.

In the cases of interest to us, such defeaters can't be counted on. If you had such defeaters, the various steps in the counterargument would no longer be compelling; you'd be able to dismiss them by saying, for example, that the source of the argument is an unreliable reporter of the relevant facts. Therefore, if you had such defeaters, the counterarguments that could easily exist couldn't easily be apparently flawless. In the situations we're concerned with, not only could there easily be arguments against your beliefs, there could easily be apparently flawless arguments.

What we have in the cases of interest are not defeaters of the sort invoked against Big Smoke, but direct positive support for your position. Ballantyne does allow this kind of defeater to save your rational belief from defeat by unpossessed evidence. If, for example, you know that "[t]he evidence I have for believing *p* conclusively establishes that *p* is true" (2015, 326) then your rational belief (and, I'll assume, knowledge) can survive unpossessed evidence you wouldn't be able to defeat in other ways. But, argues Ballantyne, while this kind of defeater may save your *cogito*-inspired rational belief that you exist, it's not enough to save the majority of beliefs of the sort that are problematized in cases like WWW. As Ballantyne concludes, "We rarely have such powerful reasons when it comes to controversial matters in philosophy or economics, but if we do, our acquaintance with the available evidence will probably be more than casual" (2015, 326).

Ballantyne's conclusion, then, proves not to be incompatible with the conclusion reached here: that knowledge can survive it being the case that there could easily be an apparently flawless argument against what you know to be true. If there is a disagreement with Ballantyne, it's a disagreement in scope. Ballantyne thinks rational belief (and, I'll assume, knowledge) survives "rarely." I'm not sure how to quantify the rarity of such survival, but I grant that I'm probably more optimistic than Ballantyne. I think, for example, that knowledge can survive even when your acquaintance with the available evidence is no more than casual. Knowledge about the Holocaust survives, for example, when confronted with the unpossessed evidence presented by mountains of Holocaust-denial literature. Knowledge that homosexuality isn't a disability and isn't something that should be prevented or "cured" survives finding out that an otherwise excellent philosopher has presented arguments to the contrary in a keynote address.[20] Knowledge that vaccines don't cause autism survives the unread evidence offered by anti-vaxxers.

Knowledge, in these and similar cases, can survive not only when there could easily be apparently flawless arguments to the contrary, but when there *are* apparently flawless arguments to the contrary. That’s because, when it comes to controversial propositions, that there is an apparently flawless argument that not-*p* is not, in general, very strong evidence for not-*p*, especially when you lack the facility required to reliably evaluate that argument. There are often apparently flawless arguments for both sides of controversial propositions. But arguments on at most one of the sides have true conclusions. So, a very high proportion of those apparently flawless arguments—especially those you lack the facility to reliably evaluate—have false conclusions.

---

[20] See (Hackett 2016). For a response to Swinburne's argument see (Littlejohn 2016). For the examples of the demand to engage with Swinburne's arguments before dismissing them, see the discussion in the Preface to this book.

If that there is an apparently flawless argument for not-*p* is not in general very strong evidence for not-*p*, then that there merely could easily be an apparently flawless argument for not-*p* is even worse evidence for not-*p*. The mere fact that you have some *prima facie* evidence for not-*p* does not preclude your knowing that *p*. What then should we say about your overall support for *p* when, other than the ease of there being an apparently flawless argument that not-*p*, you have extremely strong support for *p*? At the very least, this: it is not guaranteed that what is otherwise extremely strong support for *p* is inevitably reduced to a level insufficient for knowledge. On the contrary, often your strong support for *p* will swamp the quite weak evidence for not-*p*.

Wouldn’t the same argument that shows that the fact that there could easily be an apparently flawless argument that not-*p* is not very strong evidence for not-*p* also show that the fact that you have seemingly strong evidence for *p* is not very strong evidence for *p*? It might in some cases. But, of course, your evidence for *p* is not just that you have seemingly strong evidence for *p*. You also have the first-order evidence itself. And, the point here is that this first-order evidence, if knowledge-level strong, will often swamp the weak evidence for not-*p* provided by the fact that there could easily be an apparently flawless argument for not-*p*.[21] Knowing *p* is consistent with it being the case that there could easily be an apparently flawless argument that not-*p*.

Now suppose you become familiar with an apparently flawless counterargument; you find that a counterargument is apparently flawless in the sense that you come to understand the details of the counterargument but, after reasonably lengthy attempts, you fail to expose a flaw. With respect to many such arguments, this shouldn't be very exciting news to you: there's nothing too surprising about finding out that something that could easily be, in fact is. It's not that finding out that there is such a counterargument can never make a difference to epistemic status. As David Christensen says, "it makes a difference whether my friend actually does disagree with me, or whether I just know she could have disagreed with me" (2007, 208). In general, evidence that could easily exist can be epistemically more forceful when you find out it actually exists. But at least in a lot of cases, the fact that it could easily exist means that it won't make enough of a difference, when it is actually discovered, to destroy knowledge if that knowledge existed in the first place.

As an analogy, consider the Lady Galadriel's remarks to Sam Gamgee in J. R. R. Tolkien's *The Lord of the Rings* after Sam looks in the Mirror of Galadriel and, seeing mischief and destruction taking place in his home, the Shire, commits to go home. In response, Galadriel points out to Sam, "You did not wish to go home without your master before you looked in the Mirror, and yet you knew that evil things might well be happening in the Shire" (Tolkien 1966, 378).

---

[21] For more on asymmetries between your positive evidence for *p* and the evidence provided by counterarguments to *p*, see Chapter 3.

Galadriel has it right. There is something odd about giving up an intention that you formed knowing that something might very well be happening, just when you find it is actually happening.[22] Maybe sometimes the difference in certainty between the easy possibility and the actuality can make a difference. But it won't always make a difference, if your reasons for forming the intention were strong enough. The same goes for knowledge. If you have knowledge even though there could easily be an apparently flawless argument somewhere out there, then at least sometimes your knowledge will survive encountering one of those arguments.

Of course, you did not expect *this* argument.[23] Perhaps the specific content of the argument can, at times, destroy knowledge even if the mere fact that you find the argument apparently flawless cannot. But this is often not what happens when you become familiar with novel, apparently flawless counterarguments. It is not, for example, what happens when you become familiar with apparently flawless arguments that there is no motion. It's not what happens at magic shows.

That the illusionist is not really burning to death and that there is motion, while true, are not controversial. But the controversiality of the matter shouldn't make a difference. As long as you often can be epistemically quite well situated with respect to controversial truths—as long as you can often know them—the weighing should be favorable to maintaining belief that *p* in the face of apparently flawless arguments that not-*p*, even if those arguments are relevant counterarguments. Put another way,

---

[22] One way to explain the oddity is in terms of weakness of will. This would be in keeping with Richard Holton's suggestion that "weak-willed people are irresolute; they don't persist in their intentions; they are too easily deflected from the path they have chosen" (Holton 1999, 241).

[23] One way that the specific content of the argument can destroy your knowledge is if the argument presents evidence for not-*p* that you'd be very surprised to be misled by. This can happen if you have expertise that makes you a reliable evaluator of the evidence presented by the argument, or if the argument presents evidence that is of the same kind as the evidence you have for your view. For more on why these factors can destroy knowledge, see Chapter 3.

> Epistemic support strong enough to allow you knowledge despite the fact that there could easily be an apparently flawless relevant counterargument is sometimes strong enough to allow you knowledge despite the fact that there is an apparently flawless relevant counterargument.

In slogan form: *unsurprising arguments are often impotent*. I give a more formal argument for this principle in Chapter 3.

The principle is not in conflict with the standard solution to the so-called "dogmatism paradox" suggested by, e.g., Gilbert Harman (1973), Carl Ginet (1980), Roy Sorensen (1988), and Earl Conee (2001). According to the dogmatism paradox (usually, following Harman, attributed to Saul Kripke), if you know, say, that Tom stole the book, then you know that any evidence (argument) that Tom did not steal the book is misleading. If skepticism is false, you can know that Tom stole the book. Therefore, you can know ahead of time that any evidence that Tom did not steal the book is misleading. What happens if, knowing this, you subsequently get evidence that Tom did not steal the book? Are you thereby justified in ignoring that evidence? Seemingly not. You don't get license to ignore all future evidence that Tom did not steal the book just by now having enough evidence to know that Tom did steal it.

On the standard solution, you can indeed know, before acquiring evidence that Tom did not steal the book, that the evidence is misleading, even if you are not thereby justified in dismissing that evidence when it comes along. That's because, when the evidence comes along, you lose knowledge that Tom stole the book, and so lose knowledge that the evidence is misleading. As Harman puts the point,

> Since I now know that Tom stole the book, I now know that any evidence that appears to indicate something else is misleading. That does not warrant me in simply disregarding any further evidence, since getting that further evidence can change what I know. In particular, after I get such further evidence I may no longer know that it is misleading. For having the new evidence can make it true that I no longer know that Tom stole the book. (1973, 149)

When it comes to the dogmatism paradox, there is no presumption that there could easily be evidence that Tom did not steal the book. There is only a presumption that such evidence is possible. More importantly, there is no presumption that there could easily be a relevant counterargument against the claim that Tom stole the book: there is no presumption that there could easily be evidence that would convince lots of people that Tom did not steal the book and would convince those people even if they had your evidence that Tom did steal the book. Once that presumption is added in, it's much harder to see how knowledge could standardly be lost, if it is there in the first place.[24] Knowledge may very well be destroyed by discovering surprising evidence. This is consistent with it being impossible (or difficult) for knowledge to be destroyed by the discovery of unsurprising evidence. It's this latter impossibility (or difficulty) that is required here: if you should already suspect, before exposure, that the evidence is there, then exposure to the evidence will often not change whether you know.

One final proviso: we need to allow for the possibility both that there can be close calls when it comes to knowledge and that exposure to an apparently flawless relevant counterargument, even if it leaves knowledge intact, always reduces your strength of epistemic support.[25] If both of these things are the case, then if you just barely know a controversial proposition, exposure to any apparently flawless relevant counterargument will reduce your strength of epistemic support below the level required for knowledge. Therefore, it's not the case that whenever you know a controversial proposition, your knowledge can survive exposure to apparently flawless relevant counterarguments. But if you can know controversial propositions, some of this knowledge won't be of the close call variety. Therefore, it will sometimes be the case that, if there could easily be an apparently flawless relevant counterargument, finding the apparently flawless argument won't weaken your epistemic support below the level required for knowledge.

---

[24] Is it possible for you to know that Tom stole the book once it is no longer surprising that there is evidence that Tom did not steal the book? Yes, if what's unsurprising is that there is *prima facie* evidence that Tom did not steal the book. Knowledge is compatible with having *prima facie* evidence to the contrary, so knowledge is compatible with it being unsurprising that there is *prima facie* evidence to the contrary.

[25] Though I want to allow for its possibility, I do not endorse this latter claim. It is plausible that your first confrontation with Zeno's arguments reduces strength of epistemic support not one iota for the proposition that there is motion.

Therefore, we have two premises that can figure in an argument for dogmatism^{FL}. The first follows from the Principle of Modesty:

1. If you know some controversial lay proposition, then there could easily be an apparently flawless relevant counterargument.

The second follows from the principle that unsurprising arguments are often impotent:

2. If there could easily be an apparently flawless relevant counterargument against a proposition, then sometimes it won't make a difference to whether you know the proposition if there is an apparently flawless relevant counterargument.[26]

It follows from 1 and 2 that, if you know some controversial lay proposition, it sometimes won't make a difference to whether you know that proposition if there is an apparently flawless relevant counterargument. And it follows from this that

3. If you know some controversial lay proposition, sometimes you will know that proposition even if there is an apparently flawless relevant counterargument.

If the antecedent of this conclusion is necessarily false—if it is impossible to know controversial lay propositions—then 3 (and 1 and 2) is trivially true. Is it possible to know controversial lay propositions? Is the mere fact of significant disagreement about *p* guaranteed to undercut whatever positive epistemic support you have for *p*? Some who work on the epistemology of disagreement might be interpreted as arguing that you can't know anything when there is significant disagreement. As Richard Feldman argues,

> I see that another person, every bit as sensible and serious as I, has an opposing reaction. Perhaps this person has some bit of evidence that cannot be shared or perhaps he takes the evidence differently than I do. It's difficult to know everything about his mental life and

---

[26] This assumes that belief is sometimes preserved after exposure to the argument.

> thus difficult to tell exactly why he believes as he does. One of us must be making some kind of mistake or failing to see some truth. But I have no basis for thinking that the one making the mistake is him rather than me. And the same is true of him. And in that case, the right thing for both of us to do is to suspend judgment. (2007, 229)

Others who take similar lines include David Christensen (2007) and Adam Elga (2007). Following Lackey (2010), call those philosophers who think confrontation with significant disagreement standardly requires significant revision of belief "conformists." Assuming that the revision should be significant enough, conformists will also think that confrontation with significant disagreement standardly entails loss of knowledge.[27]

If conformism is committed to the claim that you cannot reasonably believe—and hence cannot know—propositions about which there is significant disagreement, this seems like a difficulty for conformism. At least, *so say conformists themselves*. Elga calls it "the problem of spinelessness"—"the problem that a [conformist] view on how to respond to disagreement will recommend suspension of judgment on virtually all controversial issues" (2007, 492). Christensen, though not emphasizing controversiality, likewise considers it a significant worry for conformist views if they always require giving up extremely well supported beliefs whenever faced with disagreement:

> The first sort of hard cases are ones where an agent begins with extremely high rational confidence in her belief. In various such cases, it seems wrong to hold that she should revise her belief much at all, even if the agent's friend disagrees sharply, and even if, before discovering the disagreement, she would have considered the friend her epistemic peer on the sort of issue in question. (2011, 8)

---

[27] We shouldn't be too quick to saddle any of these authors—Feldman included—with the view that knowledge is impossible in the face of disagreement. In addition to the reasons offered in the body of the text, below, some of the authors (e.g. Christensen and Elga) require merely that you move your credence toward the credence of the disagreeing interlocutor. This movement need not result in suspension of belief if, for example as Christensen points out, "my friend has a degree of confidence barely sufficient (given the context) for rational belief that not-*P*, but that I have a degree of confidence much greater than that required for believing *P*" (2007, 214).

One way for conformists to respond to such cases is to bite the bullet and say that in all such cases, conformism requires significant belief revision. But this is not what actual conformists do. What they do is find ways to accommodate within their view the result that you can reasonably maintain strong confidence even when faced with disagreement:

> we can have reasonable beliefs (but perhaps not knowledge) about some political, scientific, philosophical, or religious matters, even if disagreement sometimes undermines our justification for those beliefs. (Feldman 2006, 222)
>
> contrary to initial impressions—the equal weight view does not require one to suspend judgment on everything controversial (Elga 2007 494)
>
> The personal information I have about myself . . . provides a perfectly reasonable basis for continuing to think that our shares of the bill are much more likely to be \$43 than \$45, despite discovering the disagreement of my heretofore-equally-reliable friend. (Christensen 2011, 9–10)
>
> Conciliationism does not require agnosticism in response to disagreements about the most basic moral beliefs . . . Conciliationism doesn't require us to suspend judgment about such matters. (Vavova 2014, 302)

Allowing for reasonable belief in the face of disagreement is still short of allowing for knowledge in the face of disagreement (and Feldman's parenthetical explicitly separates the issues). But if conformists can allow reasonable or justified belief in controversial propositions, there is no principled reason why they can't also save knowledge. If controversial propositions can be believed with the strength of justification sufficient for knowledge—knowledge-level justification—then those propositions can satisfy the other conditions on knowledge: they can be true and unGettiered, for example. So the only way to insist that controversial propositions can be justifiedly believed but not known would be to insist that controversial propositions can be justifiedly believed, but not knowledge-level justifiedly believed.

Why think that? The degree to which a belief is justified is a function of the positive support the belief enjoys and the strength of the countervailing evidence. Again, there is no principled limit on the degree of positive support you might have for controversial propositions. And *if* conformists can make plausible the position that reasonable belief can be maintained in the face of significant disagreement, then they have found a way to minimize the strength of the countervailing evidence—the evidence, that is, that there are seemingly reasonable people who disagree. If they have found such a way, then depending on the situation, there is no principled bar to that countervailing evidence being weakened to the extent that the belief is knowledge-level justified and, hence, potentially known.

Can conformists allow reasonable belief to be maintained in the face of significant disagreement? They think they can. One way is to allow you to fail to regard the disagreeing party as an epistemic peer. Elga, for example, maintains that in "messy, real-world cases,"

> one's reasoning about the disputed issue is tangled up with one's reasoning about many other matters. . . . As a result, in real-world cases one tends not to count one's dissenting associates—however smart and well-informed—as epistemic peers. (2007, 492)

Similarly, Katia Vavova argues that once we disallow facts about the disputed issue itself to play a role in evaluating whether those who disagree with me are my epistemic peers,

> I have no reason to think you are better, or even as well, positioned as I am on this matter. . . . I am unable to make any judgment about our comparative reliabilities. In the absence of this, I have no idea how to take your opinion into account. In such cases, Conciliationism remains silent. (2014, 310)[28]

A different strategy is suggested by the conformist Tomas Bogardus. According to Bogardus, in the problematic cases, you can "just see" that the disagreeing party's view is wrong and thereby acquire evidence that the disagreeing party lacks:

> In philosophy-speak, we might say Smith comes to have *knowledge from direct acquaintance* in the problematic cases. . . . A relevant piece of evidence is intellectually obvious to Smith; she has unmediated cognitive access to the truth of a pertinent proposition. Her knowledge does not rely on any report, indication, or representation. (2009, 331)

---

[28] For a reply to the kind of defense of conformism suggested by Vavova and Elga, see (Fritz forthcoming).

Christensen, alternatively, emphasizes the possibility of knowing that your belief-forming methods were not unreliable in various ways that, for all you know, plague the disagreeing party. He licenses the following reasoning:

> The belief my friend announced was incompatible with the one at which I arrived. This is strong evidence that one of us did not arrive at his or her belief in a highly reliable way, or that one of us is not sincerely announcing his or her belief. I can eliminate (via personal information) many of the ways that I could have failed to use a reliable method, as well as the possibility that my announcement was not sincere. But I cannot eliminate analogous possibilities for my friend. So it's likely that she did not sincerely announce a belief that was formed by a highly reliable method. (2011, 10)

What are the ways one might fail to have used a reliable method? Christensen mentions exhaustion, psychotic attacks, drugs, insincerity, and lack of care. While the first three are relatively temporary interferences with reliability that do not realistically affect large numbers of people who disagree with you, insincerity and lack of care may be more pervasive. We might add to the list racism, wish fulfillment, Hume's "passion of surprize and wonder," iconoclastic motivations, and other psychological pressures to form beliefs that do not track the truth. To the extent that when your belief is not formed on the basis of methods influenced by such factors, it will be more accessible to you that it isn't than that the ill-formed beliefs of others aren't influenced by such factors, Christensen's strategy will allow you to believe reasonably in the face of disagreement from one you would otherwise regard as an epistemic peer.

Finally, both Christensen and Elga mention another way to legitimately disregard a disagreeing party. In order for significant belief revision to be mandated, it's not enough that your evidence be neutral or indeterminate on the subject of whether the disagreeing party is an epistemic peer. You must have strong positive evidence that the disagreeing party is an epistemic peer. (Vavova (2014, 320 ff.) makes a similar point about moral disagreement).

Elga is more implicit on this issue. In the text immediately prior to the passage cited above, Elga says that “in the clean cases one is in a position to count one’s associates as peers *based on reasoning that is independent of the disputed issue*. But in the messy real-world cases one is rarely in a position to do so” (2007, 492). Why, according to Elga, can you maintain reasonable belief in the face of messy real-world disagreement? It is because, in messy real-world cases, you are so rarely in a position to conclude that a subject is an epistemic peer.

Christensen is quite explicit about the requirement of positive evidence of epistemic peerage in his rejection of the following principle of belief revision:

> (A) Insofar as the dispute-independent evaluation *fails to give me* good reason for confidence that I’m better informed, or more likely to have reasoned from the evidence correctly, I must revise my belief in the direction of the other person’s. (2011, 15)

In place of (A), Christensen recommends (B):

> (B) Insofar as the dispute-independent evaluation *gives me* good reason to be confident that the other person is equally well-informed, and equally likely to have reasoned from the evidence correctly, I must revise my belief in the direction of the other person’s. (2011, 15)

When it comes to the controversial lay propositions for which you have the strongest support, it is not obvious that dispute-independent evaluation would provide strong positive evidence that the disagreeing parties are your epistemic peers. At the very least, it doesn’t follow from the mere fact that a proposition is a controversial lay proposition that dispute-independent evaluation would provide such evidence. On the contrary, with respect to many of the relevant issues—Holocaust denial, for instance—it seems that there would be positive evidence that the disagreeing parties were not epistemic peers, especially if Christensen’s “personal information” were allowed to count as relevant evidence.

Do these strategies save conformism from the objection that conformism doesn't allow reasonable belief in controversial propositions? For the purpose of defending dogmatism^{FL}, no answer is required. The point is not that any of the conformists' strategies for allowing reasonable strongly held belief in the face of apparent peer-disagreement are successful. The point is that conformists feel the need to have the strategies in the first place: they do not shrug and insist that reasonable belief in controversial lay propositions is impossible. They think it's important that conformism not have the consequence that reasonable belief in controversial propositions is impossible.

It might be that reasonable belief in—and, thus, as I argue above, knowledge of—controversial propositions is possible but uncommon. A full defense of the stronger conclusion would require a case-by-case argument that a significant proportion of controversial propositions are potentially known. Such a thorough examination of controversial propositions is, to quote the seer, "beyond the scope of this book," though I consider two candidates in Chapter 4. But if knowledge of controversial lay propositions is possible, there is nothing in principle standing in the way of it being the case that you often can know controversial lay propositions, because there is nothing in principle standing in the way of you having a very strong support for controversial lay propositions.[29]

---

[29] I include moral propositions among those controversial propositions you can know. This might seem to commit to metaethical cognitivism. But many noncognitivists have stories to tell about moral knoweldge. In particular, many noncognitivists have stories to tell about when you should act on moral beliefs. What's important about knowing a proposition, for my purposes, is that when you know a proposition, you should act as if it is so (see Chapter 6 for details). Therefore, as long as a noncognitivist allows that you sometimes should act on your moral beliefs and sometimes not, the larger conclusions of this book follow for moral judgments as well as factual controversial beliefs. Allan Gibbard, for example, seems to allow for this position: "we can depend ultimately on our own lights, even when we disagree with others, and legitimately count what we end up with as knowledge. If I am right in what I have been saying, then some of this will be knowledge of good and evil" (2009, 18). Thanks to Josh Stein for discussion.

Because you can know (and not just barely) controversial lay propositions, from 3 we can derive:

4. Sometimes you know some controversial lay proposition even if there is an apparently flawless relevant counterargument.

Arguments are apparently flawless just in case, after understanding the details and making reasonably lengthy attempts, you find each step compelling and fail to expose a flaw. Therefore, 4 entails that knowledge can survive dismissal of relevant counterarguments, even when you understand the details, find each step compelling, and, after reasonably lengthy attempts, fail to expose a flaw. That's dogmatism^{FL}.

None of this is to say that it won't often be of interest to investigate relevant counterarguments—even the ones closed-minded dismissal of which knowledge survives. Nor is it to say that studying such arguments can't teach us interesting things about the relevant objects of study. Zeno's arguments, after all, are of interest even though they pose no threat to knowing that there is motion. By figuring out where they go wrong, or even trying to, you might figure out legitimate and illegitimate uses of the concept of infinity, or you might refine your concept of motion, or you might simply find engagement with the arguments fun. You are not, of course, required either intellectually or morally to investigate Zeno's arguments, either in order to retain belief or knowledge that there is motion, or for any other reason (unless it's your job). Some people are interested in such things and some aren't. The same might be said for closed-mindedly dismissed arguments in more controversial domains, though the situation is not so clear. There might be moral requirements to fully investigate arguments for false moral beliefs; this is the focus of the second part of this book. But there is no requirement to fully investigate such arguments for the purpose of retaining knowledge and, hence, justified belief about the relevant propositions.

2.4 CONCLUSION

Whether knowledge survives closed-minded dismissal of a counterargument depends on two factors: the nature and strength of support you have for your position, and the nature and strength of support provided by the counterargument. I will not discuss in much depth when support for your position is strong enough and of a kind to count as otherwise knowledge-level. That discussion is a discussion of epistemology proper, and its resolution is independent of issues in the epistemology of disagreement and what the appropriate reaction is to apparently flawless counterarguments. It's hard enough to know when support counts as knowledge-level even for propositions for which there are no relevant counterarguments—like the proposition that here's a hand or that Moscow is in Russia. The remaining question is when the support provided by a relevant counterargument is of a nature and strength that allows for closed-minded dismissal. I turn to this subject in Chapter 3.

3

The Epistemic Efficacy of Amateurism

> [C]urrently popular theories of knowledge have the surprising consequence that stupidity can enhance, and intelligence diminish, one's prospects for knowledge. So if any of these theories is correct, Socrates may have known less than others precisely because he was wiser than they. . .
>
> But this conclusion should not be construed as a counterexample to currently popular theories. In some cases at least, it seems reasonable to believe that Watson knows more than Holmes. The blunt man of solid, uninspired common sense, being untroubled by subtleties, may know what's what, while the more sensitive, finely tuned intelligence is distracted by nuances. (Catherine Elgin (1988, 297 and 310))

3.1 COINS AND COUNTERARGUMENTS

When can knowledge survive closed-minded dismissal of a counterargument? I open this chapter by giving away the answer: knowledge can survive when the counterargument is too sophisticated for you. More precisely, knowledge can survive when you can't reliably evaluate the counterargument because you lack expertise in the evidence types and argumentative methods invoked by the counterargument. In this way, amateurism can be epistemically efficacious.

This is not to deny that expert opinion is, in general, more reliable than and epistemically superior to lay opinion. To an extent, John Hardwig is right that

> Within her area of expertise, an expert's opinion is better than a non-expert opinion. By "better," I mean more reliable. . . . Areas in which expert opinion exists and is available are areas in which one ought not to make up one's own mind—without first becoming an expert. (1994, 84–85)[1]

In addition to a general tendency toward reliability, there might be other reasons to develop expertise in the methods and evidence types employed by counterarguments against what you know. You might need to have expertise with those methods and evidence types in order to explain why what you know is true, to understand the underlying nature of the items about which you have knowledge, or to convince those who disagree with you. But you don’t need to have expertise with the methods and evidence types employed by counterarguments against what you know in order to retain that knowledge. And sometimes, contra Hardwig, it is your very lack of expertise that allows you to retain knowledge while an expert does not, even about the subject matter about which the expert is an expert.

---

[1] See also (Hardwig 1985, esp. 339–343).

The answer is easier than the path to the answer. The initial climb especially has some rocky terrain, but things smooth out in section 3.3 and the path passes by a nice scenic overlook in section 3.4. The first step on that path uses a well-known and uncontroversial consequence of Bayesian epistemology: surprising evidence has greater confirming power than unsurprising evidence. It is unsurprising for a coin to come up heads on a single flip, even if the coin isn’t weighted toward heads. That is why such a result only very slightly confirms that the coin is weighted toward heads. But unless the coin is weighted toward heads, it is very surprising that it would come up heads on each of 1,000 flips. That is why 1,000 heads in a row strongly confirms that the coin is weighted toward heads. Surprising coin flips have greater confirming power than unsurprising coin flips.

As go coins, so go arguments. If you shouldn’t expect an argument to be apparently flawless unless it were sound, then if the argument is apparently flawless, it’s probably sound. If you should expect the argument to be apparently flawless even if it’s misleading, then if it’s apparently flawless, it may well be misleading. In such cases, the apparent flawlessness of the argument does not strongly confirm its soundness, and so does not strongly confirm the truth of its conclusion. Upshot: in such cases, knowledge that the conclusion is false can survive the argument being apparently flawless.

How unsurprising must it be for an argument to be apparently flawless in order for knowledge that its conclusion is false to survive? To answer that it is necessary to do a little math. Assume that you start out—prior to finding an argument apparently flawless—knowing that the conclusion of the argument is false and, so, knowing that the argument is misleading. Knowledge can survive if, after finding the argument apparently flawless, you still can know that the argument is misleading; for you know that the conclusion of the argument is false if and only if you know that the argument is misleading.

To retain knowledge that an argument is misleading after finding it apparently flawless, the probability that the argument is misleading must remain knowledge-level high after finding the argument apparently flawless; the probability that (M) the argument is misleading conditional on (AF) the argument is apparently flawless[2] must remain (k) knowledge-level high. Formally,

$$P(M|AF) \geq k.$$

The relevant instance of Bayes' Theorem tells us that this probability can be calculated as follows:

$$P(M|AF) = \frac{(P(AF|M))(P(M))}{P(AF)}$$

As it gets less surprising that the argument would be apparently flawless even if it were misleading, P(AF|M) approaches and can even surpass P(AF). When P(AF|M) is at least equal to P(AF) then the ratio between the two (on the right side of the equation) is at least 1. In that case P(M|AF) will be no smaller than P(M). Because you know prior to investigating the argument that it is misleading, the prior probability that the argument is misleading is very high:

---

[2] P(M) is not determined by the objective probability that the token argument is misleading, but the probability that an argument with the salient features of the token argument is misleading. P(AF) is not determined by the objective probability that the token argument is apparently flawless, but the probability that an argument with the salient features of the token argument is apparently flawless.

knowledge-level high. That is, $P(M) \geq k$. Therefore, when $P(AF|M) \geq P(AF)$, $P(M|AF)$ is no smaller than k; knowledge can survive the argument being apparently flawless. If $P(M)$ is sufficiently above knowledge-level, so that even small decreases in probability leave $P(M|AF)$ knowledge-level high, then knowledge can survive the argument being apparently flawless even if $P(AF|M)$ is somewhat less than $P(AF)$.

This delivers the first principle in an argument for the epistemic efficacy of amateurism:

> (Inertness of Unsurprising Error) If it is sufficiently unsurprising that a counterargument would be apparently flawless if it were misleading, then knowledge can survive the counterargument being apparently flawless

where, again, the degree of unsurprisingness is "sufficient" if the ratio of $P(AF|M)$ to $P(AF)$ is at least 1 or, if lower than 1, not so much lower that it can't be compensated for by the degree to which $P(M)$ is above the threshold for knowledge.

While the Inertness of Unsurprising Error lessens the impact of apparently flawless counterarguments, the news for the dogmatic is not wholly good. Similar reasoning shows there are cases in which you shouldn't have equanimity when face with counterarguments that are apparently *flawed*. The prevalence of confirmation bias suggests that the prior probability that you will find flaws even in sound counterarguments is relatively high. You are guilty of confirmation bias when you evaluate a counterargument in a way that protects your current beliefs. You can wrongly think there are flaws in sound arguments, and you can wrongly think that you've exposed a flaw in fallacious or misleading arguments, as well. Because of this, the probability that an argument is apparently flawed conditional on its being misleading will in some cases be not too much higher than the prior probability that the argument is apparently flawed. In those cases, by the relevant application of Bayes' Theorem, the probability that the argument is misleading conditional on it being apparently flawed is not too much different from the prior probability that the argument is misleading. In short, an argument being apparently flawed, when the ratio is close enough to 1, sometimes won't allow you to significantly increase your confidence that the argument is misleading.[3]

3.2 HIGHER-ORDER EVIDENCE

Here's a worry about the Inertness of Unsurprising Error. The principle says only that the mere fact that a counterargument is apparently flawless sometimes can leave knowledge intact. But when a counterargument is apparently flawless, your total evidence usually includes more than the fact that the counterargument is apparently flawless.[4] You also have the evidence provided by the counterargument itself. If the counterargument provides good evidence—even if it's good but misleading evidence—then your total evidence might destroy knowledge even if the mere fact that a counterargument is apparently flawless leaves knowledge intact.

The solution to this worry is found in the literature on higher-order evidence. Higher-order evidence, most generally, is just evidence about evidence, while lower-order evidence is evidence that's not about evidence. The general construal of higher-order evidence is precisified in various ways. David Christensen (2010) uses the term "higher-order evidence" to mean evidence about how your cognitive faculties are functioning. Tom Kelly (2010) uses the term "higher order evidence" to mean evidence about whether your evidence is good. I construe

---

[3] Thanks to Catarina Dutilh Novaes for suggesting I discuss the issue of apparently flawed counterarguments.

[4] Here's Roger White making a similar point:

> Normally when I draw some conclusion from premises, whether by application of logic, statistical inference, or on the basis of the explanatory virtues of a theory, or what have you, I don't take my evidence to consist solely in the fact that my conclusion *seems* right to me. My evidence consists in the premises from which I reasoned. (2010, 602)

"higher-order evidence" broadly enough to include both these senses, but also to include evidence that there is evidence.[5]

Sometimes your higher-order evidence and lower-order evidence supplement each other. If you correctly and competently solve a moderately difficult logic puzzle, your lower-order evidence—the facts of the puzzle—confer justification on the conclusion. If a logician testifies that the facts of the puzzle are good evidence for the conclusion you've reached, then you have higher-order evidence that the conclusion is true and your justification should increase. In this case, both your higher-order evidence and lower-order evidence contribute to the justification of your conclusion. But they don't always do so.

According to Christensen, you should sometimes "bracket"[6] parts of your evidence; there are times when parts of your total evidence do not make a difference to the justification of your belief even though, on its own, that part of your evidence would be relevant to the truth of the belief. To illustrate, here is a large circle:

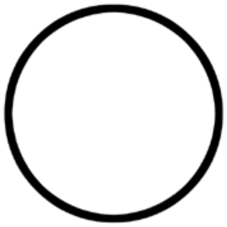

You have perceptual evidence that there is a large circle on the page. That you have perceptual evidence that there is a large circle on the page is additional evidence that there is a large circle on the page; it is distinct from the perceptual evidence itself.[7] But the support your total evidence confers on the proposition that there is a large circle on the page is exactly what is conferred by your perceptual evidence. You should bracket the higher-order evidence—that you have perceptual evidence that there is a large circle on the page—because your higher-order evidence doesn't contribute support to the proposition that there is a large circle on the page over what is contributed by your lower-order, perceptual evidence. Your higher-order evidence piggybacks on your lower-order evidence.

---

[5] "Evidence" I also construe broadly. Arguments that confer justification on *p* are evidence for *p*. Therefore, evidence that there is an argument that confers justification on *p* is evidence of evidence for *p*—is higher-order evidence for *p*. Because apparently flawless arguments for *p* are evidence for *p*, evidence that there is an apparently flawless argument for *p* is higher-order evidence for *p*.

[6] I use the term "bracket" to refer to cases in which the presence of the evidence you should bracket does not make a difference to the doxastic attitude you should have. This is consistent with the second of Maria Lasonen-Aarnio's conditions in her "two-tiered theory of justification": "A doxastic state is epistemically rational only if one lacks evidence that it is flawed" (2014, 325). Lasonen-Aarnio introduces the second tier to accommodate intuitions in cases like Christensen's.

There are also cases in which you should bracket lower-order evidence. In Christensen's "Drugs" example, an experimental subject is asked to figure out the answer to a simple logic puzzle. The subject has unknowingly taken a drug that leads to severe impairment of ability to solve logic puzzles in 80% of subjects; they nearly always get the answers wrong. Though the subject is among the 20% unaffected, after the subject is informed that they have been given the drug, the subject's confidence in their answer should drop dramatically. According to Christensen, in this case the subject should bracket the lower-order evidence. For,

> in the case where I'm immune, it is not obvious why my total evidence, after I learn about the drug, does not support my original conclusion just as strongly as it did beforehand. . . . The undoubted facts support my answer in the strongest possible way—they entail my answer—, and this kind of connection cannot be affected by adding more evidence. (2010, 195)

In Christensen's case, the subject's higher-order evidence is in tension with the bracketed lower-order evidence. There are also cases in which you should bracket lower-order evidence that is in accord with higher-order evidence. Consider a variant of the Drugs case. You receive, unknowingly, a drug that has the following effect on 80% of people: it makes them upgrade the strength of good arguments a small but significant amount. Fairly strong arguments seem very strong. Very strong arguments seem decisive. You are, also unbeknownst to you, one of the 20% of people on whom the drug has no effect. You encounter a decisive argument for a surprising conclusion and come to have an appropriate degree of justification for that conclusion and proportion your confidence accordingly. What should your confidence be after you discover that you have been given the drug? As in Christensen's unmodified case, your confidence should decrease. It should decrease, but should still be high. You have higher-order evidence that your seemingly decisive argument is likely only very strong. Here, the higher-order evidence is in accord with your lower-order evidence. Both suggest the truth of the surprising conclusion. But you should nonetheless bracket the lower-order evidence.[8]

---

[7] After all, that you have perceptual evidence that there is a large circle on the page might be evidence I have that there is a large circle on the page, even if I lack perceptual evidence that there is a large circle on the page.

That there is an apparently flawless counterargument is higher-order evidence that the conclusion of the argument is true. But it can be much weaker evidence than the lower-order evidence provided by the counterargument itself. Considerations about your unreliability in evaluating the counterargument can, in such cases, have much the same impact as the drug: it can make it the case that you should bracket the lower-order evidence and allocate credence in line with the higher-order evidence. When you have reason to believe that you are unreliable in your evaluation of lower-order evidence, this higher-order evidence—that there is an apparently flawless counterargument—is not terribly strong. When such not-terribly-strong evidence comes up against the otherwise-knowledge-level lower-order evidence provided by your arguments, it can often lose.

---

[8] Some (for example, (White 2010, esp. 602–605)) worry that Christensen's verdict about these kinds of cases undervalues the fact that the subject has performed a competent deduction from the evidence. Surely, the thought is, the fact that the subject has performed a competent deduction from the evidence makes an epistemic difference. But there is room to deny Christensen's verdict in his specific case, while granting that you should sometimes bracket the lower-order evidence provided by apparently flawless counterarguments. For there is no presumption in these cases that drawing the conclusions of those arguments involves competent deduction. On the contrary, if you lack the training to reliably evaluate arguments of that sort, then your deduction of the conclusion from the premises will precisely not be competent.

One might think that symmetrical considerations should have you bracket, not only your opponent's apparently flawless argument, but also the otherwise knowledge-level lower-order evidence provided by your argument. If general considerations of fallibility require you to bracket your opponent's lower-order evidence that not-*p*, they should require you to bracket your lower-order evidence that *p*. If you do that, you'll be left with two competing bits of evidence: that there is an apparently flawless argument (yours) that *p*, and that there is an apparently flawless argument (your opponent's) that not-*p*: stalemate. The evidence for *p* is guaranteed to beat out the evidence for not-*p* only if the evidence for *p* is the strong lower-order evidence, while the evidence for not-*p* is the weak, higher-order evidence. So, why are there cases in which, though you should bracket your opponent's lower-order evidence that not-*p*, you need not bracket your lower-order evidence that *p*?

3.3 THE EPISTEMIC EFFICACY OF AMATEURISM

To favorably resolve the matter, there needs to be a symmetry breaker that requires you to bracket your opponent's lower-order evidence or argument while bracketing your higher-order evidence. One possible symmetry breaker involves not features of the evidence or argument itself, but features of the person or group putting forth or convinced by the evidence or argument.[9] For example, taking off from Ballantyne's "Big Smoke" example, if you know that those convinced by some argument that smoking is not a health hazard are also in the pay of Big Smoke, you might discount that argument as long as you know you don't have similar cognitive pressures in the other direction.

---

[9] Thanks to Nima Khodabandeh for making clear that this possibility needs some discussion.

If someone with a selfish interest in convincing you that *p* merely asserts that *p*, you may have reason in some cases to discount their testimony.[10] But in the relevant cases, you are faced not simply with a naked assertion, but with an argument. The quality of the argument, it might be thought, is independent of the selfish motivations of those putting forth the argument. Should you discount apparently flawless counterarguments merely because those putting forth and convinced by the arguments are known to be biased in some way? I incline toward thinking that at least in some cases you should. But I don't have a decisive argument to that effect, so my interest in the rest of this chapter is in the features of the evidence or argument itself that might license you in discounting it when it is apparently flawless.

The requisite symmetry-breaker is suggested by Christensen's Drugs example. In that case, according to Christensen, you should bracket the lower-order evidence because you have reason to believe you can't evaluate the evidence reliably. Perhaps in his specific case this verdict can be questioned because, in fact, you are competent when it comes to that evidence type; the drug doesn't work on you.[11] But if you not only have reason to believe you can't evaluate the evidence reliably, but you in fact aren't competent with that evidence type, then Christensen's verdict should be even more appealing. It's not general considerations of fallibility that warrant bracketing the lower-order evidence provided by your opponent's argument. It's that you tend to be less competent with respect to certain kinds of arguments than others; you tend to find certain kinds of arguments apparently flawless even if they're misleading.

Such arguments need not be fallacious. For example, suppose highly complicated arguments tend to be apparently flawless even if they are misleading: complexity can mask misleadingness.[12] It's not that complexity in an argument invalidates the argument. But it might make it that the support for the argument's conclusion provided by your total evidence after you are exposed to the argument is not much greater than the support provided by the fact that argument is apparently flawless.

---

[10] See (Ballantyne 2015a) for a discussion of the limits on this kind of "debunking reasoning."

[11] Again, see (White 2010, 602).

All that's left to do is figure out what features make arguments tend to be apparently flawless even if they are misleading. One possibility involves the contrast between theoretical vs. perceptual evidence. Knowledge might survive arguments that involve highly abstract principles or complicated mathematical formulae, but "seeing is believing." For example, your immediate perception of moving objects completely trumps any theoretical arguments that there is no motion. If we're looking for a substantive rule regarding what kinds of counterarguments knowledge can survive, we might start by placing perceptual arguments in a safer place than theoretical ones. Perhaps if you have perceptual knowledge, your knowledge can survive purely theoretical arguments to the contrary because it is less surprising for purely theoretical arguments to be apparently flawless, even if misleading, than it is for perceptual arguments to be apparently flawless even if misleading.

This is too hasty. First, a blanket prejudice against theoretical arguments smacks of George W. Bush's repeated charge in the first 2000 U.S. presidential debate that Al Gore's calculations about Bush's tax cut plan were just "Washington fuzzy math."[13] Second, it's not easy to separate theoretical from perceptual components of arguments. Magic shows seem to involve a tension between theoretical arguments about physical possibility and perceptual evidence that, say, a rabbit appeared out of nowhere—a tension in which the theoretical arguments win. But the perceptual considerations import theoretical ones: they aren't just reports of sense data. And the theoretical ones are based on a lifetime of perceptual training. The suggestion that perceptual considerations always trump theoretical ones is not promising.[14]

---

[12] Moti Mizrahi (2015) says that "[i]t appears that many philosophers are often bedazzled by complex arguments," (1) though he may not mean that they are bedazzled in the sense that complex arguments are often apparently flawless. He goes on to say that the bedazzlement often remains even when the philosophers admit that "the arguments are often not of the highest argumentative standard" (1).

[13] See transcript at http://www.debates.org/index.php?page=october–3–2000-transcript.

We need a different symmetry-breaker. A second possibility is that arguments with certain kinds of content are particularly tempting so that you're particularly apt to overlook flaws in those arguments even if they're there. For example, perhaps you're particularly apt to be fooled by arguments for racist or sexist conclusions or with premises that posit racial or sexual differences or which dehumanize some group of people.[15] If so, comparing tempting contemporary arguments with a certain content to past (debunked) arguments with that same content should be convincing, as it is in this argument from Josephine Ross:

> Although future generations were often offered as justification for repressive laws, this argument was a front to make those who wished to prevent interracial sex and marriage appear high-minded. Opponents of same-sex marriage rely on this same argument today. They claim that it is better to raise children with one mother and one father, and they use concern for children and future generations as the reason to deny marriage rights. (Ross 2002, 267 ff.)

---

[14] We might even think that in some domains—like basic arithmetic—theoretical or *a priori* evidence always trumps perceptual evidence. Mill (1850, Book III, Ch. III, §5) famously claims that knowledge of arithmetical truths derives from observations of aggregates of objects, a view about which Frege famously quipped "What a mercy, then, that not everything in the world is nailed down" (1953, 9). Those sympathetic to Frege might argue that, not only isn't observation required for knowledge of basic arithmetical truths, no experiential evidence could require us to rationally deny at least simple arithmetical truths. (Thanks to an anonymous referee for Oxford University Press for recommending I incorporate this example.) Such a stance is questioned by Quine (1980). Its denial seems to be the lesson of the philosophical literature on the relationship between the *a priori* and the empirical that was produced in the 1980s and 1990s (by, for example, Noah Lemos (1984), Al Casullo (1988), Donna Summerfield (1990), and Aron Edidin (1984)), the primary verdict of which is that the *a priori* admits of empirical revision. Though some—Hartry Field (1996) and Paul Boghossian (1997), for example—argue that, properly understood, *a priori* justification is not empirically revisable, it seems like a mistake in this context to rely overmuch on the unrevisability either of *a priori* or empirical justification.

[15] You might also tend to fear new technology, so end up unwarrantedly tempted by arguments that some new item of technology will be dangerous or should be banned. Thanks to Todd Dilling for suggesting this possibility.

Call this general argument strategy, in honor of Jeff Lockhart (2017), the “here-we-go-again” induction.[16] The idea is that, if arguments with certain kinds of content have been apparently flawless, only to be shown later to be misleading, then you shouldn’t take too seriously new arguments with that same kind of content.

While my hope is that some suitably restricted form of this induction can be made plausible (because who doesn’t want to dismiss *The Bell Curve* without having to read it?[17]), the worry is that it overgeneralizes and undercuts much that we want to claim to know. Many people, including me, who are convinced by Ross’s argument, will not be convinced by the here-we-go-again induction in other contexts:

> The same argument that was used to justify slavery is now used to justify abortion—slaves weren’t technically humans, so slavery was ok. Today some proponents of abortion assert that abortion is ok because fetuses aren’t technically human. Perhaps 100 years from now, we’ll see the evil in abortion that we now recognize in slavery.[18]

Likewise, if we adopt the here-we-go-again induction when it comes to arguments for racist conclusions, it’s going to be hard to avoid discounting the conclusions of scientific arguments in the face of the so-called “pessimistic induction.” The pessimistic induction—like the here-we-go-again induction—is an inductive argument: apparently flawless scientific arguments have often turned out misleading in the past, so we should be leery of conclusions of contemporary scientific arguments, as well. Of course, there are differences between the pessimistic induction and the here-we-go-again induction. The here-we-go-again induction discounts arguments because of their content. The pessimistic induction discounts arguments because of the methods used in those arguments. It’s not clear, however, that such differences favor the here-we-go-again induction over the pessimistic induction rather than the other way around. There may be ways that the pessimistic induction fails while the here-we-go-again induction—in some restricted form—succeeds. But for the time being the search continues.

---

[16] The expression “here we go again” occurs in the title of Lockhart’s critique of an op-ed piece advocating innate biological differences between men and women. Lockhart uses, at one point, a version of the here-we-go-again induction, favorably quoting the last two sentences of this longer selection from Beth Hess:

> It is no accident that theories of biological and psychological determinism, rooted as they are in reifying distinctions, capture the imagination of intellectual elites, even though each in turn has been discredited. For two millennia, “impartial experts” have given us such trenchant insights as the fact that women lack sufficient heat to boil the blood and purify the soul, that their heads are too small, their wombs too big, their hormones too debilitating, that they think with their hearts or the wrong side of the brain. The list is never-ending. . . . (1990, 81)

[17] In *The Bell Curve*, Richard Herrnstein and Charles Murray conclude, about racial differences in test IQ, “It seems highly likely to us that both genes and the environment have something to do with racial differences” (Herrnstein and Murray 1994, 311). For those hoping to deny that it is likely that genes have any part to play in the lower tested IQ scores of black test-takers (as opposed to whites), one must find some way to discount Herrnstein and Murray’s largely statistical argument. If you lack statistical acumen, one way to discount their argument is via an induction over past arguments that have purported to show such genetic differences and have since been debunked. That is the way suggested in this part of the body of the text. A second way—argued for in the rest of the chapter—is via your very lack of statistical acumen; if you lack statistical acumen, it will be unsurprising to find their argument apparently flawless even if it is misleading. This second strategy can allow you to discount their argument even on the condition that you read the book and find their argument apparently flawless.

[18] Posted by an anonymous contributor at https://answers.yahoo.com/question/index?qid=20070601105312AAPwdx4.

I’ll go back to square one. Here’s an incomplete list of arguments I wouldn’t be surprised to find apparently flawless even if they’re misleading: Bayesian arguments for intelligent design, history-based arguments that the Holocaust didn’t occur, statistical arguments about how many deaths are caused in the name of various religions, meta-analyses that purport to show that psychic phenomena exist, arguments that some university was wrong to revoke an honorary degree, strategic arguments that Israel is justified in retaliating in ways that risk civilian lives, legal arguments that restrictions on same-sex marriage are Constitutional. I know the conclusions of some of these arguments are false. For others I have confidence that the conclusions are false, though I feel that I fall short of knowledge. For others I lack confidence either way. But in all these cases, I wouldn’t be surprised if the arguments were apparently flawless even if misleading.

With respect to all these arguments, I lack some knowledge or expertise that would allow me to competently appraise the arguments. I am not adept in probability theory, Bayesian analysis, or the other technical skills required to grasp most recent and sophisticated arguments for intelligent design. I'm not a historian, so I don't know enough historical facts to adequately counter arguments for Holocaust denial. Statistics don't lie, but liars use statistics, and I lack sufficient statistical acumen to responsibly evaluate arguments about how many deaths are caused in the name of various religions. I am also, at best, an amateur when it comes to evaluating the kinds of studies for which meta-analyses are appropriate, so I'm the perfect dupe for meta-analyses purporting to show the existence of psychic phenomena. Objections to the revoking of honorary degrees are usually made on the basis of selected quotes and facts about the recipient's personal projects, many of which seem very worthy (hence, the initial granting of the honorary degree). But I'm not familiar enough with the overall work or life of most recipients to be aware of defeaters—if there are any—to those inevitably inductive cases. Nor am I familiar enough with the strategic, historical, or moral situation in Israel and the West Bank to evaluate arguments purporting to show that Israel is justified in responding militarily when missiles are launched against it. Finally, I'm not adequately versed in Constitutional Law to cite precedents that work against whatever legal argument is used to establish the Constitutionality of restrictions to same-sex marriage.

My limitations include limited factual knowledge and methodological acumen. Of course, if I have spent significant time with an argument, understood it, and have worked with the argument, I will hopefully have learned whatever facts are included in the premises of the argument itself—the activities of the recipient of the honorary degree, say, that *prima facie* warranted granting them an honorary degree in the first place. What I won't get from the argument is background knowledge that is relevant to evaluating the argument. If the argument for granting the honorary degree stems from evidence about things the recipient has said and work they've done, then for it to be surprising that I wouldn't be able to figure out where the argument goes wrong (if it does), I'd have to know a lot about things the recipient has said and work they've done. I'd have to know, for example, that the domain of facts suggested by the argument doesn't leave out any facts that run counter to the argument. I can't know this by evaluating the argument in isolation.

The same goes for arguments that rely on formal methods—like various statistical or Bayesian maneuvers. To reliably evaluate those arguments, I have to be familiar not only with the statistical or Bayesian maneuvers they use, but also with potential statistical or Bayesian maneuvers that would counter the argument. I have to be familiar enough with Bayesianism or statistical arguments generally to know that there are no such counter-maneuvers. Bayesian dilettante that I am, this is not something I am in a position to know.

Therefore, if some evidence type is essential to the soundness of an argument, in order for it to be surprising for me to find the argument apparently flawless even if misleading, I'd have to be well versed in the evidence type generally, not merely the specific subset of that evidence type that is presented in the argument. If I'm not well versed in the evidence type generally—if I lack the relevant knowledge base or methodological acumen appropriate to the evaluation of that evidence type—it will be unsurprising for the argument to be apparently flawless, even if it's misleading.

One consequence is that, the more kinds of evidence, methods, and formal principles an argument brings in to support its conclusion, the more background knowledge you need for it to be surprising that the argument is apparently flawless. This is not simply because, the more evidence types, methods, and formal principles are brought in, the more complicated the argument is. It's that the more evidence types, methods, and principles are brought in, the more likely it is that you don't have the knowledge or methodological acumen needed for it to be unsurprising that you evaluate the argument incorrectly. If you find it frustrating to argue with a person who is not convinced by your sound argument, despite the fact that they can't tell you what's wrong with it, they just may be responding correctly to an evidence type, method, or formal principle with which they are insufficiently familiar.

So, take this as our criterion of unsurprisingness:

> (Principle of Expertise) It is often sufficiently unsurprising that a misleading argument is apparently flawless if the argument invokes evidence types, methods, and principles with respect to which you lack relevant background knowledge or requisite methodological acumen.

The Principle of Expertise doesn't simply offer a condition for it to be somewhat unsurprising that a misleading argument is apparently flawless. It gives a condition for it to be "sufficiently" unsurprising. Sufficient for what? For knowledge to survive the counterargument being apparently flawless. We saw earlier how unsurprising it needs to be. It needs to be no less surprising that a misleading argument is apparently flawless than that an arbitrary argument with relevantly similar features would be apparently flawless. Barring that, your knowledge must be sufficiently over the threshold of probability that it can survive the reduction in probability necessitated by the relevant application of Bayes' Theorem.

If you lack the relevant background knowledge or the requisite methodological acumen to evaluate an argument—so that it is sufficiently unsurprising that an argument is apparently flawless even if it is misleading—then the argument, by itself, cannot put you into a position to know that its conclusion is true. That's because, if it would be unsurprising for a source to lead you astray, then you can't know that the source hasn't led you astray. This is a much weaker—thus, more plausible—principle than one often attributed to Descartes in his Third Meditation:

clear and distinct perception can't give you knowledge if it can ever lead you astray.[19] The principle here says only that clear and distinct perception—and any other source—can't give you knowledge if it wouldn't be surprising for it lead you astray.

Therefore, a positive test for whether you have sufficient expertise with an evidence type invoked by an argument (or other putative knowledge source) is whether that argument (or other putative knowledge source) alone is sufficient for knowledge. If it is, then you have the relevant background knowledge and requisite methodological acumen to have sufficient facility with the argument. Only if you can't get knowledge on the basis of an argument because you lack sufficient facility with the evidence types invoked in the argument does the Principle of Expertise entail that it is sufficiently unsurprising that a misleading argument is apparently flawless.

Is the Principle of Expertise true? Does satisfying the condition often result in an appropriate ratio of P(AF|M) to P(AF)? I think so. Lots of counterarguments lack obvious flaws and have the sophistical virtues required to make them seem compelling to those who offer them. These are, after all, arguments that are convincing to your opponents. Therefore, the kinds of completely obvious errors that plague some counterarguments will often not be present. The counterarguments will stand or fall entirely on the basis of successes or failures in the use of a kind of evidence with which you lack expertise. Nor are the successes or failures going to be elementary enough to be obvious to someone with only a layperson's expertise. In these kinds of cases, it seems just as likely for novices to be duped by misleading but clever arguments as it is for novices to accurately respond to non-misleading arguments.

---

[19] Descartes' explicit claim is about certainty, not knowledge: "In this first item of knowledge there is simply a clear and distinct perception of what I am asserting; this would not be enough to make me certain of the truth of the matter if it could ever turn out that something which I perceived with such clarity and distinctness was false" (1641/1985, 24).

Amateurs, no less than experts, are subject to confirmation bias. Amateurs will tend to find non-existent flaws in sound and misleading arguments alike. Perhaps this means that, for amateurs (and perhaps for experts), the prior probability that some counterargument is apparently flawless is low. The prior probability that a counterargument is apparently flawless is a function of the probability that the argument is apparently flawless conditional on its being misleading, and the probability that the argument is apparently flawless condition on its being non-misleading. If the prevalence of confirmation bias makes it improbable that both misleading and non-misleading counterarguments will be apparently flawless, the resultant probability that a counterargument is apparently flawless will be similarly low.[20]

None of this creates difficulties for the argument here—that misleading arguments can be just as likely to be apparently flawless as non-misleading arguments. All it does is lower the prior probability that the counterargument is apparently flawless. This might seem to make it easier for the relevant ratio to be at least 1, since the denominator of that ratio—$P(AF)$—is low. But the factor that makes the denominator low also makes the numerator—$P(AF|M)$ low. Still, in many cases, the relevant ratio can be at least, and sometimes greater than, 1.[21] Even when not, the relevant ratio need not be so far below 1 that the difference can't be made up by a degree of knowledge-level probability significantly above the threshold.

---

[20] Again, thanks to Catarina Dutilh Novaes for discussion on these issues.

[21] The argument in this section is an argument that, for amateurs confronted with relevant counterarguments, the probability that the argument is apparently flawless conditional on it being misleading is no lower than (or not much lower than) the probability that an argument is apparently flawless conditional on it being non-misleading. This conclusion can be depicted formally as

1) $P(AF|M)=P(AF|not\text{-}M)$

The Principle of Expertise requires a different equality:

2) $P(AF|M)=P(AF)$

Here's the proof that (1) entails (2). The Total Probability Theorem tells us that

3) $P(AF)=P(M)P(AF|M)+P(not\text{-}M)P(AF|not\text{-}M)$.

Substituting $P(AF|M)$ for $P(AF|not\text{-}M)$ (in accord with (1)) yields

4) $P(AF)=P(M)P(AF|M)+P(not\text{-}M)P(AF|M)$,

which is equivalent to

5) $P(AF)=(P(AF|M)(P(M)+P(not\text{-}M))$.

Because $P(M)$ and $P(not\text{-}M)$ sum to 1, (2) is equivalent to (5), which is entailed by (1). Therefore, (2) is entailed by (1).

We now have the resources to argue for a principle governing how knowledge can survive when confronted with an apparently flawless counterargument:

> P1) If it is sufficiently unsurprising that a counterargument would be apparently flawless if it were misleading, then knowledge can survive the counterargument being apparently flawless. (from the Inertness of Unsurprising Error)
>
> P2) It is often sufficiently unsurprising that a misleading argument is apparently flawless if the argument invokes evidence types, methods, and principles with respect to which you lack relevant background knowledge or requisite methodological acumen. (from the Principle of Expertise)

Therefore,

> C) Often, if a counterargument invokes evidence types, methods, and principles with respect to which you lack relevant background knowledge or requisite methodological acumen, then knowledge can survive the counterargument being apparently flawless.

This conclusion hearkens back to Tom Kelly's observation that played a role in Chapter 2: that if the best explanation for your failure to expose a flaw in a counterargument is your "own cognitive limitations, then [you] should remain unmoved in the face of the argument" (2005, 5), though the conclusion here is about knowledge rather than justified belief. Call the conclusion the *Novice Knowledge Principle*.

The Novice Knowledge Principle does not say that whenever you lack sufficient expertise, then knowledge survives the counterargument being apparently flawless. It says only that often if you lack sufficient expertise, knowledge can survive. There might be other barriers to the survival of knowledge, other than the counterargument being apparently flawless.[22] We've already seen that knowledge that *p* only survives in the relevant cases if you should bracket the lower-order evidence against *p* while you shouldn't bracket the lower-order evidence for *p*. In addition, perhaps epistemically vicious dismissal of a counterargument can destroy knowledge even if, given your level of expertise, it wouldn't be surprising for the counterargument to be apparently flawless even if it is misleading. You might dismiss the counterargument for the wrong reasons—because you were moved by problematic affective factors or because you didn't adhere to various procedural norms.[23] If dismissing a counterargument for the wrong reasons is incompatible with knowing that the counterargument is misleading then, though the Novice Knowledge Principle rightly entails that your knowledge could have survived (had you dismissed the argument for the right reasons), your knowledge might not in fact survive. We shall also see, below, that the mere fact that there is a relevant counterargument—apparently flawless or not—can provide knowledge-destroying evidence even if the fact that the argument is apparently flawless leaves knowledge intact.

---

[22] Thanks to Todd Dilling for making me see that this clarification is necessary.

According to the Novice Knowledge Principle, it can be easier in one sense to retain knowledge if you lack the relevant expertise than if you have it. Perhaps this is another way to construe Catherine Elgin's (1988) "epistemic efficacy of stupidity." It's not exactly what Elgin meant by the expression; Elgin meant (among other things) that those who are less able to distinguish fine categories are also less apt to fall into error by putting a member of one category into the wrong nearby category. Here it means that because amateurs should be less surprised when an argument is apparently flawless even if it's misleading, amateurs might more easily keep knowledge they already had.

---

[23] The arguments in this chapter and Chapter 6 can be taken to show that the following isn't one of the relevant procedural norms in all contexts in which you know that a counterargument is misleading: take significant time to evaluate the counterargument. For, if you know that the counterargument is misleading, then you know that taking significant time to evaluate the counterargument will either reveal where the argument's misleadingness lies (and so will re-confirm what you already know) or will fail to reveal where the argument's misleadingness lies (and so might incorrectly indicate that the argument is not misleading).

This might seem problematic. But it shouldn't. First, even if it's in one sense easier for amateurs to retain knowledge, it might be harder for them to have knowledge in the first place (at least of controversial matters). Second, experts in an area, though they should be more moved by apparently flawless misleading counterarguments than amateurs are, won't find nearly as many misleading counterarguments apparently flawless. So it's not like there is a whole range of cases in which experts should be moved by misleading arguments that amateurs shouldn't be moved by.

But, third, it's the correct result. Experts should be moved by apparently flawless arguments more than novices are moved by apparently flawless arguments. The novice has a fallback excuse for not being moved by the apparently flawless argument: "I wouldn't be able to tell that this argument is misleading even if it is." The expert has no such fallback; if there were an obvious objection or response to the argument they, unlike the novice, would know about it. Therefore, the expert would—much more likely than the novice—be able to tell that the argument were misleading if it were. But they can't. This is why, from a third person perspective, we should take more seriously arguments in which an expert can't expose flaws than in which a novice can't expose flaws. If such arguments should be more convincing to those observing the expert's failure, they should be more convincing to the expert themselves.

Here's a potential counterexample to the Novice Knowledge Principle, inspired by a case from Todd Dilling's unpublished (2017):

> THE CASE OF THE MISLEADING X-RAY
> You are injured on the job and a surgeon diagnoses you as having a stress fracture in the scaphoid bone in your wrist that requires expensive and temporarily incapacitating surgery. This testimonial evidence is good enough for you to know that you have a fractured scaphoid. Your employer requires you to get a second opinion from the company doctor. While waiting in the company doctor's office for them to tell you their opinion, you notice your X-Rays sitting on their desk. Curious, you look them over. Your scaphoid is circled and accompanied by a little note: "No visible fracture." Indeed, you

> see no discoloration or line that might indicate a break. You have no special facility in evaluating X-Rays.

It might seem that the Novice Knowledge Principle allows you to retain your knowledge that you have a fractured scaphoid. You started out knowing that you have a fractured scaphoid. The X-Ray provides you with an apparently flawless argument that you don't have a fractured scaphoid. You find it compelling that there is no visible fracture on the X-Ray. You find it compelling that if there is no visible fracture on the X-Ray, then you don't have a fractured scaphoid. But you lack expertise to reliably evaluate X-Ray results; you don't know about the sorts of considerations that might override the inference from lack of visible fractures to the absence of a fracture in the bone. The evidence that grounds your knowledge—the testimony of your surgeon—is of a different sort from the evidence invoked by the apparently flawless counterargument. Therefore, by the Novice Knowledge Principle, your knowledge survives the X-Ray evidence.

Intuitively, this is the wrong result. The X-Ray evidence should reduce your confidence below what is required to know that you have a fractured scaphoid.[24] If you are called away from the office prior to hearing the official opinion of the company doctor, then until you return, you do not know that you have a fractured scaphoid. If your friend asks, "Well? Do you have a fractured scaphoid?" to answer truly you have to say, "I guess I don't really know." This isn't because of the company doctor's testimony; you weren't there long enough to hear that testimony. The only testimony you had from the company doctor was that there was no visible fracture on the X-Ray. You could see that for yourself.

---

[24] After all, you shouldn't just act as if you have a fractured scaphoid and schedule your expensive and temporarily incapacitating surgery. Thanks again to Todd Dilling for noting this.

The Novice Knowledge Principle does not entail that you retain your knowledge after seeing the X-Ray. Again, the Novice Knowledge Principle provides a sufficient condition for when knowledge can survive apparently flawless counterarguments. It doesn't provide a sufficient condition for when knowledge does survive. One factor that can get in the way is the mere fact that there is a would-be[25] relevant counterargument—apparently flawless or not.[26] The fact that there is X-Ray evidence that would convince lots of people that you don't have a fractured scaphoid is surprising, given your knowledge that you have a fractured scaphoid; conditional on your having a fractured scaphoid, the existence of a relevant counterargument is improbable. Therefore, the existence of a relevant counterargument is evidence that you don't have a fractured scaphoid, and it's evidence that you shouldn't bracket. It can require you to reduce your confidence and, if that reduction is sufficient, destroy your knowledge.

The situation is different when it comes to controversial propositions you know. When it comes to controversial propositions, it is unsurprising to find an argument that would convince lots of people. Therefore, you should expect the existence of relevant counterarguments. Because you know that the conclusions of those counterarguments are false, their existence is unsurprising even conditional on their being misleading. Therefore, their existence is not good evidence of the truth of their conclusions and, so, their existence does not require significant confidence-reduction. Your knowledge can survive unscathed.

It might seem like the case of the misleading X-Ray is closer to the case of controversial belief. The second opinion is being provided by the company doctor. Therefore, you might suspect, they'll be doing their best to provide an argument that suggests you don't have a fractured scaphoid. It might be certain for you going into the office that they will provide—sound or no—an argument designed to fool a layperson like you into believing that you don't have a fractured scaphoid. If so, then it is unsurprising that there is a would-be relevant counterargument, even if the counterargument is misleading. Therefore, the existence of the would-be relevant counterargument doesn't significantly confirm that you don't have a fractured scaphoid and you shouldn't reduce your confidence in a way incompatible with knowing that you have a fractured scaphoid.

---

[25] It's a "would-be" relevant counterargument, not a straight-up relevant counterargument, because it doesn't convince lots of people. It only would convince lots of people if they were to encounter it.

[26] There are various ways for this fact—that there is a would-be relevant counterargument—to destroy knowledge. One way—presented in the body of the text—is for that fact to count as countervening evidence. A second way is for that fact to make it that you shouldn't bracket the lower-level evidence presented by the counterargument. Perhaps you should bracket lower-level evidence only when it is sufficiently unsurprising for there to be a would-be relevant counterargument. I am neutral between these two mechanisms.

This is the correct result. If you should be certain that you will find a relevant counterargument prior to encountering it, then finding such an argument should not reduce your confidence that you have a fractured scaphoid. Your confidence should be reduced only to the extent that you'd be surprised to find a relevant counterargument conditional on the counterargument being misleading. If your knowledge is justified to a sufficient degree above the threshold for knowledge, your knowledge can survive non-zero but minimal reductions in justified confidence. When it comes to controversial propositions, you should be significantly unsurprised to encounter relevant counterarguments. So the mere fact that there is a relevant counterargument is a minimal confidence-reducer.

A second consequence of the Novice Knowledge Principle is that evidence of an unfamiliar form has automatically diminished probative force. Suppose that for the first time you are taking a class in advanced mathematics and your instructor presents you a novel proof for a difficult theorem, using inference rules that—while you've gotten to the point where you understand them—you haven't gotten facility with yet. Your position with respect to that kind of evidence means that it might be unsurprising for you to find the argument apparently flawless even if it is misleading. Therefore, you should not be as convinced as you would be by the proof were you to have facility in the general class of rules that constitute that kind of evidence generally.

Again, this is not problematic. When you take a class in some newish subject matter, arguments presented in class in terms of that subject matter should have less confirming power than they otherwise would, even if those arguments seem completely convincing to you on first blush. This doesn't mean that you should be leery of the various conclusions. If the mathematician professor tells you that the argument is sound and that the conclusion is true, you should still believe and can know that the argument is sound and the conclusion is true. But the knowledge-providing evidence is not the lower-order proof, nor the higher-order evidence that you find the proof apparently flawless. The knowledge-providing evidence is the sincere testimony of the mathematician professor, perhaps corroborated by a textbook. This kind of evidence—testimonial evidence provided by an expert in a subject about which there is little disagreement about basic proofs—is a kind of evidence with which you are familiar and do have expertise, just in virtue of living in a world of testifiers.[27]

Not everyone has the expertise required to reliably distinguish trustworthy from untrustworthy testifiers, at least not in all domains, and the internet may be making it worse. According to a 2016 report of a study on how well students from middle school to college do at distinguishing fake from real news stories,

> we would hope that middle school students could distinguish an ad from a news story. By high school, we would hope that students reading about gun laws would notice that a chart came from a gun owner's political action committee. And, in 2016, we would hope college students, who spend hours each day online, would look beyond a .org URL and ask who's behind a site that presents only one side of a contentious issue. But in every case and at every level, we were taken aback by students' lack of preparation. (Stanford History Education Group 2016, 4)

---

[27] Josh Stein points out that higher-order evidence can increase the evidential value of lower-order evidence.

The degradation of the ability to reliably distinguish trustworthy from untrustworthy testifiers might seem to present a worry for the Novice Knowledge Principle. If you can't reliably evaluate testimonial evidence that, say, there is human-caused global warming then the Novice Knowledge Principle might seem to incorrectly imply that you can retain knowledge—or, at least, knowledge-level justification—that there is no human-caused global warming despite the testimonial consensus to the contrary. But the Novice Knowledge Principle only has this consequence if 1) your lack of facility with testimonial evidence doesn't result from problematic affective factors or failures to adhere to relevant procedural norms and 2) you can have knowledge-level justification for controversial propositions without the benefit of testimonial evidence. Arguably, in many cases, at least one of these conditions isn't satisfied.[28]

One might think that whenever you lack the expertise to evaluate a counterargument, then you can't have otherwise knowledge-level support and you'd better just withhold judgment. This is perhaps in keeping with Clifford's pronouncement that if you don't have time "for the long course of study which would be necessary to make [you] in any degree a competent judge of certain questions, or even able to understand the nature of the arguments," then you "should have no time to believe" (1886, 347). Distorting a bit, "If you don't have the expertise to consider the objection, you don't have the expertise to believe."

The Cliffordian conclusion is false. You don't just have the expertise to believe; as argued in Chapter 2, you can have the expertise to *know*. The considerations in this chapter give us an explanation why. Your lacking the expertise to evaluate a counterargument doesn't need to impact the strength or nature of your support for what you believe. Your belief might be based in a kind of evidence with which you have the relevant expertise. For example, the experts on some topic might have evidence that you lack the expertise to adequately evaluate. But if the experts agree, you should conform your belief to their consensus—say, on the existence and general causes of global warming. Even if you lack the expertise to evaluate global warming data, you have the expertise to evaluate the testimonial argument.

---

[28] You also can be epistemically criticizable for your lack of facility with testimonial evidence.

If an iconoclast presents a scientific argument in favor of global warming denial, you may lack the expertise to adequately evaluate their lower-order case. But you don't have to bracket your lower-order evidence, which does not employ the kind of evidence they employ in their lower-order case, but rather employs lower-order testimonial evidence with which you do have expertise. Again, a positive test for whether you have sufficient facility with an evidence type or knowledge source is whether that source, alone, could provide you with knowledge. Since scientific testimony—barring defeaters—can provide you with knowledge, you have sufficient facility with that source.

The Novice Knowledge Principle also allows knowledge to survive many arguments for conspiracy theories. In the case of conspiracy theories, sources of knowledge that there is no conspiracy behind some event often include two primary ones: consensus of recognized experts and the repeated and manifest failures of pro-conspiratorial arguments in that domain. Taken together or separately, these two sources of knowledge often resist counterarguments, when they do come along, whose errors you can't expose. Daniel Pipes[29] (1997) includes reliance on "obscurity" of evidence (40) and "overabundant learned factoids and pedantic references" (41), as well as “reluctance to divulge information” (40) in his long checklist of features that arguments for conspiracy theories tend to manifest when they’re not to be trusted. All of these features make it the case that you are less likely to have facility in the evidence types and methods used in the argument for the conspiracy theory. Though Pipes doesn’t give an argument that these features are problematic, the Novice Knowledge Principle explains why they would be.

---

[29] Pipes has published and been reported as expressing some problematic attitudes about Muslims—for example that "The Palestinians are a miserable people and they deserve to be" (Kelley 2001, 57). (Pipes denies that he ever said such a thing. But he does not deny saying that "Islamists constitute a small but significant minority of Muslims, perhaps 10 to 15 per cent of the population. Many of them are peaceable in appearance, but they all must be considered potential killers" (Pipes 2001, 22). Nor does he deny the percentages, instead qualifying the remark by saying, "I now find that 'these words ring archaic' because I now recognize that some Islamists are not potential killers" (Pipes 2014).) Still, he may have gotten right some of the reasons why some arguments for conspiracy theories shouldn't move us.

The Novice Knowledge Principle does *not* allow your knowledge to survive apparently flawless counterarguments when your knowledge is based on a principle like “Absence of evidence is evidence of absence.” Suppose that lack of evidence that a proposition is true constitutes evidence that the conclusion is false. For example, the lack of evidence for the existence of a mammal with 37 legs is evidence that there is no such mammal. Even if such a basis is sufficient for knowledge, if your knowledge is grounded only in this way and you are confronted by an apparently flawless argument that there is such an animal (for example, perhaps you see one), you lose your knowledge that there is no such animal.

The reason why you lose your knowledge is that the evidence that grounds your knowledge is the lack of evidence for the alternative. Your evidence is, in part, that there is no visual evidence for the existence of the animal. This evidence confers knowledge only if you have sufficient facility with the evaluation of visual evidence. But evaluation of the counterargument—the visual experience of the animal—requires just the same kind of facility. Therefore, you have no more facility in the methods used to argue for your conclusion than you do in the methods used in the counterargument. So the Novice Knowledge Principle can’t be used to favor your conclusion at the expense of the conclusion of the counterargument.

A further limitation of the Novice Knowledge Principle is that it does not explain why you get to dismiss trick proofs that 1=0. These proofs get their force precisely from the fact that they seem to rely on an apparatus with which most of us have facility—elementary algebra. Likewise, experts can find themselves stumped in a range of cases by various arguments: it's certainly possible for the world's greatest mathematician to happen upon a novel and astounding argument that 1=0, or for the world's greatest physicist to construct an argument that there is no motion in which even they can't identify a flaw. In neither case need the expert reduce their confidence. The Novice Knowledge Principle does not explain why they don't need to.

There may be other principles that allow us to conclude, in these cases, that it is unsurprising even for the experts that the argument is apparently flawless even if it is misleading.[30] Perhaps trick arguments generally have tricky trappings. But it might also be that, in the end, in many of these cases it is surprising for the experts that the argument is apparently flawless if misleading. It doesn't follow that confidence revision is required or that knowledge is destroyed. The Novice Knowledge Principle expresses what is often a sufficient condition for the survival of knowledge, not a necessary one. There may be other ways for your total evidence to remain knowledge-level even if there is an apparently flawless counterargument.

I've been considering cases in which your positive support for your position involves an evidence type with which you have sufficient expertise, so that you should bracket your higher-order evidence while bracketing the lower-order evidence provided by your opponent's counterargument. This isn’t to say that you always should bracket your higher-order evidence for *p* and the lower-order evidence provided by your opponent’s apparently flawless argument for not-*p.* Sometimes you should bracket your lower-order evidence for *p*.

---

[30] Arguments might tend to be apparently flawless, for example, if they support your preconceived opinions, though this will less likely be operant when it comes to counterarguments, and does not provide a sufficient condition for bracketing the evidence provided by such arguments, because you're often able to reflect critically on arguments that support your views. The same goes for some other possible criteria—for example, that the counterargument is for an exciting and novel conclusion (again, think Hume's "passion of surprize and wonder"). What's distinctive about the criterion on offer in the body of the text is that you can't learn to better evaluate arguments that invoke evidence with which you lack the sufficient expertise without acquiring the requisite expertise, in which case the criterion no longer applies.

This can happen especially when you are first delving into a new area of research. If you’re like me, you’ve had the experience of being overly excited by some argument you’ve come up with when first thinking about a subject. Particularly prone to these experiences are newcomers to University, who have yet to learn the bitterness of the decisive counterexample. Once you become jaded, you learn to mistrust initial stabs at argument, even if you have spent some time with them and they’re apparently flawless. You think, “This argument looks pretty good, and I hope it’s sound. It *seems* sound, and I’ve checked it over. But I’m just not quite there yet. I don’t know what more I need, but so far all I can conclude is that, well, I’ve got an argument whose steps seem compelling and in which I cannot find a flaw. No more than that.”

The Novice Knowledge Principle explains why the experienced have this reaction. You mistrust initial forays that reveal apparently flawless arguments because you *should*. You recognize that your initial forays are often marked by lack of expertise with the kind of knowledge base and intellectual skills needed to adequately evaluate your arguments. Therefore, you recognize that you should bracket the lower-order evidence provided by those forays until you have the expertise you need to competently evaluate your argument.

3.4 THE NOVICE KNOWLEDGE PRINCIPLE AND THE CASE FOR THE OTHER SIDE

The Novice Knowledge Principle offers a new explanation for why it is important in academic papers and in arguments generally to present a sympathetic case for the opposition’s side. One natural reason for thinking it’s important to present a sympathetic case for the opposition is purely rhetorical: in many contexts you stand a better chance at convincing others if you demonstrate that you are sympathetic enough to their view that you understand where they are coming from. You also can stand a better chance at convincing others if you respond to any objections they might have.

A second natural reason for thinking it's important to present a sympathetic case for the opposition is more self-regarding. Writing does more than present in clear form what we already thought or believed. It doesn't just reveal our cognitive structures; it improves them.[31] The writing process helps us learn what to think and what is problematic about our thinking[32] and how best to revise and extend[33] our thinking. One way to achieve these goals and to make sure that you are on the right side of the issue is to treat the case for the opposition fairly, and ensure that you have responses to the most devastating objections for the other side. To return to a passage I quoted from Mill in Chapter 2,

> He who knows only his own side of the case, knows little of that. His reasons may be good, and no one may have been able to refute them. But if he is equally unable to refute the reasons on the opposite side; if he does not so much as know what they are, he has no ground for preferring either opinion. The rational position for him would be suspension of judgment, and unless he contents himself with that, he is either led by authority, or adopts, like the generality of the world, the side to which he feels most inclination. (Mill 1859/2003, 104)

We've seen that Mill's conclusion in full generality is false: you don't need to be able to refute the reasons on the opposite side in order to retain knowledge of your view. But you may often need to have refuted some of the most compelling reasons on the opposite side in order to retain knowledge of and justified belief in your position. One way to ensure you can do this is by writing it down.

---

[31] See (Menary 2007).
[32] See (Stoehr 1967, esp. 419–421).
[33] See (Oatley and Djikic 2008).

A third natural reason for emphasizing a sympathetic case for the opposition is that by doing so you assure readers unfamiliar with the issue you're discussing that the topic and your thesis are controversial. It explains to the reader why your thesis matters and why the issue is academically important—because there are reasons to deny what you're saying and that, barring the argument you're about to present, a reader might expect that your thesis is false. This is why presentation of the case for the other side can serve the job of what Gordon Harvey, the former Associate Director of the Harvard Expository Writing Program, has called the "motive":

> the intellectual context that you establish for your topic and thesis at the start of your essay, in order to suggest why someone, besides your instructor, might want to read an essay on this topic or need to hear your particular thesis argued—why your thesis isn't just obvious to all, why other people might hold other theses (that you think are wrong). (Harvey 1996)

The Novice Knowledge Principle provides a fourth reason for presenting a quite sympathetic case for the other side. It's not simply to answer objections that a hostile reader might have in mind; it's to train the reader to have facility with the kind of evidence you're using in your argument. If the reader comes away from your paper feeling like your argument—as it stands—could be convincing but that they don't have the resources they would need to adequately evaluate it, the reader will rightly be unsurprised that the argument is apparently flawless even if it's misleading. Thus, they will rightly bracket the lower-order evidence provided by your argument and won't be convinced to the degree your argument might warrant. A sufficient case for the other side will give the reader sufficient background knowledge and sufficient facility in the skills needed to evaluate the kind of evidence you're offering. The best cases for the other side make it surprising that your argument would be apparently flawless unless it were sound.

This is not an easy skill to master because it's very hard to give a novice reader a justified (and true) feeling that they have sufficiently mastered the skills and background knowledge needed to adequately evaluate the kind of evidence you invoke in your positive argument.[34] The only way to do it is by honestly presenting devastating objection after devastating objection to your argument, so that the reader gets background knowledge of the most dangerous objections out there and, what is even more difficult, gets a justified belief that the objections you consider are sufficiently exhaustive of the best objections.[35] The overall case for the other side consists of convincing cases against a paper's thesis—general reasons for thinking that the thesis is false—and more focused compelling objections against specific steps in the argument for that thesis. What the most talented practitioners of the craft seem to have in common is tenacity: the trait of never giving up on their opponent's argument; just when their opponent's case seems refuted, they consider another version of the case that withstands the refutation.[36] Just when you think

---

[34] In the Translator's Introduction to Gilles Deleuze's (2008), Hugh Tomlinson and Barbara Habberjam note that, though Deleuze's strategy is to show

> how the problems which arise in each of the first two Critiques, problems which are often seen as decisive objections to the Kantian philosophy, are recognized by Kant and dealt with in the third Critique. . . . [I]t is also remarkable, at first sight, that such a work should be written by, of all people, Gilles Deleuze. It is difficult to think of two philosophers more apparently opposite. . . (xiii)

As they quote Deleuze as remarking, "I wrote [the book on Kant] as a book on an enemy, in it I was trying to show how he works, what his mechanisms are. . ." (xiii) Thanks to Eric Walker for bringing this book to my attention.

[35] Indeed, William Lycan (2009) has published an entire paper consisting of nothing but the case for the other side. He notes that though "I have been a materialist about the mind for forty years" and that "[m]y materialism has never wavered. Nor is it about to waver now" and that "I have no sympathy with any dualist view, and never will," his "purpose in this paper is to hold my own feet to the fire and admit that I do not proportion my belief to the evidence. The dualism I shall defend is Cartesian, 'substance' dualism" (551). He calls the paper "an uncharacteristic exercise in intellectual honesty" (551). But if the argument here is correct, there is a further use for the paper; the more you train your readers in the case for the other side, the more surprising it will be for them once you publish an argument for your position that they find apparently flawless. Thanks to Andrew Cullison for directing me to this paper.

[36] For example, in G. Thomas Couser's (2001) defense of Oliver Sacks against an array of critics, he first spends some time responding to what he calls Alexander Cockburn's "hit-and-run attack" (2) before saying, "And yet. And yet there are some troubling aspects to Sacks's work. One could argue that Sacks's representation of his subjects is all the more invasive because his commentary establishes difference where the eye cannot easily detect it" (4). Only then does Couser move on to increasingly more sophisticated sorts of concerns.

Couser's "And yet. And yet . . ." is particularly illustrative of the sort of persistence required in a thoroughgoing case for the other side, but in general it's difficult to find single quotes from individual papers that adequately represent the degree to which an author expresses the case for the other side. Part of what it is to be adept at presenting the case is exhaustivity—rejecting a facile version of the case for the other side, fixing the case so it avoids the facile rejection, dealing with the revamped case, presenting an even more sophisticated version, etc. For other particularly skilled examples of the case for the other side, see (Carroll 2001), (Thomson 1994), and (Schroeder 2008). Thanks to Landon Schurtz, Robert Streiffer, and an anonymous commenter on the philosophy blog Daily Nous for these suggestions. The downside of the kind of tenacity necessary for an exhaustive case for the other side is that the resulting texts can be extremely long. See, for example, Dorothy Edgington's ninety-five page they've established their view, they consider an even more difficult obstacle the view needs to overcome.

Anyone who adequately spells out the case for the other side ends up ensuring that their readers can come to know that what they are arguing for is true. This may, of course, be why proficiency at developing the case for the other side ends up so rhetorically effective. If you leave your reader unconvinced that your argument would be apparently flawless even if it were misleading, then you can make your argument as apparently flawless as possible, and it won't do much good. If you can convince a reader that your argument wouldn't be apparently flawless unless it were sound, and you can also make your argument apparently flawless, then your reader will likely be more convinced by your argument.

Still, the rhetorical benefit is distinct from the epistemic one. You don't just want your readers to believe your conclusions. You want your readers to come to be justified in believing them and even know them on the basis of your arguments. You may even have moral and social obligations to ensure that your reader has the resources needed to become a knower on the basis your argument.[37] You can't do this unless you ensure that they aren't justified in saying to themselves, "I just don't know enough about the issue to trust my reactions to what, admittedly, seems to be a compelling argument." This is so whether readers start out with otherwise knowledge-level evidence that your thesis is false, or don't have pretheoretic opinions either way.

3.5 CONCLUSION

---

classic (1995). She is similarly persistent in chasing down possible options for the other side in her somewhat shorter (1986).

[37] Thanks to Josh Stein for pointing this out.

In this chapter I have focused on what kinds of arguments knowledge can survive. The answer is that your knowledge can survive arguments you wouldn’t be surprised were apparently flawless even if they were misleading. This is because, whether or not you know that the conclusions of those arguments are false, you shouldn’t give much credence to the conclusions of arguments you wouldn’t be surprised were apparently flawless even if they were misleading.

Because a counterargument being apparently flawless more often is compatible with knowing that the counterargument is misleading for novices than for experts, it might seem that novices can more easily be rightly closed-minded than can experts. This result certainly holds after both novices and experts have spent significant time with a counterargument and found it apparently flawless. For closed-mindedness with respect to a counterargument is, in part, the unwillingness to reduce your confidence in response to that counterargument even if you find the counterargument apparently flawless. When a counterargument is apparently flawless for a novice, they can more often than experts continue to know that the counterargument is misleading. Therefore—as I argue in more detail in Chapter 6—they can, more often than experts, remain rightly unwilling to reduce their confidence in response to apparently flawless counterarguments. Experts, who more often will lose knowledge that apparently flawless counterarguments are misleading will in those cases wrongly be unwilling to lower their confidence in response to the counterargument. Novices should tend to be more closed-minded with respect to apparently flawless counterarguments than experts are.

It might seem that such a result also holds prior to the expert or the novice spending significant time with the counterargument. Because the novice can more reasonably expect to retain knowledge that the counterargument is misleading if it is apparently flawless, they should even now—prior to examining the argument—be less willing than the expert to reduce their confidence conditional on it being apparently flawless. For the novice knows that if it is apparently flawless, this will more likely be because they lack the skills and background knowledge to adequately evaluate the counterargument. On the other hand, the expert knows that if the counterargument ends up being apparently flawless, this will not likely be because they lack the skills and background knowledge to adequately evaluate the counterargument; it will more likely be because the counterargument really is misleading. Therefore, one might think, experts should be willing, prior to examining the counterargument, to reduce their confidence conditional on the argument being apparently flawless.

This second conclusion should be resisted. It is not true that, prior to examining the counterargument, the expert should tend to be more open-minded than the novice. It is true that, once the counterargument is examined and it is apparently flawless, the expert can't as easily chalk this fact up to lack of skill and background knowledge. But prior to examination of the argument, the expert *can* know that the counterargument is misleading precisely because it would be—unlike for the novice—very surprising for the counterargument to be apparently flawless. For the expert, there couldn't easily be apparently flawless relevant counterarguments to what they know. Therefore, prior to examining the counterargument, the expert can know that the counterargument is misleading. Therefore, the expert also knows, prior to examining the counterargument, that if after examination, the counterargument is apparently flawless, it will have to be because of some error or incapacity of theirs. It can't be the result of the soundness of the counterargument because, they know, the argument is misleading.

Therefore, even though once a counterargument is apparently flawless, the expert can more easily lose knowledge that the counterargument is misleading, it's not true that prior to examining the counterargument, the expert more easily lacks knowledge that the counterargument is misleading. The result is that, prior to examination, experts need not be willing to lower their confidence in response to the counterargument, conditional on the counterargument being apparently flawless. Experts in that situation—as well as knowing novices—should be closed-minded with respect to the counterargument. Or so I argue in Chapter 6.

4

Psychic Phenomena and the Existence of God

> I'm just a caveman. I fell on some ice and was later thawed by some of your scientists. Your world frightens and confuses me! Sometimes when I fly to Europe on the Concorde, I wonder, am I inside some sort of giant bird? Am I gonna be digested? I don't know, because I'm a caveman, and that's the way I think! . . . But whatever world you're from, I do know one thing—in the 20 years from March 22nd, 1972, when he first ordered that extra nicotine be put into his product, until February 25th, 1992, when he issued an inter-office memorandum stopping the addition of that nicotine, my client was legally insane. (Phil Hartman as Unfrozen Caveman Lawyer)[1]

4.1 SOURCES OF KNOWLEDGE

There are many sources of knowledge: perceptual, testimonial, intuitive, mnemonic, inductive, deductive, and statistical. Each general source can be specified in more detail. Perceptual sources can be visual or auditory. Testimonial sources can be casual, formal, spoken, or written. Particular instances of knowledge based in those sources can rely on the deliverances of instruments—like hearing aids—, or on certain kinds of training—as when viewing autostereograms:

---

[1] See http://snltranscripts.jt.org/95/95pcavemanlawyer.phtml.

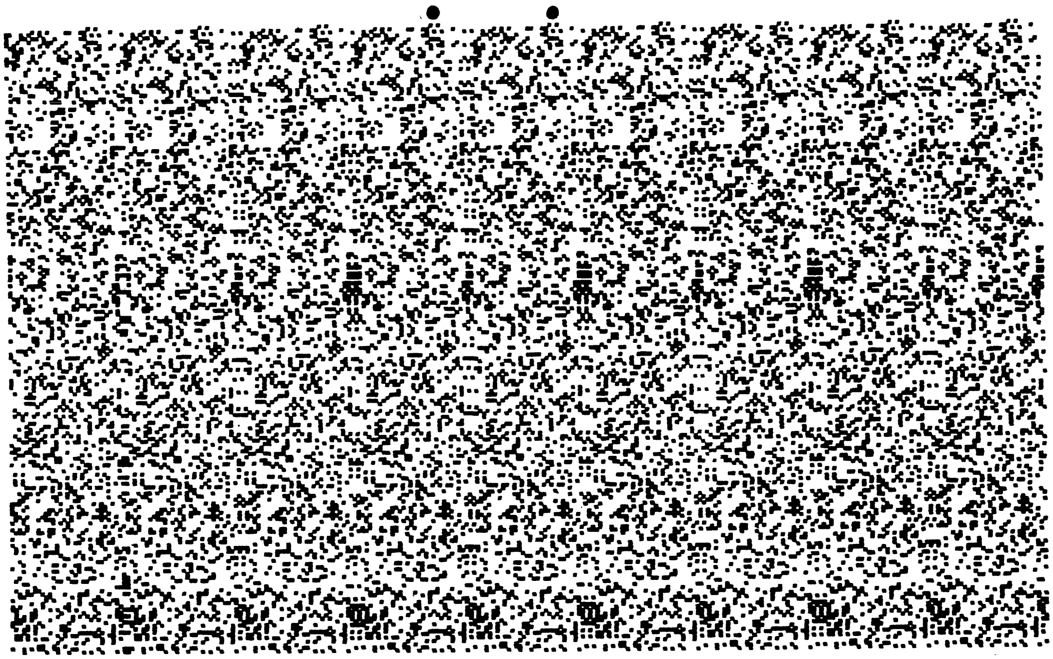

Figure 1. Autostereogram of a checkerboard in depth. Cross the eyes slightly so that the two solid fixation dots appear as three dots in a line. Focus on the center one until depth is perceived in the rest of the display. (Tyler and Clarke 1990, 183)

They can require minimal depth of understanding—as in straight-up calculation of a grade point average—or deeper understanding—as in the statistical evidence employed by studies of vaccine effectiveness. Some bases for knowledge might be default knowledge conferring or be constituted by your worldview.

As go the sources of knowledge, so go sources exploited by counterarguments. The evidence exploited by counterarguments can be statistical, mnemonic, testimonial, intuitive, perceptual, inductive, or deductive. Counterarguments can rely on default knowledge conferring foundations or fundamental worldviews. Like the sources of your would-be knowledge, counterarguments can exploit all these resources at different levels of specificity and in ways that require different kinds of tools, training, or depth of understanding.

Some of these sources provide knowledge that is more fragile—less resistant to counterevidence—than others. Seeing that it is raining provides knowledge more resistant to testimonial evidence that it isn't raining than does knowledge that it is raining provided by an internet weather map. If the previous chapter's Novice Knowledge Principle[2] can allow even the most fragile forms of knowledge to survive apparently flawless counterarguments, then that is strong evidence that all knowledge can survive apparently flawless counterarguments when the appropriate conditions are satisfied. Accordingly, I explore in this chapter two examples in which apparently flawless counterarguments would fail to destroy the output of what might seem to be the most fragile knowledge-producing source.

4.2 FELT OBVIOUSNESS

Some things feel so obvious, you know them to be true. Here's a vivid and horrifying example from Eleonore Stump:

> A young Muslim mother in Bosnia was repeatedly raped in front of her husband and father, with her baby screaming on the floor beside her. When her tormentors seemed finally tired of her, she begged permission to nurse the child. In response, one of the rapists swiftly decapitated the baby and threw the head in the mother's lap. This evil is different, and we feel it immediately. We don't have to reason about it or think it over. As we read the story, we are filled with grief and distress, shaken with revulsion and incomprehension. The taste of real wickedness is sharply different from the taste of garden-variety moral evil, and we discern it directly, with pain. (1994, 239)

Reading this story, you know the act to be wrong, and horribly so. You don't have to commit to knowing that the rapist was blameworthy; perhaps there are equally horrifying circumstances of the rapist's upbringing that are somehow exculpatory. But you know the situation is awful and that the acts done in the situation are awful as well. You don't need to reason from a prior moral theory; the intuition about the case is a constraint on moral theory.

Nor must your knowledge be unrevisable; perhaps there are further details you could discover—bizarre science-fiction-thought-experiment type details—that would rightly change your mind about the case. Still, the felt obviousness of a proposition sometimes provides knowledge that the proposition is true. Call this view *Minimal Conservatism about Knowledge* (MCK):

---

[2] As a reminder, the Novice Knowledge Principle says that, often, if a counterargument invokes evidence-types, methods, and principles with respect to which you lack relevant background knowledge or requisite methodological acumen, then knowledge can survive the counterargument being apparently flawless.

> (MCK) If it feels obvious to you that *p* then you *prima facie* know that *p*.

You *prima facie* know that *p* when, unless there are defeaters, you know that *p*. The plausibility of MCK depends on the range of defeaters it allows to defeat the *prima facie* knowledge conferred by felt obviousness. We should be liberal about allowable defeaters because there are lots of cases in which felt obviousness is not sufficient for knowledge. This can happen if

- the other conditions of knowledge aren't satisfied (for example, the conclusion is false)
- you have good reason to believe that your feelings of obviousness are unreliable about the issue in question
- there is a knowledge-destroying salient counterargument against *p*
- your feeling of obviousness is the result of unreliable belief-forming faculties
- *p* shouldn't feel obvious to you
- your belief that *p* conflicts with other attitudes you have
- objective prior probabilities require that you have low confidence in *p*

Because of the wide range of allowable defeaters, it doesn't follow from MCK that if you disagree with someone, neither of you have good arguments at the ready, and you each take your respective positions to be "just obvious," that either both of you or neither of you have knowledge (or even are in a position to know). One of you, for various, possibly externalist-friendly reasons, can know while the other doesn't. This makes MCK seem in some ways stronger and in some ways weaker than Michael Huemer's phenomenal conservatism:

> (PC) If it seems to *S* as if *P*, then *S* thereby has at least prima facie justification for believing that *P*. (Huemer 2001, 98)[3]

---

[3] In later work, Huemer modifies the account to make the seeming that *p* sufficient, barring defeaters, for S to thereby have "at least some degree of justification for believing that *p*" (Huemer 2007, 30). The reason for adding the "at least some degree of" is "to make clear that one need not have full justification for belief merely by having, for example, a weak and wavering appearance" (30). In the cases of interest in this chapter the feeling of obviousness is not "weak and wavering." If phenomenal conservatism is true, in the absence of defeaters, such felt obviousness should suffice for outright justification.

Phenomenal conservatism is a view about justification. MCK is a view about knowledge. In that way, MCK might seem stronger than phenomenal conservatism. But because of the wide range of defeaters MCK allows, MCK is no less plausible than PC. For if externalist conditions prevent a *prima facie* justified belief from counting as knowledge, it's still the case, according to MCK, that the *prima facie* justified believer *prima facie* knows.

Allowing this wide range of defeaters makes MCK in another way seem weaker than phenomenal conservatism. Huemer's phenomenal conservatism is an internalist view; the defeaters Huemer allows to defeat the *prima facie* justification provided by it seeming to *S* that *P* are all internal defeaters.[4] In that sense, MCK is closer to Berit Brogaard's "Sensible Dogmatism," according to which

> the mental states that confer prima facie justification on belief are states to which you have introspective access but the factors that determine whether those states confer justification is not something to which you have introspective access. (Brogaard 2013, 278–279)[5]

This is not to say that we should be committed to externalism—just that there is a weak version of phenomenal conservatism that is consistent with externalism. MCK is like that, except about knowledge rather than justification.[6]

Liberally interpreting what counts as knowledge-destroying defeaters makes MCK on one reading trivial: if it feels obvious to you that *p*, then you know that *p* except when you don't. MCK is not trivial if it is stipulated to include an implicit existential claim: that there are some cases in which there are no appropriate defeaters—in which it feeling obvious to you that *p* is sufficient for knowing that *p*. Phenomenal conservatism benefits from a similar supplementation. For PC to be not only plausible but substantial, there have to be some cases in which its seeming to S as if *p* results not only in *prima facie* justification, but outright justification. MCK is no less plausible in this regard than PC. If seemings sometimes result in outright justification, felt obviousness sometimes results in outright knowledge.

---

[4] See (Huemer 2006).

[5] It's also similar to Peter Markie's Qualified General Dogmatism, which holds that "seemings are a source of prima facie justification only when appropriate background conditions are met" (Markie 2013, 250). Though Markie wants the appropriate background conditions to be consistent with internalism, the requirement of background conditions might allow externalists to adopt the general schema of Markie's view.

[6] Allowing external constraints on which feelings of obviousness confer knowledge allows MCK to avoid certain objections to phenomenal conservatism, for example Clayton Littlejohn's (2011) worry that phenomenal conservatism is committed to the possibility that you can be justified in believing that you shouldn't do something even though, in fact, you should.

Here's an objection to the claim that felt obviousness sometimes results in outright knowledge: felt obviousness is fallible; falsehoods can feel obviously true. Perhaps the mere fallibility of felt obviousness destroys knowledge. This could happen if the fallibility of felt obviousness were a Gettierizing feature. Since, it might be said, we could easily go wrong by relying on felt obviousness, reliance on felt obviousness makes the truth of our beliefs too much subject to luck.

Seemings can't result in outright justification—as PC claims they can—if the fallibility of our seemings is a Gettierizing feature. We are all aware of the fact that our seemings are fallible. But awareness of Gettierizing features renders the resultant beliefs unjustified. Awareness of the prevalence of barn facades renders unjustified the belief that you are looking at a barn. Awareness that the sheep-shaped object you're looking at is a rock renders unjustified the belief that there is a sheep in the field. Therefore, if mere fallibility of felt obviousness destroys knowledge of what feels obvious to you, the universal awareness of the fallibility of felt obviousness should destroy justification of what feels obvious. Because PC requires that awareness of the fallibility of felt obviousness does not destroy outright justification, any advocate of PC needs to agree with MCK that the mere fallibility of felt obviousness does not destroy knowledge.

Therefore, if we grant with PC that felt obviousness results in *prima facie* justification and sometimes results in outright justification, we should grant that felt obviousness results in *prima facie* knowledge and sometimes results in outright knowledge. The existential claim is confirmed by example.[7] In addition to Stump's case of horrific action, it's also just obvious that no magic is going on at magic shows, that your lost car keys didn't just vanish into thin air, that the world is more than five minutes old, that three is less than eight, that my eight-year-old daughter is not the Zodiac Killer, and that there are fewer than 1,000 speckles on this hen:

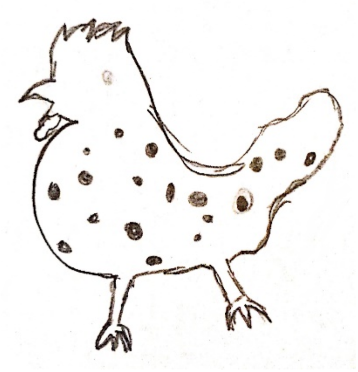

.

Some of the examples are arguably examples of *a priori* knowledge—as in Stump's example and your knowledge that three is less than eight. In others the resultant knowledge is at least partly empirical, as in the example of the speckled hen. Sometimes felt obviousness stands in for the epistemic source that generated the feeling of obviousness and transfers epistemic standing onto the proposition you know—as in your knowledge that no magic is going on in magic shows, which presumably is the result of a lifetime of experiences that you aren't immediately able to bring to mind. In some examples—most plausibly your knowledge that your keys didn't just vanish and that the world is more than five minutes old—perhaps the knowledge is grounded in your worldview, or is a framework judgment, or a Wittgensteinian hinge, or a "blik" in R. M.

---

[7] This case-based argument for MCK is somewhat similar to Jessica Brown's (2013) case-based argument for the principle she calls *Immediacy*: "(Sometimes) when S has the intuition that p, the belief which is S's intuition is immediately justified" (74).

Hare’s (1955) sense (if such things can be construed loosely enough to allow them to count as knowledge). MCK takes no stand on the reasons why obviousness can be sufficient for knowledge. It claims only that, however felt obviousness is sufficient for knowledge, it sometimes, barring defeaters, is.

This leaves MCK open to a second worry—not that it is trivial, but that it leaves no epistemic work for the feeling of obviousness to do. MCK says that there are cases in which it is just obvious to you that *p* and you know that *p*. Because it allows that, ultimately, you know that *p* because of a lifetime of experience, or because *p* is a fundamental part of your worldview, the feeling of obviousness itself might seem to be epistemically inert. What provides knowledge is your lifetime of experience or the fact that *p* is a fundamental part of your worldview, not that it feels obvious to you that *p*. MCK, the worry goes, is no more epistemologically revealing than a principle that says, “Sometimes it’s Tuesday and, also, you know that *p*.” According to MCK, sometimes it feels obvious to you that *p* and, also, you know that *p*.

Felt obviousness does epistemic work in all these cases because in all these cases felt obviousness is an essential part of the basis for your belief. Take the case in which your knowledge is most clearly empirical: the case of the hen with fewer than 1,000 speckles. You could, of course, count the speckles and reason to the conclusion that they number fewer than 1,000. But you don’t have to. When asked, “How do you know without counting that there are fewer than 1,000?” you’ll answer, “Just look! It’s obvious.” You cite the obviousness of the conclusion when asked for your basis.[8]

---

[8] When you cite felt obviousness as your basis, you need not be claiming that your belief is an inference from the fact that the content feels obvious to you. Felt obviousness can be a phenomenal justification-maker as opposed to a premise used in an inference.

You can be wrong about the bases of your beliefs. You might sincerely report that your reason for condemning Trump's missile assault on a Syrian airbase is that it risks alienating Russia, when in fact your reason is that you voted against Trump and so have an interest in concluding that any decisions he makes are wrongheaded. Much research in cognitive psychology has the lesson that motives and methods are opaque.[9] Therefore, the fact that you cite the felt obviousness of your conclusion when asked for your basis may not reveal your actual basis.

Still, in many cases, it's not clear why you'd misreport felt obviousness as your basis. It's not like you have some self-motivated interest in misreporting your actual basis in the case of the speckled hen. Even if there were, it seems plausible that your self-motivated interests would create a feeling of obviousness when there otherwise would be none. It's certainly no less plausible than that you wrongly report that the felt-obviousness is partly the basis of your belief. Furthermore, the basis of your belief may not require a causal relation between the basis and the belief it bases. This is demonstrated in Keith Lehrer's (1971) Gypsy Lawyer case. It may be sufficient for basing that you'd sincerely cite a state or proposition when asked to justify your belief. Therefore, I assume without further argument that in many of these cases felt obviousness is an essential part of the basis for your belief.

If felt obviousness is an essential part of the basis for your belief, it had better be part of what confers epistemic status on the proposition you believe. This follows from a premise Huemer uses in his self-defeat argument for phenomenal conservatism: "if one's belief that *p* is based on something that does not constitute a source of justification for believing that *p*, then one's belief that *p* is unjustified" (Huemer 2007, 40). The premise depends on a distinction between doxastic justification—the justification that accrues to beliefs—and propositional justification—the justification that accrues to propositions. On the construal I favor, a proposition is justified for you just in case you should believe it, while your belief in that proposition is doxastically justified only if your belief is based in the right way on whatever justifies the proposition. This is standardly referred to as the “basing requirement” for doxastic justification. Here’s Declan Smithies’ statement of the requirement:

---

[9] See, for example, (Vazire and Carlson 2011), (Wallis 2008), and (Nisbett and Wilson 1977).

> One’s belief that *p* is doxastically justified if and only if one *has* propositional justification to believe that *p* and, moreover, one *uses* one’s propositional justification in believing that *p* on the basis of one’s propositional justification to believe that *p*. (2011, 26)

To sum up the standard picture, doxastic justification = propositional justified + proper basing.

The standard picture is more than a claim about entailments. It’s a metaphysical claim: propositional justification comes first. Your belief is justified *because* it’s a belief in a proposition with propositional justification and it’s properly based.[10] The standard view is also two-directional: propositional justification plus proper basing is both necessary and sufficient for doxastic justification. Neither the metaphysical claim nor the sufficiency direction is required to show that felt obviousness is epistemically relevant to the knowledge it’s sometimes sufficient for. All that’s needed is a minimal conception of doxastic justification: that if your belief that *p* is doxastically justified, then your belief that *p* is properly based—based on whatever it is that propositionally justifies *p*.

---

[10] See (Kvanvig and Menzel 1990) for an argument for the metaphysical claim. Smithies (2012, 6) also endorses the order of explanation required by the metaphysical claim. Alvin Goldman (1979, 21) and John Turri (2010, 320) deny the metaphysical claim. But even Goldman and Turri either are committed to or leave open the bare entailment claim that if your belief that *p* is doxastically justified, then *p* is propositionally justified for you and properly based. Turri offers counterexamples to the other direction of this bare entailment claim—to the claim that propositional justification plus basing suffices for doxastic justification. The direction he leaves open is all that is needed for felt obviousness to have epistemic force.

This minimal conception of doxastic justification entails that, in the above examples, felt obviousness is a source of propositional justification for what you know. Focusing on the hen example, suppose by hypothesis that you know that the hen has fewer than 1,000 speckles and that an essential part of your basis is that it's just obvious to you that it does. By the minimal conception of doxastic justification, your belief is properly based; it is based on something that makes it justified for you that the hen has fewer than 1,000 speckles. Therefore, once it's granted that you know that the hen has fewer than 1,000 speckles and that your basis is (in essential part) that it's just obvious it does, it follows from the minimal conception of doxastic justification that felt obviousness is a *proper* basis for belief in that case.

Sometimes the felt obviousness of your position is supplemented by other kinds of evidence. For example, when it comes to the denial of psychic phenomena, non-believers—sometimes called "goats" (as opposed to "sheep") (Blackmore 1992, 368)—supplement their feelings of obviousness with the fact that there is no single decisive study that demonstrates the existence of psychic phenomena and that whenever believers are called upon for a demonstration, they fail. Atheists supplement their feeling that God's non-existence is obvious with arguments about flying spaghetti monsters, the Problem of Evil, the existence of multiple competing religious views, and the fact that a person's religion tends to be a product of contingencies of their upbringing.

As we saw in the last chapter, some of these arguments—if they rely on the principle that absence of evidence is evidence of absence—won't ground the kind of knowledge that resists being unable to figure out what's wrong with a relevant counterargument. References to flying spaghetti monsters won't help sustain knowledge in the face of relevant counterarguments, because they rely for their force on there being no relevant counterarguments in need of resistance. These kinds of supplements to felt obviousness won't help show how atheistic or goat knowledge can resist the inability to figure out what's wrong with theistic arguments or arguments for psychic phenomena.

Other arguments—like the Problem of Evil—may offer more hope, though they also have a number of problems for those looking to avoid reliance on felt obviousness. First, they may themselves rely on felt obviousness in the making of their case—for example, the felt obviousness (to the atheist) that there is no plausible secret greater good that could make the suffering in the world worth it. Second, some of the arguments may undermine atheistic and goat knowledge just as much as they undermine the relevant counterarguments. Worries about the role of upbringing in religious beliefs might undermine atheism as much as any specific religious view. Those raised in atheist homes tend to be atheists just as those raised in Christian homes tend to be Christians.

Therefore, I assume that sometimes the feeling of obviousness is not only sufficient for knowledge, barring defeaters, it's also sometimes necessary. My attention is on two ways that obviousness might manifest itself in putative knowledge—the atheist's putative knowledge that there is no God and the goat's putative knowledge that there is no such thing as psychic phenomena.

4.3 *PRIMA FACIE* GOAT AND ATHEISTIC KNOWLEDGE

Psychic phenomena, or *psi*, are often divided into four categories: 1) psychokinesis—the ability to intentionally manipulate objects outside your body without touching them or touching any intervening things, 2) precognition—the ability to reliably predict future contingent events without relying on facts about the present or past, 3) telepathy—the ability to directly transfer information between your mind to another's, and 4) clairvoyance—the ability to immediately know spatially remote contingent facts.[11]

C. D. Broad, though convinced of the existence of psi, argues that the fact that psi conflicts with what Broad calls "basic limiting principles" means that, prior to an experiment demonstrating the existence of psi, philosophy at least can

> afford to ignore psychical research; for it is no part of its duty to imitate the White Knight by carrying a mousetrap when it goes out riding, on the offchance that there might be mice in the saddle. (1949, 292)

The basic limiting principles, according to Broad, include "general principles of causation" among which are principles denying the possibility that future events can bring about past events and the possibility that causes can have immediate effects at a remote spatial distance (293), principles that impose "limitations on the action of mind on matter" (293), and principles that impose "limitations on ways of acquiring knowledge" (294).

If Broad is right that your defeasible license to deny the existence of psi depends on your license to believe the basic limiting principles, then it looks like—prior to any pro-psi experimental results—not only can philosophy "afford to ignore" psi, you can know it to be non-existent. For, according to Broad, the basic limiting principles are in a very strong position epistemically: they either “seem to be self-evident” or “are so overwhelmingly supported by all the empirical facts which fall within the range of ordinary experience and the scientific elaborations of it . . . that it hardly enters our heads to question them” (291).

---

[11] This taxonomy is from Dean Radin (1997), though the definitions are mine. James Alcock points out that one difficulty with defining "psi" is that, to exclude uncontroversial ways in which, say, minds can communicate with another (for example, through speech), definitions of psychic phenomena normally include clauses like, "in the absence of any 'normal' mechanism that could account for it" (2003, 33). I've tried to avoid such clauses in my definitions, but only by using phrases like "directly transfer" or "immediately know." Where I do not invoke these kinds of mystery clauses, I suspect there likely are counterexamples that would satisfy the definitions, but which we'd not want to categorize as psi. Such counterexamples would require "normal mechanisms" fixes.

Alcock thinks that defining psi in contrast to "normal" means psi ends up being "just a label. It has no substantive definition that goes beyond saying that all normal explanations have apparently been eliminated" (2). I don't think this is a decisive problem. It would be impressive enough if we could show that there was such a thing as mind-to-mind transfer of information not through any ordinary means. If it could be proven that one person could successfully pass information to another, in a separate room, without making a sound or indeed moving the air between the rooms, without sending any radio waves and without access to any piece of technology that might effect the transfer, that would be enough for me. To have a robust scientific theory of psi, we might need more. But to have a result that destroys at least my worldview, the purely negative definition is enough.

Broad thinks that there have been experiments demonstrating the existence of events that are incompatible with these basic limiting principles. For example, of “Dr. Soal’s experiments on card-guessing with Mr. Shackleton,” Broad says,

> There can be no doubt that the events described happened and were correctly reported; that the odds against chance-coincidence piled up to billions to one; and that the nature of the events, which involved both telepathy and precognition, conflicts with one or more of the basic limiting principles. (292)

There have been numerous criticisms of Soal’s experiments,[12] but my point here is simply that even someone with sympathies to psi as strong as Broad’s grants that prior to the experimental results, the felt obviousness of the limiting principles channels “overwhelming” empirical support. If their felt obviousness channels overwhelming empirical support, then in the absence of relevant counterarguments, you can know to be false claims that conflict with them—for example, that there is such a thing as psychic phenomena.[13]

It is less plausible that the basic principles that motivate atheistic denial of God’s existence channel overwhelming empirical support. The principles that atheists would have to invoke would need to exclude the possibility of transcendental entities. It’s hard to see how experiential faculties whose ambit is contingent features of the world could get us access to facts about entities that transcend that world. At any rate, I don't assume that fundamental atheistic principles have the kind of overwhelming empirical support that Broad's limiting principles do.

---

[12] Most of the criticisms involve allegations of fraud. For example, Christopher Scott and Phillip Haskell (1973) charge that one of Soal’s participants—“Mrs. G. A.,” later identified as Gretl Albert—reported observing Soal changing numbers on the score sheets—from ones to fours and fives. Scott and Haskell (1974) provide confirming evidence for Albert’s allegations, pointing out that on scoring sheets for those sessions in which Albert was a participant (and for two other participants as well), there were a greater than expected proportion of fours and fives and a lesser than expected proportion of ones. Betty Markwick (1978) notes additional failures to properly randomize card sequences and, worse, provides evidence of after-the-fact introduction of data points that were scored “hits.”

[13] There is another way you might be able to know, barring defeating evidence, that there is no such thing as psi—by reasoning from the fact that, as George Price put it, psi is “incompatible with current scientific theory” (1955, 360). Psi-sympathist Jean Burns notes that, unlike psi, “in presently known physics all transfer of information involves a signal” (2003, 8). This is distinct from the method of felt obviousness, because it involves an argument from a theoretical physical claim.

Still, as we've seen, the kind of felt obviousness that generates knowledge need not have empirical support at its base. Felt obviousness can generate *a priori* knowledge, and can perhaps have its source in fundamental worldviews or framework judgments. Of course, it is possible for something to feel obvious to you even though you don't know it to be true. You may fail to satisfy some necessary externalist conditions. There may be such a thing as objectively irrational prior probability assignments, and you might assign objectively irrational prior probabilities in such a way that it generates a feeling of obviousness about propositions that aren't true. Under what conditions felt obviousness generates—barring defeaters—what would otherwise be knowledge is a difficult problem and the subject matter of the rest of epistemology. What's important is that felt obviousness *can* generate knowledge. The question I'm interested in is whether, *if* it does, that knowledge can survive the inability to figure out what's wrong with relevant counterarguments.

For atheists to be in a position to have counterargument-safe knowledge, there must be ways to generate the atheist's felt obviousness that do not rely on the principle that absence of evidence is evidence of absence. Fortunately, there are. The atheist might incorporate the absurdity of God's existence into their framework judgments or fundamental worldview. The atheist might have a strong *a priori* intuition—of the sort that you have in response to Stump's example of morally horrific action. The atheist's feelings of obviousness might manifest objectively correct prior probability assignments. I assume that some of these sources are sufficient—barring relevant counterarguments—for knowledge.

Barring relevant counterarguments, then, goats and atheists can be in a position to know that there are no psychic phenomena and that God does not exist, respectively. Were there no relevant counterarguments, goats and atheists could know these things, and could know them without relying on the principle that absence of evidence is evidence of absence. But there are relevant counterarguments and some are hard to answer. Do they defeat the knowledge-level support that goats and atheists can otherwise have?

4.4 GOAT AND ATHEISTIC KNOWLEDGE, UNDEFEATED

4.4.1 The Fine Tuning Argument

Keith DeRose, though he agrees with MCK, thinks that the existence of relevant counterarguments—and those who present them—means that felt obviousness does not suffice for atheistic knowledge (nor, in all likelihood, for theistic knowledge). In a 2014 interview in the *New York Times* "Opinionator" column, DeRose says that while "I, in fact, think some of our most important and interesting knowledge comes not through anything like arguments, but from just rightly rejecting as bizarre things that so strike us," this knowledge is undermined when "the matter strikes others differently." In such cases, apparent bizarreness of the alternative is not sufficient for knowledge: "you need a good argument to claim to know they're wrong." At the very least

> To know that God does (or doesn't) exist, you have to show that there are no arguments for atheism (or for theism) that a reasonable person could find plausible. But to support that claim you would have to have better critiques of all those arguments than I've ever seen. (DeRose in (Gutting 2014))

As we've seen, the fully general version of DeRose's contention—that to know something you need good critiques of the arguments against it—is false. But that doesn't mean that he's wrong about the specific case he's discussing: belief about God's existence in which the only basis for knowledge is the felt absurdity of the alternative.

When your basis for knowledge is simply the felt absurdity of the alternative, do you then need good critiques of relevant counterarguments? Here's a simple argument from the conclusion of the previous chapter to the claim that you don't. According to the conclusion of the previous chapter, you can retain knowledge in the face of a relevant counterargument in which you can't expose a flaw when 1) you lack sufficient facility with the methods and evidence-types employed by the relevant counterargument, but 2) you have facility with the methods and evidence-types employed in your putative knowledge source. We all have facility with the method of felt obviousness, so atheists who base their purported knowledge of God's nonexistence on felt obviousness have facility with the methods involved in their basis. As long as they lack sufficient facility with the methods employed by some theistic argument they can't expose a flaw in, they can retain knowledge that God does not exist.

This simple argument may not seem to do the trick. DeRose's point stands: God's existence can strike the theist as just as obvious as God's nonexistence strikes the atheist. This fact undercuts the evidential force provided by the felt obviousness, to the atheist, of atheism. It's an undercutting defeater, because it renders untrustworthy strong feelings of obviousness about God's existence. Why suppose that what the theist feels is obvious is any less trustworthy than what the atheist feels is obvious?[14]

---

[14] Perhaps atheists or theists have independent reasons for thinking that the other's intuitions are biased, but DeRose is skeptical about the possibility:

> If we knew that adherents to other religions came to hold their beliefs in some way that discredits them (say, through brainwashing), we might still know those beliefs are wrong on the basis of how bizarre they seem to us. Of course there are probably some individual believers who have come to hold their beliefs in a way that discredits them. But we don't know enough about many believers to discredit their beliefs. So I don't think we can know they're wrong just because their beliefs strike us as bizarre. (DeRose in (Gutting 2014))

If this sort of consideration undermines the atheist's position, it does so at the cost of rendering inert the theistic arguments. It's not the theistic arguments that undermine what would otherwise be atheistic knowledge. It's the fact that some theists find the existence of God as obvious as some atheists find the non-existence of God, in the same way that Stump says that, for her,

> The mirror of evil becomes translucent, and we can see through it to the goodness of God. . . . [S]tart with a view of evil and deep taste of the goodness of God, and you will know that there must be a morally sufficient reason for God to allow evil—not some legal and ultimately unsatisfying sort of reason, but . . . a reason in which true goodness is manifest. (1994, 242)

Stump's "tast[ing] and see[ing]" (241) that there is an all-good God might seem to undercut atheistic knowledge, if that atheistic knowledge is based on "tasting and seeing" that God doesn't exist. If it does, it does so no less because Stump's belief isn't based on argument. It's the theist's existence, not their arguments, that arguably undermines what would otherwise be atheistic knowledge.

At the risk of avoiding what might seem to be the crux of the matter, I do not take a stand on this contested issue in the epistemology of disagreement. I take no stand on conformism or non-conformism, whether permissivism or uniqueness is true, whether fundamental framework judgments can allow for knowledge, or whether there are objective or externalist parity-breakers that favor one member of groups of disagreeing peers. The question I'm interested in is this: do the theistic arguments—as distinct from the theist's existence—undermine atheistic knowledge more than does the mere fact that theists disagree with atheists about what is obviously true? On the assumption that atheistic knowledge based on sheer felt obviousness can survive the existence of theists who feel it's just obvious that God exists, can atheistic knowledge survive the theistic counterarguments, if the atheist can't figure out what's wrong with them?

It can, provided the atheist is sufficiently ill-versed in the methods or evidence-types invoked by the theistic arguments so that it would be unsurprising that they'd find those arguments apparently flawless even if they're misleading. I use as my running example the fine-tuning argument. According to the fine-tuning argument, the various constants that show up in the laws of nature—for example, the gravitational constant—would only have to differ by tiny amounts for it to be impossible for the universe to sustain life or, indeed, for matter to hold together at all. L. Stafford Betty and Bruce Cordell cite physicist Don Page as calculating the odds of the formation of our universe at ten billion to the 124^{th} power to one against (Betty and Cordell 1987, 416).

Of course, the universe had to turn out *some* way. Whatever way it turned out, the odds were astronomically against it turning out that way. If you drop 150 toothpicks, they will end up arranged on the floor in a certain way. The odds that they would have come to be arranged that way are very low. This is no evidence that the arrangement was intelligently designed. We can't infer from something being unlikely that it was designed that way. But the fine-tuning argument is stronger than that. Ours is not just an unlikely universe; it's a universe of special interest to us. It's more like dropping 150 toothpicks on the floor and having them spell out, in perfect font, "I am trapped in a box of toothpicks"—a precise arrangement no less initially probable than many other precise but haphazard arrangements.

More relevantly, consider a version of John Leslie's (2002, 13 ff.) Firing Squad Story. You are facing a firing squad. There are one hundred sharpshooters. Each rifle, you know, has a one in two chance of misfiring, unless tampered with. Some arrangement of firings and misfirings has to occur, and whichever arrangement occurs has no greater prior probability than any other particular arrangement: it's no less probable that they all misfire than that gun #3 fires while the rest of the guns misfire, or that some specific jumbled pattern occurs, or that they all fire. Still, if all the guns misfire, so that you are left alive, you are well justified in believing that the guns were tampered with.[15]

Unlikely results with a certain kind of special interest confirm intelligent design more than merely unlikely results. But not all results of special interest do. If I enter a regular lottery and the numbers I pick come up, that's no evidence that the lottery was rigged, even though the result is of special interest to me. But if I enter a regular lottery and the lottery commissioner (who, through an unfortunate legislative oversight, is allowed to buy a ticket) wins, that is evidence that the lottery was rigged. If the lottery commissioner's numbers come up, it is more probable that the election is rigged than if my numbers come up, even though the prior probability of each set of numbers coming up is the same, and even though both results are of special interest. Here's why (again, it is necessary to indulge in a little math):

Let *r* be the proposition that the lottery is rigged, *me* be the proposition that my numbers come up, and *lc* be the proposition that the lottery commissioner's numbers come up. We need to compare the probability that the lottery is rigged conditional on my numbers coming up to the probability that the lottery is rigged conditional on the commissioner's numbers coming up. Here are the relevant instance of Bayes' Theorem:

$$\mathrm{Pr}(\mathrm{r}|\mathrm{me}) = \frac{\mathrm{Pr}(\mathrm{r})\ \mathrm{Pr}(\mathrm{me}|\mathrm{r})}{\mathrm{P}(\mathrm{me})}$$

$$\mathrm{Pr}(\mathrm{r}|\mathrm{lc}) = \frac{\mathrm{Pr}(\mathrm{r})\ \mathrm{Pr}(\mathrm{lc}|\mathrm{r})}{\mathrm{P}(\mathrm{lc})}$$

---

[15] Betty and Cordell offer an example involving a lottery machine that spits forth your nine-digit Social Security Number (416).

I assume that the prior probability that the lottery is rigged is very low and so that the prior probability that my numbers come up is very close to the prior probability that the commissioner's numbers come up. But conditional on the lottery being rigged, it is much more probable that it is rigged in favor of the commissioner than in favor of me. After all, I'm a nobody whereas the commissioner is the one who would be in charge of rigging the lottery. Therefore, the numerator of the second instance of Bayes' Theorem is much higher than the numerator of the first instance, though the denominators are close to equal. Therefore, the probability that the lottery is rigged conditional on the commissioner's numbers coming up is much higher than the probability that the lottery is rigged conditional on my numbers coming up.

There might be quibbles with the empirical claims; perhaps as a matter of empirical fact, when commissioners rig lotteries they tend to disfavor themselves. If so, then it might be that the lottery is more probably rigged in my favor than in the commissioner's. If I win, I should be more confident that the lottery was rigged than if the commissioner wins. I don't think these empirical claims are true. Regardless, the probability of intelligent design conditional on some evidence e is proportional to the probability of e conditional on intelligent design.

Like Leslie's Firing Squad Story, what seems distinctive about our universe is that it ends with our existence. It is plausible that the probability that the universe favors creatures like us conditional on it being intelligently designed is much higher than the probability that it favors creatures like us conditional on it not being intelligently designed. Therefore, the results we observe in the universe are more like the case of the commissioner's numbers coming up than the case of my numbers coming up. Our existence tends to confirm the hypothesis of intelligent design. Because the prior probability of the universe favoring our existence without intelligent design is extremely low, our existence greatly favors the hypothesis of intelligent design.

Here are three worries about this argument, all targeted at the claim that the prior probability of the universe favoring our existence is extremely low. The first has to do with the difficulties assigning probabilities to the values of the physical constants. As Timothy McGrew, Lydia McGrew, and Eric Vestrup point out,

> Probabilities make sense only if the sum of the logically possible disjoint alternatives adds up to one—if there is, to put the point more colloquially, some sense that attaches to the idea that the various possibilities can be put together to make up one hundred percent of the probability space. But if we carve an infinite space up into equal finite-sized regions, we have infinitely many of them; and if we try to assign them each some fixed positive probability, however small, the sum of these is infinite. (2001, 1031)

According to the fine-tuning argument, the prior probability that some constant takes any specific value is no greater than that it takes any other specific value. Because there is an infinite numerical region from which to take the physical constants, that probability must be very low. If the probability is zero, since it is equally probable that the constant takes any other value, the sum of those probabilities is still zero. If it's greater than zero, then since it is equally probable that the constant takes any other value, the sum of those probabilities is infinite. Therefore, if it is coherent to assign probabilities to the possible values of the constants, there must be different probability assignments corresponding to different values. We have no way to exclude the possibility that the values of the constants are the values the constants had high prior probabilities of having.

The second worry has to do with the possibility of life developing in a universe with constants quite different from ours. As Betty and Cordell put the point, following (Carr and Rees 1979),

> intelligent life might have evolved out of a very different kind of universe, and not just the one we know, with its particular governing constants. Of course, this "life" would have almost nothing in common (at least physically) with the biological life forms to which our universe has given rise. But what, except our own narrow experience, is

> guiding us when we limit life to conditions such as our universe provides? (Betty and Cordell 1987, 417)

It might be improbable for an undesigned universe to contain homo sapiens. But it might not be improbable for an undesigned universe to contain some kind of intelligent life. This worry undercuts the fine-tuning argument, if the argument depends on the improbability of any thinking life at all arising in the universe.

Finally, various commenters have raised the possibility that a vast multiplicity of universes exists. If this array of universes arises by chance, it is still likely that one of those universes contains beings like us. There are various ways this could happen, perhaps via Hugh Everett's (1957) "many worlds" hypothesis or John Wheeler's (1973) contention that the universe "oscillates" through creative and destructive phases. But however it happens, it's unsurprising for there to be a world in which beings like us exist.

What should we think about these worries? Should the atheist rest content that the fine-tuning argument has been disposed of? Should these worries license the atheist in maintaining strong confidence in their atheism? Betty and Cordell

> admit that the theory of life's emergence from conditions greatly dissimilar to what we know on the Earth, however unlikely this scenario may be, and the theory of an infinite number of universes, however farfetched it may seem, are at least not intrinsically absurd. (418)

On that basis, they search elsewhere for a convincing teleological argument. And I agree that the atheist should be happy to recognize that these are difficulties for the argument. But if the atheist is forced into taking one of these three options, that's a victory for the theist. We wouldn't have thought that to be a rational atheist you have to think either that the physical constants are necessarily what they in fact are, or that life could have evolved in radically different universes, or that many possible universes exist. We wouldn't have thought that an atheist was forced into one of these three positions. As Betty and Cordell say in the above passage, these possibilities are "farfetched" and "unlikely." So, if the atheist has to accept one of the three positions, that's a victory for the fine-tuning argument.

Put another way, many atheists' confidence that there is no God is stronger than their confidence that there are multiple universes, that life could have developed in a universe with different laws, and that the physical constants couldn't easily have differed from their actual values. Their confidence that there is no God is stronger than any of these hypotheses taken singly and stronger than the disjunction of the three. They don't take themselves to know that the disjunction of the three hypotheses is true. On the contrary, each step of the fine-tuning argument might remain more compelling to them than the truth of any of the potential solutions, and of their disjunction. But they might nonetheless take themselves to know that there is no God. Therefore, their knowledge that the fine-tuning argument is misleading had better not depend on the truth of the disjunction of these three solutions.

There are other responses to the fine-tuning argument.[16] But no matter how many disjuncts are added to the disjunction, an atheist might not take themselves to know that the disjunction is true, if it doesn't feel obvious to them that it is true, and if they don't have sufficient facility with the methods and evidence-types involved. Nor are the trustworthy experts in sufficient agreement about the truth of any of these proposals to grant atheists testimonial knowledge that the disjunction is true. After all, even Elliott Sober, who *does* have the requisite background knowledge and methodological acumen, can—by his admission—misdiagnose the problem with the fine-tuning argument:

> I argued in my 2004 paper that the prisoner [standing before the firing squad] is in the grip of an observational selection effect and so, when the relevant likelihoods are used to evaluate his evidence, a likelihood equality arises and he must conclude that his survival

---

[16] See, for example, (Sober 2009).

> does not discriminate between ID and Chance. Weisberg (2005) and several other patient friends have persuaded me that I was mistaken. But I still think there is an observation selection effect in the fine-tuning argument. If the prisoner is right, how can the fine-tuner be wrong? Careful attention to time indices draws our attention to the following:
>
>> Pr(I observe at t_3 that I am alive|the firing squad decides at t_1 that it will spare my life when it fires at t_2 & I am alive at t_2 when the squad fires) > Pr(I observe at t_3 that I am alive|the firing squad decides at t_1 that it will fire in randomly chosen directions at t_2 & I am alive at t_2 when the squad fires).
>
> There is a difference between the firing squad and fine-tuning. I missed it in Sober (2004), thinking that there is an observation selection effect in both. Weisberg missed it too, thinking that there is an observation selection effect in neither. The prisoner's being alive at t2 does not screen off ID and Chance from his observing at t3 that he is alive. But my being alive at t2 does screen off ID and Chance from my observing at t3 that the constants are right. (Sober 2009, 86)

This response and the fine-tuning argument itself are not the kinds of arguments that someone lacking the relevant methodological acumen and background knowledge has sufficient facility to reliably evaluate.

Of course, given the conclusion of the past few chapters, it is the atheist's lack of facility with the methods and evidence-types involved in the fine-tuning argument that saves their knowledge that there is no God, if they have that knowledge in the absence of theistic arguments. It should be no surprise to the untrained atheist that they would find each step of a relevant counterargument employing the methods and evidence-types of the fine-tuning argument compelling even if it were misleading. The evidence-types and methods employed in the fine-tuning argument are evidence invoked in theoretical physics and methods involved in Bayesian updating. Many atheists have, at most, a dilettante's skill in the techniques of Bayesian updating and no more than a *New York Times*-headline amount of background knowledge in theoretical physics. When it comes to the methods involved in the fine-tuning argument, they are indeed the Unfrozen Caveman Lawyer.

By the criterion in the last chapter, atheists are only able to retain knowledge that there is no God in the face of a counterargument that employs methods and evidence-types with which they lack sufficient facility if their basis for knowledge involves methods or evidence-types with which they have sufficient facility. Do atheists have sufficient facility with the source of what would otherwise be their knowledge that there is no God?—Do they have sufficient facility with felt obviousness? As we saw in the last chapter, a sufficient condition for having sufficient facility with a putative knowledge source is that the source—by itself—would deliver knowledge. If it is unsurprising that you'd be moved by the source even if the source were misleading, then you couldn't know on the basis of this source. But, we've determined, the atheist's feeling of obviousness that God doesn't exist, by itself—that is, in the absence of relevant counterarguments—can deliver knowledge. Therefore, atheists have sufficient facility with felt obviousness, even when they lack sufficient facility with the evidence-types and methods invoked in the fine-tuning argument.

We might think that, though felt obviousness is simpler than theoretical physics and Bayesian confirmation theory, it's more difficult to have facility with its use. What makes obviousness justification conferring may be beyond our ken. There may be externalist conditions that must be satisfied for felt obviousness to confer knowledge. There may be objective facts about when a proposition should feel obvious to you and when it shouldn't. To have facility with the use of felt obviousness, it might seem, you would need to know what it is that makes felt obviousness knowledge conferring, and when those conditions are satisfied and when they're not.

Neither is required. First, you don't have to know what it is that makes felt obviousness knowledge conferring in order to have sufficient facility with felt obviousness, just as you don't have to know what makes perception knowledge conferring in order to have sufficient facility with perception. Second, if you don't need to know what the conditions on knowledge conferring felt obviousness are, then you can hardly be expected to know when those conditions—whatever they are—are satisfied.[17] Stump's example of horrific wrong action confirms these conclusions. You can continue to know that the actions are wrong without having an understanding of how felt obviousness confers knowledge or whether you satisfy various externalist conditions.

What goes for atheists arguably goes for theists. Some theists find it just obvious that God exists. They "taste and see" the goodness of God, or know via religious experience. If these or some other means of coming to believe in God's existence otherwise provides knowledge—if they satisfy the relevant externalist conditions—then theists can be in a position to resist atheistic arguments against God's existence. They can retain their knowledge that God exists in the face of apparently flawless atheistic arguments, provided they lack sufficient expertise to reliably evaluate those arguments.

Whether a theist lacks sufficient expertise to reliably evaluate some atheistic argument is case-specific. It depends both on the nature of the argument and the qualifications of the theist. We might think, for example, that the ignorance most atheists are in with respect to the subtle mechanisms invoked in the fine-tuning argument does not carry over to some theists confronted with the Argument from Evil. Many theists have sufficient facility with the Argument from Evil to reliably evaluate it—at least as much facility as they have with respect to their sources of knowledge of God's existence. This is why theists in particular struggle so sincerely with the argument. That's not to say that such theists lose their knowledge in the face of the Argument from Evil—just that those theists can't make use of the Novice Knowledge Principle to explain why they can retain knowledge that God exists if the Argument from Evil is apparently flawless.

---

[17] If the conditions are genuinely externalist, you precisely don't need to know whether they're satisfied.

4.4.2 Statistical Arguments for Psi

Dean Radin reports that the first recorded experiments on psi date back to the second millennium BC's Shang Dynasty, whose "Oracle Bones"—tortoise shells whose cracks allegedly foretold future events after being tossed into the fire—were used to record both the predictions and what actually happened (2009, 55). This hardly satisfies George Price's requirement on a sufficient test for psi:

> What is needed is one completely convincing experiment—just one experiment that does not have to be accepted simply on the basis of faith in human honesty. We should require evidence of such nature that it would convince us even if we knew that the chief experimenter was a stage conjurer or a confidence man. (1955, 363)

Recent studies use more laboratory-approved methods. For example, Darryl Bem's (2011) study on precognition includes nine experiments, all but one of which show significant influence of future events on present actions. In his second experiment on precognitive avoidance of negative stimuli, subjects correctly preferred one of two pictures prior to the pictures being revealed; they were precognitively averse to negative pictures. One hundred fifty Cornell undergraduates over 150 sessions picked the correct picture 51.7% of the time as opposed to the expected 50% of the time (Bem 2011, 413).

There has been difficulty replicating this and the rest of Bem's experiments. Eric-Jan Wagenmakers (2014) handily summarizes the recent attempts. The left hand column includes the name of the replicator and the experiment they were attempting to replicate, the middle column shows the size of the replicating experiment and the right hand column shows the effect size demonstrated:

Experiment	**N**	**ES**
Galak Exp 1	112	–0.113
Galak Exp 2	158	0
Galak Exp 3	124	0.110

Experiment	**N**	**ES**
Galak Exp 4	109	0.170
Galak Exp 5	211	0.050
Galak Exp 6	106	–0.029
Galak Exp 7	2469	–0.005
Ritchie Exp 1	50	0.016
Ritchie Exp 2	50	–0.222
Ritchie Exp 3	50	–0.041
Robinson	50	–0.120
Subbotsky Exp 1	75	0.282
Subbotsky Exp 2	25	0.302
Subbotsky Exp 3	26	–0.412
Traxler Exp 1	48	.060
Traxler Exp 2	60	-.346
Wagenmakers	100	–0.022

As Wagenmaker points out, though he "did not conduct the meta-analysis . . . from eyeballing the numbers it seems that there is nothing there whatsoever."

The failure to replicate is not limited to Bem's experiments. Alcock notes that "parapsychologists have *never* been able to produce a successful experiment that neutral scientists, with the appropriate skill, knowledge, and equipment, can replicate" (2003, 35). Psi-sympathist Jean Burns, likewise, laments that "one of the most frustrating aspects of [psi's] elusiveness is the failure to replicate large studies which in cumulative effect had given highly significant statistical evidence for psi" (2003, 24). This is a significant problem, because replicability "is perhaps the best safeguard against being taken in by results produced by error, self-delusion or fraud" (Alcock 2003, 34).

If there is failure of replication, might there be other ways to confirm the existence of psi? Contemporary pro-psi researchers don't generally point to single, decisive studies the way that we saw Broad tries to (and as Price demands). What they point to are meta-analyses. A meta-analysis uses advanced statistical techniques to generalize from multiple studies. It has the advantage of, ideally, providing a larger population. It has the disadvantage of—if some of the input studies have bad data—tainting the results of the meta-analysis while hiding the bad data of the garbage input study: garbage in, garbage out. It is also subject to the “file drawer problem”—the problem that meta-analyses might leave out studies with results unfavorable to the target conclusion. Therefore, meta-analyses require care in implementation. Still, with such care, “rigorously conducted systematic reviews and meta-analyses are essential for evidence-based decision making” (Nordmann, Kasenda, et al. 2012, 1).

The widest scale pro-psi meta-analyses cover studies that employ a procedure for testing telepathic response known as the *ganzfeld procedure*. The procedure requires two people: a receiver and a sender. The receiver and sender are isolated from each other, so that neither can get any cues from the other. The sender looks at a card and mentally “sends” the image to the receiver. The receiver is given a choice of four cards on one of which is the picture that is being sent. If the experiment is well conducted, the scientists showing the sample cards to the receiver do not know which card is being sent. The most sophisticated ganzfeld experiments—called *autoganzfeld* experiments—use computers to randomize cards and automate the process in ways that remove the possibility of “sensory leakage,” the phenomenon that occurs when the experimenter unwittingly cues information about the cards to the receiver.

In Daryl Bem and Charles Honorton’s (1994) meta-analysis of eleven autoganzfeld experiments covering a total of more than 240 subjects and 329 trials, the overall hit rate was over 32%, as opposed to the expected 25% (11). A more recent meta-analysis by Storm, Tressoldi, et al. shows similar results: over the twenty-nine studies included in the meta-analysis, they “found a 32.2% hit rate” (2010a, 475). Nor, argue their authors, are these meta-analyses subject to the various pitfalls to which meta-analyses are frequently subject. Bem and Honorton point out that in their meta-analysis, “the 11 studies just described comprise all sessions conducted during the 6.5 years of the program. There is no ‘file drawer’ of unreported sessions” (1994, 10). Storm, Tressoldi, et al. note that, for their results to be insignificant, “there would have to be no fewer than 293 unpublished articles in existence with overall nonsignificant results to reduce our significant findings to a chance result” (2010a, 475). Honorton and Diane Ferrari, discussing their meta-analysis, conclude that “over 46 unreported studies must exist for each reported study to reduce the cumulative outcome to a nonsignificant level” (1989, 286). Similarly, Bem and Honorton argue, we need not worry overmuch about methodological problems with the covered studies that might allow “garbage in”: even if the meta-analysis only allows in the most rigorously controlled and secure studies, “the remaining studies are still highly significant” (1994, 7).

The pro-psi meta-analyses are published in the *Psychological Bulletin*, which is a peer-reviewed publication under the auspices of the American Psychological Association. If not for their subject matter, this would be sufficient evidence that the meta-analyses don’t commit obvious, surface blunders. And, of course, psi-critics would have been very quick to pounce on the data had they shown that there was no psi-effect. Therefore, there seems to be enough pro-psi evidence that, in Jean Burns’ words, “if psi were any ordinary phenomenon, it would probably be provisionally accepted and non-controversial” (2003, 23).

There are also, as one might imagine, extensive critiques of the meta-analyses.[18] A third danger inherent in meta-analyses is the possibility of what is called “heterogeneity” of the individual studies. Obviously, if one study is on the efficacy of an experimental arthritis treatment while another is on the efficacy of an experimental anti-diarrheal, no useful data will be gleaned by combining the data from the two studies. If the two studies are both on the efficacy of the same drug, then they stand a better chance at being usefully combined, though both studies would have to measure the drug's effectiveness at treating the same condition. When the subject matter of two studies is sufficiently similar, they are said to be "homogeneous." When not, they are heterogeneous. Meta-analyses can only be informatively applied to homogeneous studies.

---

[18] See, for example, (Hyman 2010).

We can't assume that just because two studies are on the existence of psi that they are sufficiently homogeneous. Even if the two studies are both on telepathy, they might not be sufficiently homogeneous if they don't test for telepathic response in the same way. If one study uses the ganzfeld procedure while another has subjects guess what food a sender has just eaten, there is no guarantee that the two studies are testing for the same thing. Surface similarities can mask underlying differences, and perhaps even the ganzfeld and autoganzfeld experiments are not sufficiently homogeneous.

Ray Hyman (2010) argues that the various ganzfeld experiment databases covered by previous meta-analyses are not sufficiently homogeneous and that the only way to make them sufficiently homogenous is by excluding some studies in a way that distorts the meta-analyses' conclusions. For example,

> The reliance of Storm et al. (2010) on meta-analysis masks rather than uncovers the actual situation. The consistency they find and report is a manufactured one. They create "homogeneous" databases by removing outliers. This practice makes the remaining effect sizes less variable but does not change the fact that the original populations of experiments are heterogeneous. They compound this problem by justifying the combination of databases whenever a statistical test fails to reject the null hypothesis of no difference among the effect sizes. . . .

> Some of the joining of databases borders on incoherence. After making a point that the Milton–Wiseman database is an outlier with respect to other ganzfeld databases, Storm et al. (2010) do not hesitate to combine that database with their own. (Hyman 2010, 487)

Storm and Tressoldi, et al. reply:

> Claiming that our meta-analysis “masks rather than uncovers the actual situation” (p. 487), Hyman (2010) alleges that we have “manufactured” consistency. We state that removing outliers to create homogeneous databases and/or to merge databases is standard practice in meta-analysis. We were transparent in our procedures and findings, applied standard rules, and did not invent any untoward techniques (see, e.g., Rosenthal, 1984). If the “original populations of experiments are heterogeneous” (Hyman, 2010, p. 487), then so be it—surely it is a better practice to err on the side of caution than mistakenly inflate or deflate the size of an effect. . . .
>
> On whether the Milton and Wiseman (1999) database is an “outlier,” we meant it to be counted as an exception because our population estimates of z score and effect size values included zero for that database, and only that database. Even so, we put the psi hypothesis severely to test in finding that such a low-yield database (effect size _ .013) could still be legitimately merged with other databases without adversely affecting the overall significance of increasingly larger databases. (2010b, 492)

I do not evaluate which, of this response and Hyman’s critique, is more convincing. My contention is only this: many goats are not adept at statistical analysis, so won’t be able to reliably evaluate the claims of either side. Therefore, if a goat wants to maintain their belief that there is no such thing as psi and wants to do this *on the basis* of Hyman’s analysis, they will need to do it on trust. But I suspect that many goats won’t want to depend for their rejection of psi on the ultimate validity of Hyman’s objection and the failure of Storm and Tressoldi, et al.’s reply. Just as atheists won’t want to depend on the success of the specific criticisms of the fine-tuning argument, goats won’t want to depend on the success of critiques of pro-psi meta-analyses that are beyond their ken. If it turns out that the studies covered by the various meta-analyses really are homogeneous, few goats will give up and admit to the existence of psi or lower their confidence that there is no such thing.

The success of the critiques of specific pro-psi meta-analyses, while heartwarming, is inessential if what's at issue is the retention of goat knowledge that psi doesn't exist. As with the survival of atheist knowledge in the face of the fine-tuning argument, it is the goat's lack of facility with the statistical methods required to reliably evaluate meta-analyses that saves their knowledge from those meta-analyses. It might be thought that any goat unable to reliably evaluate the original meta-analyses, the criticisms, or the responses, should withhold judgment on the existence of psi and therefore lose knowledge that there is no such thing. But given the conclusions of the previous chapters, this is not what is required.

The basis of what is otherwise goat knowledge that there is no such thing as psi is a kind of felt obviousness akin to that which guides your certainty that there is no magic going on in magic shows. You still can know, prior to the experimental results, that any meta-analysis purporting to demonstrate the existence of psi is misleading. If you lack facility in evaluating meta-analyses, your finding the steps in such analyses compelling and being unable to expose a flaw wouldn't surprise you even if the meta-analyses were misleading. Therefore, your confidence doesn't have to change in the face of those meta-analyses, as long as you have sufficient facility with the basis of what would otherwise be your knowledge that there is no such thing as psi. As I've already argued, you do; you have sufficient facility with your basis if that basis is sufficient, barring relevant counterarguments, for knowledge.

This is not to say that meta-analyses can't serve as part of a good case for changing your mind or establishing facts. Meta-analyses of studies on the health effects of cigarette smoking are part of a good case that cigarette smoking causes cancer,[19] and meta-analyses that show no association between vaccinations and autism are part of a good case that there is no such association.[20] Why should these be convincing while the meta-analyses invoked by pro-psi researchers need not be? The answer is that, when it comes to the meta-analyses performed on the health effects of tobacco and the proclivity of vaccinations to cause autism, there is a recognized consensus of experts whose opinion is trustworthy that the meta-analyses are sound. It's that testimony—not the meta-analyses themselves—that is the source of lay knowledge. If there were unanimity among trustworthy experts that the pro-psi meta-analyses were sound, that might indeed constitute knowledge-destroying evidence, because lay goats have no less facility in the evaluation of expert testimony than they do in the source of their prior knowledge that psi doesn't exist.

---

[19] See, for example, (Gandini, Botteri, et al. 2008).

When it comes to goat disbelief in psi, it's not that the source of goat knowledge is expert testimony that there is no such thing. The source of goat knowledge is felt obviousness that encodes the "overwhelming" empirical support for Broad's basic limiting principles. The role of expert testimony is that, when it comes to pro-psi meta-analyses, there doesn't exist the kind of trustworthy expert consensus in favor of the meta-analyses that would serve as a defeater for what is otherwise goat knowledge.

Psi-sympathist John Palmer, commenting on what he thinks of as goats' (in his term, "conventional theorists'") anti-empiricist sentiment, says that,

> It seems to me that in a proper empiricist science, a priori conventions should not be invoked until the empirical research process has been exhausted and a winner still has not emerged. Since [conventional theorists] are not willing to hold off this long, they effectively demote empirical evidence to a secondary status in their philosophy. (1986, 35)

By "a priori conventions," Palmer has in mind principles like Broad's—principles that Broad claims to have ultimately empirical support. So it's not clear to what extent Palmer seriously thinks that the goat's motivating principles are genuinely *a priori*. More likely, he thinks that they do—but shouldn't—stand as prior constraints on the confidence goats give to the results of explicit empirical research. As we've seen, however, if goats rightly judge that they are not reliable evaluators of that research, then lacking any other convincing evidence that the research is sound, goats would be remiss not to allow their confidence to be constrained by those "a priori" conventions.

---

[20] See, for example, (Taylor, Swerdfeger, et al. 2014).

4.5 CONCLUSION: THE POSITION OF THE EXPERT

In this chapter, I have explored how amateurs might retain atheistic or goat knowledge if confronted by arguments it is beyond their ability to reliably evaluate. But some of us have expertise in the relevant areas; we are atheistic philosophers of religion or expert psi-debunkers. Are experts in danger of losing knowledge when confronted by arguments in which they may not be able to expose a flaw? I discussed this issue in the previous chapter, but it is informative to revisit the matter in light of the specific topics of this chapter. I focus on atheism but my comments should apply just as well to the denial of psychic phenomena.[21]

Suppose an expert philosopher of religion and an atheist—an expert atheist –comes across a new collection of theistic arguments. The expert has sufficient expertise to evaluate those arguments. Because philosophers are clever, the expert might not be able to figure out what is wrong with them after reading them. If the atheist were an amateur, the epistemic efficacy of amateurism might save the atheist's knowledge. But because the atheist is an expert, the epistemic efficacy of amateurism won't help. Does discovery of the new collection destroy the expert atheist's knowledge?

---

[21] My thanks to an anonymous referee for Oxford University Press for encouraging additional discussion of this issue.

The expert atheist should not be overly concerned. According to the worry, because it is surprising for experts to be fooled by misleading arguments, experts lose knowledge more easily when confronted by the existence of counterarguments they haven't read; they can't take advantage of the epistemic efficacy of amateurism. But precisely because it *is* surprising for experts to be fooled by misleading arguments, they don't lose knowledge more easily when confronted by unread counterarguments.[22] They can retain their knowledge without reading the counterarguments because it is unlikely the arguments will be apparently flawless. Expert atheists retain knowledge for the same reason that you retain knowledge in the situations described in Harman's dogmatism paradox. Knowledge is not automatically destroyed by the mere possibility of counterevidence that would be surprising were you to find it.

Once they read the arguments and find them apparently flawless, of course, expert atheists might lose knowledge in a way that amateurs don't. As argued in Chapter 3, this is the correct result. If it isn't likely that experts would find a misleading argument apparently flawless, then if expert atheists find a theistic argument apparently flawless, they should rethink their atheism. Because it isn't likely that they would be fooled by misleading arguments, expert atheists won't be in this situation—finding a counterargument apparently flawless—as often as amateur atheists. Amateur atheists will routinely find misleading arguments apparently flawless that expert atheists don't find apparently flawless. So most arguments that allow amateur atheists to retain their knowledge will also allow expert atheists to retain theirs. Expert atheists, as easily as amateur atheists, retain knowledge in the face of unpossessed evidence. Expert atheists, as easily as amateur atheists, retain knowledge after exposure to counterarguments. Only when both find counterarguments apparently flawless are experts in more danger of losing knowledge. This, again, is how it should be.

---

[22] This hearkens back to Nathan Ballantyne's "problem of unpossessed evidence," discussed in Chapter 2.

5

The Obligation to Engage

> The definitive idea for deliberative democracy is the idea of deliberation itself. When citizens deliberate, they exchange views and debate their supporting reasons concerning public political questions. They suppose that their political opinions may be revised by discussion with other citizens; and therefore these opinions are not simply a fixed outcome of their existing private or nonpolitical interests. (John Rawls (1997, 772))

5.1 INTRODUCTION

The late twentieth and early twenty-first centuries have provided unprecedented opportunities to argue with each other. There are blogs, social media, email, and texts, in addition to more traditional modes of interaction: letters to the editor, town hall meetings, snail mail, conversations with friends, phone calls to elected representatives, public debates, invited speakers, question-and-answer sessions, the center of the public square, and the corner soapbox. It is easier both to make your position widely available and for others to communicate their counterarguments and objections to you.

If you have a position on some controversial matter of politics, religion, science, ethics, or philosophy, and you have access to or are already holding forth in some public or private forum in which that matter is being discussed, what is your obligation to engage with relevant counterarguments? To what extent should you be willing to pull back from your opinion or the expression of that opinion if you can't answer those counterarguments? In this chapter I develop a variety of arguments that we *should* argue with each other; there is at least a *prima facie* violation of an obligation in all the central cases in which we refuse to engage with relevant counterarguments, including the case that opens the preface to this book.

Though I use the language of obligation, it is enough if the arguments purport to show only that there is some normative violation or indecency in failing to engage. That's because I don’t think these arguments ultimately achieve even that minimal goal. In many central cases there is neither a normative requirement to engage with relevant counterarguments, nor a requirement to pull back from your expressed opinion if you can’t answer those counterarguments. On the contrary, there is often a normative requirement not to engage and, if you engage, not to do so open-mindedly. This is true even if you have already taken an explicit stand in defense of your position and even if the counterarguments in question come from members of your intended audience. I defer until Chapters 6 and 7 the argument that there often aren’t such normative requirements and that there are often normative requirements to the contrary. (I apply the conclusion to a specific case—invitations to academic speakers—in Chapter 8.)

There are at least three sources that might generate an obligation to engage with relevant counterarguments. First, the obligation might be an essential feature of the role or position occupied by the one who has the obligation. For example, having the obligation might partly constitute any of the following roles: commenter on a public paper, explicit advocate of a position, elected representative, or participant in argumentation. Second, engagement might be obligatory because of the positive consequences of engagement, among which are the positive epistemic consequences of obtaining truth and avoiding falsehood. Third, engagement might be required in order to respect the rights or agency of those who put forth counterarguments. The obligations generated by any of these sources could amount to things you owe to society or the world at large, to yourself, to those who put forth the arguments, to some specified audience,[1] or even to no one at all, if it’s possible to have obligations whose fulfillment you don’t owe to anyone.[2]

---

[1] Though I am mainly concerned with putative obligations to those presenting arguments and to the world at large, Christian Kock thinks that the purpose of participating in argumentation is to benefit audience members listening in, but not involved in the discussion: “the main reason why such debates are potentially meaningful is that other individuals facing such a choice (legislators and citizens) may hear, consider and compare the arguments relating to the choice” (2007, 238). The only way argumentation can serve this purpose is if each side engages with the arguments made by the other side.

Depending on how narrowly “salience” is construed, in addition to any putative obligations there are to engage with salient, relevant counterarguments, there might be obligations to engage with counterarguments not yet made salient. If “salience” is construed narrowly, there might be lots of counterarguments—particularly arguments that would be advocated by those whose voices are excluded from public spaces—that are never able to make it to salience. Focusing exclusively on putative obligations to engage with salient counterarguments, on this narrow construal, inevitably leaves out of the discussion the voices of the marginalized and excluded. To avoid this, I include in the domain of relevant counterarguments, not only counterarguments that have made it to public attention, but counterarguments that would be endorsed or advocated by members of disenfranchised groups or their representatives.

If there is an obligation to engage with relevant counterarguments, it is likely merely *pro tanto*. In David Copp’s (2007, 370) construal, you have a *pro tanto* obligation to engage just when you have a defeasible reason for engaging that, in the absence of conflicting reasons, would suffice to establish an all-things-considered obligation. Your obligations to engage with relevant counterarguments can often be defeated by considerations of time, interest, potential payoff, and qualifications. Still, defenders of engagement take pains to comment on what kinds of considerations the *pro tanto* obligation defeats. As we saw in the preface, Barack Obama thinks the obligation defeats considerations of offensiveness, sexism, and racism. Mill thinks it defeats considerations involving the certainty of the received opinion. Trudy Govier thinks it defeats considerations of frivolousness or silliness. If they're right, the *pro tanto* obligation to engage with relevant counterarguments is a *pro tanto* obligation with significant clout.

---

[2] Some of these sources might generate obligations to engage with relevant counterarguments only in some domains. For example, if the obligation has its source in whatever it is that makes exercises of democratic power politically legitimate, as the arguments I attribute to Rawls and Gutmann & Thompson might maintain, then there might only be an obligation to engage with relevant counterarguments in public domains. Of course, the domain properly classified as “public” is itself a matter of debate. As Seyla Benhabib says, in developing her Habermasian view of public discourse, “In effect, there may be as many publics as there are controversial general debates about the validity of norms” (1992, 105). I take no stance on this issue. As long as the arguments in this chapter establish an obligation to engage with relevant counterarguments in any domain, that’s enough.

5.2 IS THE OBLIGATION A ROLE OUGHT?

According to what is known as "role ethics" some obligations accrue to us in virtue of our occupying certain roles.[3] Teachers thereby have obligations to come to class, have a lecture ready, be prepared to take certain kinds of questions, and grade papers in a timely manner. Parents might thereby have obligations to vaccinate their children and provide them with healthy food and opportunities for play. Richard Feldman (2001, 87) calls these obligations "role oughts" in honor of the roles that generate them.

We might ground role ethics, as some scholars of Confucianism do, in a metaphysics of persons according to which persons are constituted by their relations to others.[4] We might also ground something close enough to role ethics in contracts between you and those you have obligations to. Most obviously, you might have obligations that have their source in explicit legal contracts. But if there are implicit social contracts, then they might generate obligations as well. I classify such obligations as role oughts, since their source lies not in the potential consequences of your actions nor in the intrinsic rights of your interlocutor or audience; they are generated by your standing as a party to the contract.

---

[3] Some (e.g. (Ames 2011)) see role ethics advocated by Confucius. Brian Johnson (2014) sees role ethics as central to understanding Epictetus

[4] See, e.g., (Bockover 2012, 180). For an overview of the issues surrounding Confucian role ethics, see (Ramsey 2016). Alternatively, we might think that part of what it is to occupy a role is to be the subject of various obligations (see, e.g., (Dare 2016, 704) and (Hardimon 1994, 334)). Anne Baril (2016, 725) calls the relevant kind of role a "normative role."

Some social roles naturally suggest themselves as generative of an obligation to engage with at least specific relevant counterarguments. If you are a formal commenter on a paper, you are obligated to engage seriously with the paper you are commenting on. Invited speakers have obligations to respectfully respond to questions from the audience during the regular question-and-answer session. Professors commenting on student work should go into depth on the students' arguments. When you participate in a public debate, you are rightly expected not to ignore the arguments of your fellow debaters. Failure to engage while occupying any of these roles seems to involve an abdication of the duties of the position.

What is more interesting is the possibility that you inherit an obligation to engage in virtue of occupying the social role of arguer. Certain ways of generating role oughts are not obviously transferable to the role of arguer. For example, Michael Hardimon limits his "role obligations" to "*institutional* roles—ones that are institutionally defined" (Hardimon 1994, 334). Though Hardimon casts the net of institutional roles widely, to include "political, familial, and occupational roles" (334), it may seem strained to include the role of arguer as an institutional role.

One way this obligation could be explained is by a general moral theory—like Iris Marion Young's (2006) "social connection model"—according to which our social roles generate obligations to those we affect or whose activities we presuppose when acting in those roles.[5] Regarding the latter, Onora O'Neill says that

> when agents *commit* themselves to the assumption that there are certain others, who are agents or subjects with these or those capacities, capabilities and vulnerabilities, they

---

[5] Young sees her social connection model as exploiting the "great insight" of social contract theory: that "social connection is prior to political institutions" (Young 2006, 105).

> cannot coherently deny these assumptions in working out the scope of ethical consideration to which they are committed. Commitments to others' ethical standing are taken on as soon as activity is planned or begun . . . (1996, 100)

The trick is to figure out what obligations are generated by which social roles. As O'Neill immediately goes on to say, "[W]hat is needed is a procedure for working out what these commitments are in a given context" (100). Why, in particular, does occupying the social role of arguer commit you to engaging with counterarguments?

Much of the work on this question has been done by those, like Ralph Johnson, who work in informal logic. Here are Johnson's comments on a remark of Martin Luther King's:

> Martin Luther King wrote: "Seldom do I pause to answer criticism of my work and ideas. If I sought to answer all the criticisms that cross my desk, my secretaries would have little time for anything other than such correspondence in the course of the day, and I would have no time for constructive work." . . . *I have argued that anyone who enters into argumentative space in effect makes a commitment to engage in such as part of the very work of argument*. (2001, 3 emphasis added)

When you "enter into argumentative space" you "make a commitment to engage" in the same way that when you take a teaching position you make a commitment to come to class with material prepared. Part of the reason has to do with goals you take on when you participate in argumentation:

> If someone raises an objection to your argument or makes a criticism of it, and you [sic] goal is to persuade your audience rationally to accept your thesis, it does not seem that you can hope to accomplish it without dealing with that objection or criticism." (2001, 2)

Similarly, Govier says, "If the arguer wants to rationally persuade his audience, he or she will have to address those objections that are influential, for whatever reason" (1997).

It might seem like Johnson and Govier are arguing for a merely rhetorical obligation to engage with relevant counterarguments. To the extent you have the goal of convincing your audience, you will be less successful if you don't engage with their concerns. But elsewhere Johnson is more careful to make clear that the goal of convincing others is not external to the practice of argumentation; it's partly constitutive of an argument that it be directed toward convincing others. (Govier may commit to a similar claim.[6]) If the point of a series of claims isn't to convince others, it's not an argument:

> [A]rgumentation has the structure that it does because of its purpose. An argument seeks to bring about rational assent, and because the issue to be dealt with is contentious, controversial, it follows that one must provide reasons! Mere assertion can't get the job done rationally. Hence we find that the purpose of argument as rational persuasion dictates the infrastructure of argument: premises and conclusion . . .
>
> The arguer must also address himself to the opposing points of view and show why his is superior. To fail to do this is to fail to discharge a fundamental obligation of the arguer in the dialectical situation. (Johnson 2014b, 78)

Jennifer Faust goes so far as to call this view of the purpose of argumentation, the "standard view":

> The primary purpose of an argument, understood in the philosophically orthodox sense, is to persuade someone of the truth of its conclusion. This aim of argumentation is so obviously and widely recognized that it is often written into the very definition of the term. . . . On this view (hereafter called "the standard view"), an argument's primary purpose is to persuade an audience to accept its conclusion. (2008, 74)

To the extent you offer something that doesn't have the purpose of rational persuasion, you're not participating in argumentation and haven't offered an argument. If this "standard view" is right, according to Johnson, arguers inherit obligations in virtue of the nature of their enterprise—one that essentially has the purpose of rational persuasion. In this, Johnson follows Hardimon: the content of role oughts are "fixed by the function of the role" (Hardimon 1994, 334).

If Johnson is right, the purpose of argumentation is rational persuasion. But rational persuasion is only achieved if the arguer deals with objections and criticisms; unless you engage with relevant counterarguments, your audience

---

[6] Govier says, "Given that the purpose of argument is rational persuasion . . ." (Govier 1997), indicating that she takes for granted that rational persuasion is indeed the purpose of argumentation. It's not entirely clear whether this is Govier's position or whether she is operating within the assumptions of Johnson's system.

> will not easily be won over, nor should they be, if they are rational. For they will know that there are objections to the arguer's position . . . They will want to know how the arguer intends to deal with this objection and that criticism. They will rightly not be persuaded until they are satisfied that the arguer can handle these objections. (Johnson 1999)

Therefore, arguers—in virtue of their participation in argumentation—inherit an obligation to engage with relevant counterarguments. Simplifying Johnson's argument:

1) Arguments and argumentation essentially have the purpose of achieving rational assent.

2) Rational assent is only achieved if the arguer engages with relevant counterarguments.

Therefore,

3) Arguers have an obligation to engage with relevant counterarguments.

Johnson draws a further conclusion: that arguers succeed in participating in argumentation only if they engage with relevant counterarguments. Much criticism of Johnson is devoted to this further conclusion, but even if it doesn't follow, Johnson may succeed in showing that arguers have an obligation—sourced in the practice of argumentation itself—to engage with relevant counterarguments.

Notice that the relevant counterarguments need not be plausible or objectively justification-conferring; they don't have to be good arguments. As Govier says, "Even if an arguer thinks that an objection is entirely silly and frivolous, he or she should respond to it if the audience seems to find it important" (1997). Johnson requires that you answer all of your audience's "serious" objections—plausible or not.[7] For an objection to be "serious", for Johnson, it must satisfy three conditions (Johnson 1999). It has to be "supported by a line of argument," though there is no requirement that the line of argument be plausible. It must be "directed to a crucial premise," which also includes no requirement of plausibility. And it must be a "strong" objection. This third condition might seem to require that the objection be plausible, until we see what "strength" of an objection amounts to. Following Govier,[8] Johnson construes strong objections as "'those which allege that the defective feature indicates that the case is false, wrong, incorrect'" (Johnson 1999, quoting Govier)[9], as opposed to objections that merely suggest "that the position may need qualification" (Johnson 1999). There is no requirement of plausibility here, either.

---

[7] This doesn't commit Johnson to arguing that you are automatically culpable for failing in this duty. You might be excused for reasons of time or priority.

Nor should there be, because the purpose of argumentation is not simply to garner rational assent from people who offer good arguments. In any case, the additional requirement that you engage not only with plausible relevant counterarguments, but any seriously intended counterarguments should not impose an especially onerous burden, given that if the counterarguments are so implausible, you arguably should be able to say so and say why.[10]

One question that arises from Johnson's discussion is the degree to which the obligation is necessary. Is it necessary or contingent that arguers have an obligation to engage with relevant counterarguments? If the obligation stems from the psychological fact that humans are only convinced when their concerns are engaged with, then the obligation is no more necessary than that psychological fact. If humans tend to be more convinced when their concerns are ignored, then to the extent an argument has the purpose of convincing others, participation in argumentation would come with an obligation to ignore the objections and counterarguments of your audience. And let's not fool ourselves into thinking that this possibility is incredibly remote. All it takes for humans to have this psychological trait is for them to think that dismissive and over-confident speakers must be on to something[11]—for them to have the attitude that if someone is so confident of themselves that they don't have to deign to engage with relevant counterarguments, then they must have powerful evidence indeed. In such a possible world, premise 2 of Johnson's argument is arguably false: assent is only achieved if arguers fail to engage with relevant counterarguments.

---

[8] Govier requires that

> an arguer . . . reply to all *dialectically significant* objections that have been raised against his or her position. What makes an objection *significant* in a dialectical sense? One factor is that the objection is logically serious; if it were to hold, the position would be refuted or the argument would be shown worthless. Another is that it is taken seriously, or seems, so far as the arguer can tell, to be taken seriously by the audience to whom the argument is addressed. (1997)

There is no requirement, here, of objective plausibility. Govier notes in an endnote (n. 25) potential difficulties with taking a more "objectivist" direction.

[9] Johnson attributes this passage to "a 1998 paper" by Govier, though he doesn't give the bibliographic information and Govier has no publications in 1998 in which the passage appears. The quote appears at (Govier 1999, 231).

[10] Arguably, but falsely; as we have seen in the first part of this book, arguments can be obviously misleading even if you can't say what's wrong with them.

But Johnson's aim of argumentation is not merely to produce assent, but rational assent. This is emphasized by David Hitchcock, who says in his commentary on Johnson,

> I therefore propose as the overall expected purpose of participants in the practice of argumentation that of arriving at a shared rationally supported position on an issue. The reference to rational support makes sense of the consideration of objections, openness to criticism and willingness to revise arguments which Johnson rightly cherishes as hallmarks of the practice. (2002, 291)

Rational assent is only achieved through reasoning—not, say, by browbeating your audience into submission. If to achieve rational assent it is necessary that you have satisfactory answers to your sincere concerns, then audiences will not rationally assent when their concerns are ignored, even if they assent when their concerns are ignored. Premise 2 will remain true, and necessarily so: the only way to achieve *rational* assent is to seriously address your audience's concerns—to engage with relevant counterarguments.

A further worry about Johnson's argument concerns the degree to which participation in a certain practice provides an obligation to fulfill purposes essential to that practice. Why does the fact that the purpose of an argument is to achieve rational assent mean that arguers have some sort of obligation to fulfill that purpose? The purpose of thievery is to steal, but thieves don’t thereby have an obligation to steal. Murderers murder, but they don’t thereby satisfy obligations by doing so. If purposes of activities don’t automatically generate obligations to fulfill those purposes when engaged in those activities, then Johnson’s conclusion doesn’t follow from his premises.

---

[11] At the time of this writing Donald Trump is the President of the United States.

In response, Johnson downgrades the kind of obligation he thinks arguers have to engage with relevant counterarguments: “The arguer’s obligations are conditional rather than categorical. The arguer must respond to criticisms, *if* the arguer wishes to remain a participant in the process” (Johnson 2001, 2). Here we run into familiar complications about wide-scope and narrow-scope obligations. Do you have Johnson’s conditional obligation just in case you have an obligation to make true the following conditional: *if you make an argument, you engage with relevant counterarguments*? Or do you have Johnson’s conditional obligation just in case the antecedent of the following conditional is true of you: *if you make an argument, you have an obligation to engage with relevant counterarguments*?

The problem with the latter is that because the antecedent of the conditional is discharged for arguers, arguers end up with a categorical—and not just a conditional—obligation to engage with relevant counterarguments. And if the obligation is grounded in the purpose of the activity, then what goes for arguers, so goes for thieves and murderers: thieves end up with categorical obligations to steal and murderers end up categorically obligated to kill. On the other hand, the difficulty with making the obligation wide in scope—with making it that to have a conditional obligation to engage with relevant counterarguments is to be such that you have a categorical obligation to make a conditional true—is that arguers acquire no new obligations in virtue of being arguers. We all have the obligation to make the conditional true, but those of us who aren't making an argument trivially satisfy the obligation because the antecedent of the obliged conditional is false of us.

Still, this latter course seems the best for Johnson to take, and is in keeping with other views that attempt to ground obligations in what's essential to practices. For example, Timothy Williamson advocates the following *rule of assertion*:

> (The knowledge rule) One must: assert *p* only if one knows *p*. (Williamson 2000, 243)

The rule expresses a conditional obligation insofar as it is an instance of a general form of rule which

> is to be parsed as 'One must ((assert *p*) only if *p* has C)', with 'only if *p* has C' inside the scope of 'One must' but outside that of 'assert'. The rule unconditionally forbids this combination: one asserts *p* when *p* lacks C. (241)

Johnson's *rule of argumentation* can be put in similar form:

> (The dialectical rule) One must: argue that *p* only if one engages with relevant counterarguments regarding *p* and one's arguments for *p*.[12]

Williamson's rule of assertion is what he calls a "constitutive rule." He illustrates constitutive rules with analogy to rules of games.[13] It's a constitutive rule of chess that bishops must move diagonally, so if you are playing chess and move a bishop in some other direction, you violate some sort of obligation. What sort of obligation is grounded by a constitutive rule? Williamson says,

> The normativity of a constitutive rule is not moral or teleological. The criticism that one has broken a rule of a speech act is no more a moral criticism than is the criticism that

---

[12] The dialectical rule may lead to trouble if conjoined with Johnson's further claim that you do not succeed in argumentation unless you engage with relevant counterarguments. Williamson argues that if a rule forbids a certain combination, that combination had better be possible "otherwise it would be pointless to forbid it" (241).

[13] The analogy with games is not perfect because, as Williamson points out, "in the ordinary sense of 'game', games such as tennis gradually change their rules over time without losing their identity; the constitute role of the rules is qualified by that of causal continuity." In the sense in which Williamson means the term, "a rule will count as constitutive of an act only if it is essential to that act: necessarily, the rule governs every performance of the act" (2000, 239).

> one has broken a rule of a game or language. . . . Nor is the criticism that one has broken a constitutive rule of an institution the criticism that one has used it in a way incompatible with its aim, whether the aim is internal or external. . . . Breaking the rules can serve both internal and external aims. Within the practice constituted by the rules, their authority does not require the backing of moral or teleological considerations. (241)

This passage tells us what the normativity of constitutive rules doesn't amount to: moral or teleological criticizability. It doesn't say what it does amount to. But there is no question that, according to Williamson, constitutive rules are normative; they mark out genuine obligations. For in breaking a constitutive rule, "one is subject to criticism precisely because one has performed an act for which the rule is constitutive" (Williamson 2000, 240).

Johnson's rule of argumentation is not quite a constitutive rule in Williamson's sense, even though both constitutive rules and Johnson's rule are grounded in features essential to the practice governed by the rule. For one thing, though Williamson, as noted, does not construe the normativity provided by constitutive rules as teleological, it is open to Johnson to do so, because adherence to the rule is what best fulfills the essential goal of argumentation—achieving rational assent. The reason why Johnson is able to construe the normativity this way, though Williamson is not, is because, for Williamson, constitutive rules are brute—explanatorily basic: "It is pointless to ask why the knowledge rule is the rule of assertion. It could not have been otherwise" (Williamson 2000, 267). In contrast, Johnson gives an explanation for why his rule of argumentation is the right rule: it's because argumentation essentially has a certain purpose—achieving rational assent.

In this way, Johnson's argument for his rule is more like arguments for norms of belief that derive those norms from the aims that essentially constitute belief—primary among which is the aim of truth.[14] For example, David Velleman says

> Let me be clearer about the relation between the constitutive aim of belief and the norm that applies to belief in light of that aim. To say that belief aims at the truth is not simply to re-express the norm stipulating that a belief must be true in order to be correct; rather, it is to point out a fact about belief that generates this norm for its correctness. (2000, 16–17)

Still, Williamson's discussion of constitutive norms opens up the possibility that Johnson's rule might be argued for in some other fashion than by the essential purpose of argumentation. This is fortunate, because there are potential counterexamples to the claim that argumentation has an essential purpose.[15] Hitchcock, for example, considers the worry that Johnson's definition

> attributes to participants in argumentative discussions a purity of intention which real-life participants . . . often do not possess. People who write letters to the editor, phone radio talk shows, send messages to Internet discussion groups and join in panel discussions are often more concerned to assert their own position than to engage in a back-and-forth interchange which might lead them to change their mind. Referees may push their own view rather than suggest ways of strengthening the argument in a manuscript. Scholars and scientists may cling stubbornly to a favoured view long after it has been decisively refuted. Similarly for judicial and quasi-judicial opinions, discussion among family and friends, and decision-focussed discussions at meetings. (2002, 291)

---

[14] Some authors (e.g. (Engel 2005, esp. 78, 82, and 90) and (Wedgwood 2002)) construe the aim of belief normatively, as the condition that must be met for a belief to be correct or proper. Construed this way, arguments from the aim of belief are closer to Williamson's argument than to Johnson's.

[15] An anonymous referee for Oxford University Press suggests—rightly—that, in many cases, philosophers participate in public argumentation, not to achieve rational assent, but as part of the activity of exploring logical space and seeing what reasons can be invoked in favor of one side or another of an issue. Josh Stein points out that you might participate in argumentation to express support for some marginalized group.

The range of argumentative purposes might be broader still. Katharina Stevens argues that "the good of argumentation is the bettering of belief-systems—furthering of knowledge, extension of justificatory inferences, gaining of information and understanding etc." (2016, 377). While rationally persuading others might be one way to improve the belief systems of your interlocutors, it is not the only way. Stevens, extending Douglas Walton and Erik Krabbe's (1995) discussion of the purposes of various dialogue types, wants to allow a wide variety of ways that argumentation can improve the belief systems of the participants. For example, following Walton and Krabbe, "Quarrels enrich the parties' understanding of each other's emotional landscapes. This, too, fits the good of bettering belief-systems" (Stevens 2016, 377 n. 378).

Hitchcock thinks Johnson might plausibly deny that the apparent counterexamples are genuine, by insisting that, say, "[e]go-involved or partisan contributors to a spoken or written exchange may pretend to be involved in an argumentative discussion, but they are really not discussing. They are debating" (2002, 292).[16] Johnson himself responds to this worry by distinguishing between central and penumbral cases of argument (2000, 168–169) and claiming that argumentation whose purpose is to achieve rational assent is fundamental: "if there is a multiplicity of purposes, that does not preclude there being some purpose that is fundamental, and that is the position I have taken here. The purpose of rational persuasion is fundamental" (2000, 174).

But is it? Frans H. van Eemeren and Peter Houtlosser note that,

> In a reconstruction of a discourse as a manifestation of a critical discussion it is assumed that the arguers aim to resolve their dispute on the merits. At the same time, however, it may be assumed that they will be intent on having their own standpoints accepted. This means that on the one hand they are committed to objectives and obligations of a dialectical kind, and on the other hand they have aims and considerations that are to be understood rhetorically. While dialectical obligations have to do with the argumentative procedures that further an abstract ideal of rationality in critical discussion, rhetorical considerations are more pragmatic and relate to the contextual adjustment of argumentation to the people who are to be convinced and their discursive standards. (2001, 3)

Arguers both try to achieve rational assent and try to get their views accepted. Johnson thinks that the former goal generates obligations, but the latter goal doesn't (you don't, for example, have an obligation to gloss over weaknesses in your argument even if doing so will achieve your rhetorical purposes). He says that this is so because the former purpose—the dialectical purpose—is more fundamental. But why? Johnson can't say that the dialectical purpose is more fundamental because it—and not the rhetorical purpose—generates obligations for the arguer. The order of explanation is supposed to go the other way around.

---

[16] Whether Johnson can plausibly adopt this strategy, Hitchcock thinks, "depends on the existence of a certain practice and the recognition within it of the norms in his definition." (2002 293) But, Hitchcock continues, fortunately for Johnson, "The practice of argumentative discussion does seem to exist, in our culture and in others; its participants seem to recognize as a norm the goal of arriving at a shared rationally supported position on an issue." (2002 293)

The debate about which argumentative purposes are most fundamental is avoided if it is simply a rule of the game of argumentation that you should engage with relevant counterarguments—if the rule of the game does not derive from the fact that argumentation comes with a certain purpose. This strategy requires a change of argument, since Johnson's argument for his rule—making essential reference as it does to the essential purpose of argumentation—won't do. Williamson offers two kinds of argument for his knowledge rule of assertion: arguments based in cases, and arguments based in linguistic data.[17] Only the case-based arguments are easily adapted for Johnson's rule of argumentation, so I focus on those.

The cases Williamson emphasizes involve lottery propositions. Because you can't know that your ticket in some indefinitely large lottery is a loser, the knowledge rule of assertion rightly (according to Williamson) predicts that you can't properly outright assert that your ticket is a loser. Case-based arguments should work the same way for rules of argumentation. Imagine a case in which an arguer fails to engage with some relevant counterargument. If, intuitively, their participation in argumentation is improper, that is evidence for Johnson's dialectical rule of argumentation.

In his (2008) and his (1999), Johnson presents an example "where I think it is clear that the arguer failed to satisfy his dialectical obligation" (2008, 151).[18] In a "private email conversation" (Johnson 1999), Johnson claims, John Searle wrote,

---

[17] The first of Williamson's linguistic arguments notes that it is often appropriate to challenge assertions by saying, "How do you know?" or "Do you know that?" The second makes use of a modified version of Moore's paradox, but involving knowledge rather than belief.

[18] In his (1999) the example appears in a section entitled, "How Not to Deal with Your Dialectical Obligations: Searle's Response."

> I did not think that [L. Jonathan] Cohen's article was worth answering directly. I answered indirectly in an article you have obviously not read called "Austin on Locutionary and Illocutionary Acts."

But in that paper, Searle's response to Cohen is wholly confined to a footnote in which Searle writes, "Cohen unfortunately seems to conclude that there are no such things as illocutionary forces. This conclusion seems unwarranted" (Searle 1968, 408). About this response, Johnson says,

> Full-stop, that's it. Searle has perhaps started the task of meeting his dialectical obligation here, it seems clear that his response falls short of being adequate. We want to know what reasons Searle has for his claim that the conclusion is unwarranted. (2008, 3)

In drawing this conclusion, Johnson does not try to first establish that Cohen's response to Searle is plausible. He does—especially in his (1999)—try to establish that Cohen's response is "serious" in the sense outlined above—that it is based in argument, directed against a crucial premise of Searle's argument, and alleges that the argument is wrong rather than just in need of qualification. But none of this requires that Cohen's objections be plausible. Therefore, if Johnson's intuitions about this case are widely shared, this amounts to a good reason to buy into his rule of argumentation: arguers don't have obligations to engage only with plausible counterarguments. They have obligations to engage with serious counterarguments. This establishes Johnson's rule the way we've been construing it: arguers have an obligation to engage with relevant counterarguments.[19]

The intuition Johnson has about this kind of case is not quite the intuition that Williamson has about the lottery case. About the lottery case, Williamson has the intuition that the assertion, “My ticket is a loser,” is improper. The intuition is validated if the knowledge norm of assertion is true and it is unknown that the ticket is a loser. A parallel intuition in support of the dialectical rule of argumentation would concern a case in which the arguer fails to engage with relevant counterarguments and about which there is an intuition that participation in argumentation is improper. But that’s not Johnson’s intuition in the Searle case. Johnson’s intuition in the Searle case is more direct: it’s that Searle violated an obligation in failing to engage with relevant counterarguments, not that Searle violated an obligation in participating in argumentation in the first place.

---

[19] Johnson distinguishes seriousness from salience, noting that "[S]alience refers not so much to the strength of the criticism as to its prominence in argumentative space. It seems quite possible that an objection might achieve salience, even though it is not a strong objection" (1999). But given what Johnson means by "strong"—that is, indicating that the argument is wrong rather than simply in need of qualification—the relevant contrast is not between strong objection and salient objection, but between salient strong objection and salient mild objection.

The reason this matters is that Johnson’s intuition—even if it provides evidence for an obligation to engage with relevant counterarguments—provides at best only weak evidence about what the source of that obligation is. It’s consistent, for example, with it being the case that there is a duty to engage with relevant counterarguments even if you are not participating in argumentation but merely taking a formal stand on some matter of controversy. But if the intuition is that the participation in argumentation is improper when there is a failure to engage, then that suggests to a greater extent that the obligation to engage has its source in some essential feature of what it is to participate in argumentation. The kind of intuition Williamson marshals regarding the knowledge rule of assertion does better at pointing to the potential source of that obligation: constitutive features of the act of assertion itself.

This does not mean that similar examples can’t be employed to marshal the appropriate kinds of intuitions. We might, for example, not only have the intuition that Searle violated an obligation in failing to engage with Cohen’s arguments, but that if Searle wasn’t going to engage with Cohen’s arguments, Searle shouldn’t have been arguing for his view in the first place. The difficulty is that this intuition isn’t nearly as strong as the basic intuition that Searle’s response to

Cohen is improper. This suggests that the obligation to engage with relevant counterarguments is not constitutive of the act of argumentation in the way Williamson claims an obligation to know that *p* is partly constitutive of the act of asserting that *p*. If Johnson's original argument relying on the essential purpose of argumentation is worrisome because of its restrictions on potential argumentative purposes, defenders of an obligation of engagement will want to look elsewhere for the source of the obligation.

We should have expected that we'd need to look elsewhere, in any case. Both Johnson and Williamson-style constitutivism have to deal with the fact that disputants can be participating in two distinct kinds of practices: argumentation, a practice defined by its dialectical purpose or its constitutive dialectical rules, and another practice, sophistry, defined by its rhetorical purpose or its constitutive rhetorical rules. Clearly, both practices are relatively common. To the extent you should participate in argumentation, Johnson and Williamson-style constitutivism will say you shouldn't gloss over weaknesses in your argument. But to the extent you should participate in sophistry, you ought to gloss over weaknesses in your argument (if that will be rhetorically effective). We want to know not just what is central to each practice; we want to know which practice we should be participating in and encouraging.

While Williamson claims—as we saw—that "it is pointless to ask why the knowledge rule is the rule of assertion," he goes on the say that "[i]t is, however, pointful to ask why we have such a speech act as assertion in our repertoire" (267). We can ask similar normative questions about argumentation and sophistry. The larger normative question will be which practice we should encourage. The answer, it seems, is that we should be encouraging the practice of argumentation over sophistry. Why should we do this? Johnson's arguments about what's essential to argumentation won't tell us. But the arguments in the next two sections of this chapter might.

5.3 THE DUTY TO MAXIMIZE GOOD EPISTEMIC CONSEQUENCES

If there is an obligation to engage with relevant counterarguments, that obligation seems importantly connected to the obligation to allow expression of opinion, though the connection might be spelled out in different ways. One potential connection that doesn't look promising is that an obligation of engagement is entailed by the free speech rights of those putting forth the counterarguments. As Caroline West points out,[20]

> Although there is a sense in which an unreceptive audience 'prevents' a speaker from communicating his or her idea, this sort of prevention of communication does not amount to a violation of the speaker's right to free speech. It should not amount to a rights-violation because . . . to allow it as such would seem to license the imposition of a range of demanding duties on individuals to attend to other speakers' ideas: duties not to turn off one's television set, or avoid demonstrations, or block one's ears, or walk out of movies or public lectures, and the like. (2003, 407)

Still, a case can be made for an obligation to engage with relevant counterarguments based on similar considerations to those used to argue for universal freedom of expression. In this section I discuss the kind of argument—based on the epistemically good consequences of unrestricted freedom of expression—invoked by Mill. In the next section I discuss the kind of argument for universal freedom of expression that focuses on the intrinsic rights of the speaker.

It might seem odd to use the language of "obligation" when discussing Mill, though Mill uses normatively loaded language, for example in describing the silencing of expression as an "evil":

---

[20] West follows feminists like Catharine MacKinnon (1996), Jennifer Hornsby (1993), and Rae Langton ((1999) and (Hornsby and Langton 1998)) in arguing that pornography can be prohibited on free speech grounds because of its silencing effect on women. But they agree even with their opponents—like Ronald Dworkin (1993, 6)—that no part of a right to free speech is a corresponding right to be heard and understood.

> The peculiar evil of silencing the expression of an opinion is, that it is robbing the human race; posterity as well as the existing generation; those who dissent from the opinion, still more than those who hold it. If the opinion is right, they are deprived of the opportunity of exchanging error for truth: if wrong, they lose, what is almost as great a benefit, the clearer perception and livelier impression of truth, produced by its collision with error. (1859/2003, 87)

Suppressing opinion assumes infallibility and reduces the chance we will get truth. Perhaps as important, it risks true received opinion degenerating into prejudice and platitude:

> [E]ven if the received opinion be not only true, but the whole truth; unless it is suffered to be, and actually is, vigorously and earnestly contested, it will, by most of those who receive it, be held in the manner of a prejudice, with little comprehension or feeling of its rational grounds. And not only this, but . . . the meaning of the doctrine itself will be in danger of being lost, or enfeebled, and deprived of its vital effect on the character and conduct: the dogma becoming a mere formal profession, inefficacious for good, but cumbering the ground, and preventing the growth of any real and heartfelt conviction, from reason or personal experience. (1859/2003, 118)

Mill says that the good of being challenged even in your obviously true beliefs prevents those beliefs from becoming “prejudice” and makes you attend to the reasons available for those beliefs. Michael Scriven sounds a similar note: “[E]ven if one believes that the old values are the best values, the reasons for—and against—them must be rehearsed by each generation or it will rebel against them” (1984, 10). Commenting on Scriven, Martin Eger observes that, “this can be done only through frank encounter with unorthodox alternatives”—“‘reasonableness’ demands the *radically* open mind” (1988, 295).

Scriven’s worry is somewhat different from Mill’s. Mill worries that without examining our true opinions they become prejudice. Scriven worries that without examining our true opinions we may “rebel against them”—that is, give them up. Neither good is achieved if the helpful expression is nowhere engaged with. If you have an obligation to allow expression of opinion, and if the source of that obligation is the good that expression visits on the opinions of yourself and others, then it's hard to see how the obligation doesn't come with some obligation to engage, lest the goods of expression be wasted.

At times, Mill indicates that what's really important is that we not prevent others from listening to recalcitrant opinion. He says,

> it is not the feeling sure of a doctrine (be it what it may) which I call an assumption of infallibility. It is the undertaking to decide the question *for others*, without allowing them to hear what can be said on the contrary side. (1859/2003, 93)

But Mill rightly isn't comfortable with a situation in which, though no one prevented others from hearing an expressed opinion, all adherents to the received opinion outright dismissed the speech—refused to even hear it out:

> Truth has no chance but in proportion as every side of it, every opinion which embodies even a fraction of the truth, not only finds advocates, but is so advocated as to be listened to. (1859/2003, 88)

As Jennifer Hornsby and Rae Langton put it, commenting on this passage, "Mill is saying that realization of the good that free speech can bring requires attentive, 'listening' audiences" (1998, 34). The source of the obligation to allow opinions to be expressed is a value that is only realized if some people engage with the expressed opinion. If this value grounds an obligation to allow free expression of opinion, it grounds an obligation to engage with the expressed opinion.

The value of free expression does not come automatically; it doesn't come just from the effects that words induce at our sensory surfaces. What it takes to satisfy the obligation depends on whether the source of the obligation to engage with a relevant counterargument is the value accrued by you—the individual who directly engages—or by the group you are a member of. If the source of the obligation is the value accrued by you, then if you engage dismissively—closed-mindedly—you do not reap the benefits. You have to engage, as Mill puts it, "vigorously and earnestly," so that you're willing to change your attitudes to your beliefs in response to those opinions, perhaps by being willing to reduce your confidence in your contrary opinion. In this way, you give some power to the counterargument to change you. In Thomas Nagel's words (favorably commenting on a view of Peirce's),

> The only way we can have any hope of advancing toward the truth is to be continually dissatisfied with our opinions, to be always on the lookout for objections, and to be prepared to drop or alter our theories whenever counterevidence, counterarguments, or better-supported alternatives present themselves. Only the willingness to change one's mind gives any ground for thinking that what one hasn't been persuaded to change one's mind about may be right, or at least on the right track. (2001, 128)

Willingness to change one's mind is cheap. As noted in Chapter 1, the most steadfast dogmatist might truthfully claim to be willing to change their mind. That's because willingness to change your mind is always conditional. If it weren't, then in being willing to change your mind, you'd already have done it. When you claim to be willing to change your mind, you're claiming to be willing to change your mind if some conditions hold. What conditions? The conditions laid out in Chapter 1: what the injunction to be open-minded requires is a willingness to reduce your confidence if, upon spending some significant time with a relevant counterargument, you can't expose a flaw and all the individual steps seem compelling.

Mill wants more out of engagement than the reduction of confidence in the face of apparently flawless counterarguments. If you find the counterargument unconvincing you still have to be willing to change. You have to be willing to allow obviously misleading counterarguments to enliven and enrich your received opinion. This can happen when you base your belief on the wrong reasons—say, the testimony of the powerful rather than good, direct evidence. In that case, engagement with counterarguments, ideally, will force you to base your belief on better reasons. If Millian considerations about the value conferred on individual engagers demonstrate that there is an obligation to engage, then, it is an obligation to engage with counterarguments open-mindedly, in that you are willing to reduce your confidence in response to apparently flawless counterarguments. And it is an obligation to be willing to reconsider your reasons in response even to misleading arguments. You have to be willing to be *vulnerable* to counterarguments in both these ways.

If individuals who engage with relevant counterarguments are willing to be vulnerable, they can radiate the fruits of that open-minded engagement to the rest of the population. But engagement can be valuable even if those who engage do so closed-mindedly. If you engage with a counterargument publicly, the public witnesses you presenting objections to the counterargument, responding to evidence against the received opinion, presenting positive arguments for your position, etc. Even if you do this closed-mindedly—without a willingness to reduce your confidence if, ultimately, you cannot expose a flaw in the counterargument—the public can evaluate your response and be willing to adjust their confidence depending on how the argument goes and can be willing to marshal new reasons for their opinion depending on how effective your responses are.

It is the failure to engage even closed-mindedly that is behind many dissident critiques of mainstream opinion-makers. For example, one pro-psi website complains that

> [w]hen a large portion of mainstream science refuses to seriously consider [pro-psi] evidence, often denigrating it without reading beyond a paper's short abstract, it's no wonder that the majority of humanity have no idea that such evidence even exists. (Subtle Energy Sciences Blog 2015)

Even closed-minded but serious consideration of pro-psi evidence could raise awareness of that evidence and allow the public to decide for themselves which case is more convincing; this is all psi-researchers claim to want. If Mill's reasoning is sound, then at the very least you have an obligation to engage closed-mindedly with relevant counterarguments, and perhaps have an obligation to engage open-mindedly with relevant counterarguments.

One epistemic good that arguably is suggested by the Millian defense of engagement with relevant counterarguments is the prevention of so-called “group polarization” and contexts that produce it—sometimes called “echo chambers.” Group polarization occurs when the average opinion or judgment of a group moves closer to the extreme of whatever direction the group’s average opinion or judgment was leaning in the first place. If the average opinion of a group is moderately pro-choice, there is a tendency after discussion for the group to become more extremely pro-choice. Echo chambers are contexts in which like-minded members of a group have limited exposure to the arguments and ideas of those outside the group. As a result, they become confirmed in whatever views they hold, and thus inoculated against any counterevidence that comes along.

The social media site, Facebook, presents a case-study of this kind of phenomenon, because Facebook, it is alleged, has two features well-suited to the creation of echo chambers and the production of group-polarization. First, social circles on Facebook tend to coalesce around a narrow range of points of view. Second, Facebook allegedly sculpts ads and posts that appear on member newsfeeds according to what its algorithms determine are members’ pre-existing political views and dispositions.[21] (See the *Wall Street Journal’s* “Blue Feed, Red Feed” site for examples of how different newsfeeds can look to conservative vs. liberal-leaning users.[22])

It is generally thought that there are two interwoven explanations for why echo chambers produce group polarization. First, according to social comparison theory, group polarization is the result of group members noticing that other members of the group on average have a certain opinion and modifying their own opinion accordingly. There are different ways this could work, but one possibility is offered by Roger Brown, who says, “To be virtuous . . . is to be different from the mean—in the right direction and to the right degree” (1974, 469). If group members find out that the group’s opinion is further from neutrality than initially thought, it will tend to move the opinion of those group members in the same direction, but even further from neutrality. Proponents of social comparison theory support their view with reference to what they call “mere-exposure” effects, according to which simply finding out what the group’s opinion is can polarize the opinions of members of the group.

---

[21] See, for example, (Miller 2016).
[22] The URL for the site can be found here: http://graphics.wsj.com/blue-feed-red-feed/.

A second explanation for group polarization, known as persuasive argumentation theory, has it that what causes group polarization is engagement with valid and novel arguments offered by the group. Because a group whose members tend to hold opinions of a certain sort will also tend to advance novel arguments that support that opinion, and because individual opinions tend to track the direction of novel arguments, as long as the group keeps offering new arguments, the direction of opinion will continue to move further toward the extreme end of the spectrum. Proponents of persuasive argumentation theory point to the fact that mere exposure effects are smaller than the effects of argumentation and that, furthermore, “Passive receipt of arguments outside an *interactive* discussion context generally produces reduced shift . . . Likewise, listening to a group discussion generally elicits less shift than actual participation” (Myers and Lamm 1976, 618).

Daniel Isenberg’s (1986) meta-analysis of experimental results on group polarization tend to confirm that both of these explanations contribute to the phenomenon. If so, then to the extent group polarization is to be avoided, it seems likely helpful to include diverse points of view and actively engaging with them. This conclusion receives experimental support. In his study on polarizing effects in jury context, Martin Kaplan concludes that “if discussed information is of *opposite* valence to the predominant information integrated in the initial response . . . the averaging process will lead to a *less* extreme response after discussion" (1977, 263). And Amiram Vinokur and Eugene Burnstein note that, in groups with equally sized subgroups that have points of view at odds with each other, "massive depolarization effects were observed" (1978, 884). Thus, it is no surprise that in his formal modeling of controversial argumentation, Gregor Betz finds that "[c]ontroversial argumentation is, all things considered, consensus conducive" and that "controversial argumentation compels proponent positions to converge, eventually" (2013, 10).[23]

Is group polarization to be avoided? Mill's argument seems to indicate that it is. When the mere fact that certain points of view are left out causes you to magnify the extremity of your position, you stand a greater chance of becoming sure of the false and denying the true. This can lead, as Cass Sunstein notes, to "feuds, ethnic and international strife, and war" (2002, 185). For "one of the characteristic features of feuds is that the feuding groups tend to talk only to one another, fueling and amplifying their outrage, and solidifying their impression of the relevant events" (185). If we can avoid group polarization by including and engaging with diverse points of view, we should do it.

Michael Lynch introduces a Rawlsian game—the "method game"[24]—designed to show what epistemic principles you should commit to. You should commit to those epistemic principles you'd have self-interest in settling on if you don't "assume that one method of belief formation is more reliable than any other" or "that one metaphysical picture of the world is any more accurate than others" (2010, 274) or "what [your] relative educational level, ethnicity, or social class" is (2012, 98). For example, "it would seem in our self-interest to favour privileging those methods that . . . were *repeatable, adaptable, public, and widespread*" (Lynch 2010).

---

[23] What is more, the convergence found is in line with the truth (despite triggering "a temporary loss of, instead of a gain in, verisimilitude") (Betz 2013, 13).
[24] See (Lynch 2013, 352), (Lynch 2010, 274), and in most detail (Lynch 2012, 95).

The method game tells us that you should commit to rules favoring open-mindedness and engagement, as long as behind the veil of ignorance you would endorse principles favoring open-mindedness and engagement. Behind the veil, you will tend to prefer that, whatever views you have when the veil is lifted, and whatever social position you're in, your voice is heard and your arguments seriously engaged with. If you are choosing rules for your group from a position of epistemic, metaphysical, and social neutrality, you will want members of the group to listen to each other and listen open-mindedly.[25] Therefore, according to the method game, you should commit to rules favoring open-mindedness and engagement. Or so it might seem.

5.4 THE RIGHTS OF DISPUTANTS

Mill's argument can be used to show that engagement with relevant counterarguments is necessary to procure good epistemic consequences for society as a whole and for yourself. An argument with less utilitarian resonance can be used to argue serious engagement is owed to those who put forth the arguments. As Ronald Dworkin says,

> It is no answer to say that if individuals have these rights, then the community will be better off in the long run as a whole. This idea—that individual rights may lead to overall utility—may or may not be true, but it is irrelevant to the defence of rights as such, because when we say that someone has a right to speak his mind freely, in the relevant political sense, we mean that he is entitled to do so even if this would not be in the general interest. If we want to defend individual rights in the sense in which we claim them, then we must try to discover something beyond utility that argues for these rights. (1977, 271)

---

[25] Lynch recommends that you "adopt a policy of using logical inference and observation in as unstinting and open-minded a way as possible" (2012, 84). This is in passing, however, and he doesn't say that a commitment to open-mindedness follows from the method game.

Dworkin, like West in the passage quoted earlier, does not think that free speech rights entail a right to be heard, denying "that the right to free speech includes . . . a right that others grasp and respect what one means to say" (1993, 38). But, as with the possibility of deriving an obligation to engage from the same kinds of considerations that motivate a right of free speech, the general idea that an obligation to engage with others derives from what we owe those others has some plausibility. We might owe something to those who put forth the arguments either intrinsically, because rational agents have an intrinsic right to be attended to, or because engaging with their opinions is something that is necessary to respect other rights they have.

In this latter camp we may place certain political liberals, who argue that in order for subjects in a democracy to be legitimately held to the laws that a majority has voted for, their voices must be taken into account in voting for those laws.[26] For example, Rawls says,

> Since the exercise of political power itself must be legitimate, the ideal of citizenship imposes a moral, not a legal, duty—the duty of civility—to be able to explain to one another on those fundamental questions how the principles and policies they advocate and vote for can be supported by the political values of public reason. This duty also involves a willingness to listen to others and a fairmindedness in deciding when accommodations to their views should reasonably be made. (1993, 217)

Similarly, Amy Gutmann and Dennis Thompson note that the only alternative to the politically illegitimate tactic of forcing conclusions down the throats of fellow citizens with whom we disagree is to engage with them in an attitude of mutual respect (Gutmann and Thompson 1990, 77). This requires a mutual acknowledgement of the moral status of the other's opinion. Mutual acknowledgement of the requisite sort requires more than mere toleration of the opinion. It requires actual attention and engagement. What is more, the engagement must be open in important ways:

---

[26] Also sharing this assumption about democratic legitimacy are discourse theorists like Seyla Benhabib, who maintain that "legitimacy in complex democratic societies must be thought to result from the free and unconstrained public deliberation of all about matters of common concern" (Benhabib 1996, 68).

> Mutual acknowledgment in this kind of moral discussion could turn out to be a merely formal, ritualistic expression of mutual respect, unless there is some real possibility that each side may be moved by the reasons the other gives. Therefore, a second principle (paralleling the requirement of integrity in action) is needed to keep open the possibility that citizens could come to adopt and act on the position of their opponents. (Gutmann and Thompson 1990, 80)

It's not clear the extent to which there is an individual obligation to engage with the arguments of those who will be relevantly impacted by enacted policies or merely a group obligation to do so. If a group enacts some policy that relevantly affects a person who was opposed to that policy, and if the person's arguments against the policy were not even taken into account, then that violates—if the liberal arguments are right—at least a group obligation not to enforce policies on a person unless they have some input in the process. There may be no individual obligation to ensure that the State does not enforce policy on subjects who had no input into the process, unless there are individual obligations to ensure that the group lives up to its obligations.

But there may be an individual obligation to engage with such arguments even beyond those individual obligations a member of a group accrues simply in virtue of being in a group with various group obligations. If you know that the policy that you are voting for—either directly if you have an immediate vote on a policy itself or indirectly by voting for a representative you know will favor that kind of policy—is opposed by someone who will be governed by the policy and if you know that this person has arguments that you have refused to hear out, you may be in violation of an obligation to this person (if, again, the liberal arguments are sound). There may be some individual obligation to take into account the arguments of affected subjects before voting in favor of policies that will affect those subjects. This is especially so if an obligation to engage with relevant counterarguments can be generated by a social contract one enters into by being a member of the state.

The liberal argument makes the case that our obligations to engage with relevant counterarguments stem from the requirements of democratic legitimacy. Alternatively, we might have intrinsic obligations to engage with a speaker's argument just in virtue of the fact that the speaker is an agent and deserves to be engaged with, irrespective of whether we might impose our will on them through any political process. This may be a palatable alternative for those—like Govier and Johnson—who think that obligations to engage with relevant counterarguments stem from something essential to the practice of argumentation. Since argumentation essentially involves other rational agents, perhaps we owe it to our interlocutors to take their arguments seriously, or perhaps it's a virtue to treat our audience with respect. As Johnson says, in proposing his "generic justification" for a "*prima facie* duty to respond to *all* the dialectical material directed at an argument," "when someone responds to something you have said, courtesy and civility dictate a response" (2001, 2). Stephen Carter makes a similar point:

> The ideal of civil listening is best captured not in a simple word but in a simple explanation. . . . I am to make your interests, joys, and sorrows my own. Not simply to understand but to try to share. . . . Surely this is often true of dialogue: simply by our openness to what others have to say, we accord them a respect that they can hardly help noticing. . . .
>
> But what does it mean to be open to what others have to say? It means, quite simply, to acknowledge the possibility that they may be right. We must approach each other, even in disagreement, with openness, not cynicism. . . .
>
> And so we can now see a rule taking shape:
>
> > *Civility requires that we listen to others with knowledge of the possibility that they are right and we are wrong.* (1999, 139, Carter's emphasis)

Here, the obligation to engage with others stems from the fact that they deserve respect, respect demands that we be civil to them, and civility demands that we listen to them with openness. There's a subtly different way that the obligation to engage with others can stem from their agency, however, and that's if agents have rights to have their views heard—not simply to be treated civilly, but to have the opportunity to transfer what they know, if they know it. This might be implied, if not stated explicitly, by Miranda Fricker in her classic discussion of testimonial injustice. According to Fricker, the primary harm of testimonial injustice—the wrong doled out to a person when the credibility of their testimony is downgraded on the basis of identity prejudice—is that the speaker is

> wronged in her capacity as a knower. To be wronged in one's capacity as a knower is to be wronged in a capacity essential to human value. When one is undermined or otherwise wronged in a capacity essential to human value, one suffers an epistemic injustice. The form that this intrinsic injustice takes specifically in cases of testimonial injustice is that the subject is wronged in her capacity as a giver of knowledge. The capacity to give knowledge to others is one side of that many-sided capacity so significant in human beings: namely, the capacity for reasons. . . . No wonder, then, that being insulted, undermined, or otherwise wronged in one's capacity as a giver of knowledge is something that can cut deep. . . .
>
> When someone suffers a testimonial injustice, they are degraded *qua* knower, and they are symbolically degraded *qua* human. (2007, 44)

For Fricker, the wronging constitutive of testimonial injustice lies not simply in the fact that the speaker is prevented from transferring any knowledge they might have. It's that the speaker is prevented from doing so on the basis of culpable attitudes (like identity prejudice). She is clear that being prevented from transferring knowledge on the basis of non-culpable attitudes—like belief based on good but misleading evidence—results in an unfortunate set of circumstances, but not an injustice.

Still, since, on Fricker's view, the inability to transmit any knowledge you might have cuts to the core of what it is to be human, the harm done to speakers if they are prevented across the board from transmitting knowledge is so deep that we might think it grounds some sort of right to be heard out. This position manifests itself, as well, in discussions of the right to freedom of expression. As Kent Greenawalt says, one justification for free speech

> that focuses on the speaker more than his listeners is the idea that the government should treat people with dignity and equality. As a matter of basic human respect we may owe it to each other to listen to what each of us has to say, or at least not to foreclose the opportunity to speak and to listen. Under this view, suppression represents a kind of

> contempt for citizens that is objectionable independent of its consequences; and when suppression favors some points of view over others, it may be regarded as failing to treat citizens equally. (1989, 153)

Mirroring Fricker, Greenawalt continues by noting that the connection between dignity and speech may be particularly tight:

> Expressions of beliefs and feelings lie closer to the core of our persons than do most actions we perform; restrictions of expressions may offend dignity to a greater degree than most other restrictions; and selective restrictions based on the content of our ideas may imply a specially significant inequality. (153)

Certainly, a person who is routinely silenced has suffered some sort of harm. Not everyone who suffers a harm has been dealt that harm as the result of a violation of an obligation. Still, it seems that if there are opportunities to at least somewhat easily hear the person out in a serious way, and if the person is never given the chance to be seriously attended to, there has not only been a harm, but a *failing* with respect to the person. If all people have the right to have their arguments engaged with by someone, then if there are people who are in a position to engage with their arguments but do not, then there seems to be a violation, at least of a group obligation.

Some people, because of various disabilities or circumstances, have a much more difficult time making their concerns or arguments heard by the larger public. This current defense of an obligation to engage with counterarguments does not leave out obligations to engage with the arguments those people might have an interest in getting heard. Rather, the current defense of the obligation to engage might impose duties on those without such obstacles to seek out such arguments, or anticipate them as best they can if it is absolutely impossible for members of those groups to make themselves heard, or to do their best to ensure that such obstacles are lessened or removed.

In contrast with the liberal defense of an obligation to engage with relevant counterarguments, it is much less clear that there is an individual obligation to engage even if there is a group obligation to do so. West thinks that, “A speaker’s interest in a receptive audience is insufficiently great even *prima facie* to justify the imposition of such duties” (2003, 407). Following West, it’s pretty clear that no person has a right that some other arbitrary person engage with their arguments. You might have a right to have your romantic partner engage with your argument, in virtue of the special relationship romantic partners bear to each other. But unless there is a special relationship between two people, it doesn’t seem that one has an automatic duty to engage with the other’s arguments simply from the fact that they’re both people and agents.

Still, many relevant counterarguments can be brought to your attention, often by an advocate of the counterargument. In this, the person has established something like a special relationship with you; you are, first of all, particularly well positioned to engage with the counterargument. Second, you have established something like an interpersonal relationship. Even if special relationships are established haphazardly or arbitrarily—even if you just happened to be in the wrong place at the wrong time—the establishment of that relationship can give you special duties that others don’t have. If you have been the seat companion of an eight-year-old child travelling alone to the big city, when at the end of the trip the adult who was supposed to be meeting the child does not show up, you may have duties to assist the child that the rest of the people on the bus do not. Likewise, if a person has brought their counterargument to your attention, and if it’s true that people in virtue of their agency have a right to be heard, this might force on you an obligation to engage that others don’t have.

This is not plausibly anything more than a *pro tanto* obligation, and it's presumably easily defeated.[27] It might be defeated, even, by your being bored by the argument, having a trivial task you need to perform, finding the argument offensive or malicious, or intellectually lazy, etc. On the other hand, if a person simply in virtue of their personhood has a right to have their arguments engaged with, then the group obligation to engage with that person's argument seems—even if only *pro tanto*—much stronger and much less easily defeated. For group obligations seemingly arise out of the existence of *pro tanto* individual obligations, even if none of the individuals have an overall obligation to take steps to fulfill that group obligation.

5.5 CONCLUSION

I've discussed three kinds of arguments for an obligation to engage with relevant counterarguments, distinguished by the source they suggest for that obligation: the role of the arguer, the epistemic goods engagement delivers to the one who engages and to the society at large, and the essential rights of those who put forth the argument. Ultimately, as I say, I think they all fail to establish a universal *pro tanto* obligation to engage with relevant counterarguments, even for those who are already participating in argumentation and have explicitly advocated for a position on some controversial and important matter.

Most especially, they fail to establish an obligation to engage with relevant counterarguments open-mindedly. On the contrary, there is a class of arguers—those who know that a relevant counterargument is misleading—who in the central cases not only lack a *pro tanto* obligation to engage open-mindedly with the relevant counterargument, but who positively should *not* engage open-mindedly with the counterargument. Things are less clear when it comes to closed-minded engagement. There are many cases in which those who know relevant counterarguments to be misleading should be engaging closed-mindedly with them. But if there is a *pro tanto* obligation, it is defeated to a greater extent than proponents of the obligation might have thought. I turn to the discussion of the obligation to at least closed-mindedly engage with relevant counterarguments in Chapter 7. In Chapter 6, I argue that in standard situations, knowers lack even a *pro tanto* obligation to engage open-mindedly with relevant counterarguments. Instead, they have an obligation not to engage open-mindedly.

---

[27] The obligation might be less easily defeated for some members of society—for example, politicians. Thanks to Josh Stein for emphasizing this point.

6

Against Open-Minded Engagement (for Some People)

> To defend one's mind against these follies a man must have an adamantine faith, so that, even if he is not able to detect the precise trick by which the illusion is produced, he at any rate retains his conviction that the whole thing is a lie and an impossibility. (Lucian)[1]

6.1 INTRODUCTION

A Google search for the expression "you're not going to change my mind" delivers, at the time of this writing, 568 thousand results. In an unscientific sampling, almost every instance expresses or is represented as expressing a culpable attitude. This is so even when the expression follows on the heels of a commitment to listen respectfully to what someone has to say. For example, Howard University professor Kelly Brown Douglas says that when she teaches courses that will eventually cover issues lesbians and gay men face in predominantly African American churches,

> students come in almost ready to pounce with, "You're not going to change my mind." They tell me on the first day, "You know, certain things are just sinful; and we'll hear what you have to say, but you're not going to change my mind." (interview in (Comstock 2001, 233))

Other examples are not hard to find:

> I often get asked why I don't believe in global warming, at least that man has recently caused it. I welcome any civil comments at the bottom of this article, but you're not going to change my mind. (Planka 2016)

and

> I'm doing my best to address your arguments and make my responses, but you're not going to change my mind. (Fora_Fauna 2015)

and 567,997 others, most of them intuitively bad.

---

[1] This passage is taken from Lucian's *Alexander; or, the False Prophet*, paragraph 17, and the translation is Joseph Jastrow's (1902, 49), though he doesn't cite the source. Though relatively pithy, Jastrow's translation is not particularly faithful to the text. For a more respectful translation, see A. M. Harmon's (Lucian c. 150 AD/1925).

Many of the declarations of intransigence, including the last two cited above, are not confessions; the authors aren't castigating themselves for being unwilling to reduce their confidence. Yet my intuitive reaction is that they should be. In the last chapter I offered arguments for an obligation to engage with relevant counterarguments, some of which suggest that engagement should be open-minded. But given the immediate negative valence surrounding declarations of closed-mindedness, one might think those arguments constitute Bernard Williams' "one thought too many" (1981b, 18). The quotes alone make it clear: open-mindedness is good and closed-mindedness is bad.

In this chapter, I argue that in many standard situations— though not in the cases above – open-mindedness is bad and closed-mindedness is good. There is sometimes no *pro tanto* obligation to engage open-mindedly. You sometimes should be unwilling to change your mind if you are unable to figure out what's wrong with relevant counterarguments and find each step compelling. This is so even when the counterarguments convince lots of people and would convince those people even if they had your evidence to the contrary; you sometimes should be closed-minded toward relevant counterarguments.

6.2 THE CASE AGAINST OPEN-MINDEDNESS

The most natural reason for thinking that you sometimes needn't engage with relevant counterarguments open-mindedly has to do with the sheer number of relevant counterarguments compared to your limited time and resources.[2] But this doesn't show that there is no *pro tanto* obligation to engage open-mindedly, nor that there is something more essentially problematic about open-minded engagement, which is what I will argue for. Even if there were all the time and resources in the world, you often shouldn't engage open-mindedly for epistemic and, sometimes, moral reasons.

---

[2] See, for example, (Cohen 2009, 59 ff), (Adler 2004b, 285), and (Govier 1997).

There are, of course, many situations in which you should engage open-mindedly with relevant counterarguments. But there are standard situations[3] in which you shouldn't engage open-mindedly with relevant counterarguments because you know they are misleading.[4] This is especially obvious when it comes to uncontroversial matters. There is something epistemically problematic about being so open to Zeno's arguments that if each step looks compelling, you are prepared to reconsider whether things move. The same goes for trick "proofs" that 1=0 and Sorites-style arguments that I'm not bald. However difficult those arguments are to resolve, there is no doubt that 1≠0 and, trust me, I'm bald.[5]

Even when it comes to controversial matters there are standard situations in which you shouldn't engage open-mindedly with relevant counterarguments because you know they are misleading. You shouldn't be willing to reduce your confidence or marshal new reasons simply because you find out what you know might well happen anyway— that you are unable to find a flaw in an argument for a false conclusion. You needn't be willing to marshal new reasons for evolution if you can't figure out where the argument for intelligent design in biology goes wrong. You shouldn't be willing to become less confident that it is impermissible to discriminate against gays and lesbians if you can't find the mistake in a clever argument that discrimination is morally permissible. The counterarguments might have the trappings of soundness, but if you know ahead of time they are misleading— as you sometimes do— you shouldn't be willing to be vulnerable to them.

---

[3] A situation is "standard" if there is nothing cognitively or conatively unusual. In particular, you neither fail to see obvious connections between premises and conclusions of theoretical or practical arguments nor, especially, do you have deviant preferences— preferences, for example, to have high confidence in the conclusions of misleading arguments. Discussion of the need for this proviso is below.

[4] You might know that the counterargument commits a fallacy or has false premises without knowing that the conclusion is false or what fallacy it commits. Some of the arguments invoked in this chapter can be modified to show that you shouldn't engage open-mindedly with apparently flawless arguments you don't know are misleading but you do know are fallacious or have false premises. See note 7. I don't have any new argument to offer against an obligation to engage when you fail to know that the counterargument is misleading or fallacious or has false premises.

[5] It may be required that you be open-minded to some degree, even toward arguments for uncontroversially false conclusions. For example, in response to Zeno's arguments that there is no motion, it might be required that you reduce your confidence that you understand what motion and change fundamentally amount to. If being willing to change at least some attitudes in response to a counterargument amounts to some degree of open-mindedness toward the argument— as I noted in Chapter 1 was suggested to me by Evangelian Collings— then this would amount to a requirement to be open-minded to some degree.

I'm not just arguing that you shouldn't be willing to here-and-now reduce your confidence when you know a counterargument is misleading. This minimal claim might simply be a consequence of the justified belief condition on knowledge (provided that the degree of belief knowledge justifies you in having is sufficiently high). But the minimal claim would not establish that you shouldn't engage open-mindedly with counterarguments you know are misleading. You engage open-mindedly when you are willing, not to here-and-now reduce your confidence, but conditionally willing to do so— when you are willing, *before* and *during* engagement, to reduce your confidence if, *after* engagement, you find the steps compelling and are unable to locate a flaw. It's in that sense you that you often shouldn't be open-minded when you know you're right.

The argument has two premises:

> (The knowledge premise) There are standard situations in which you know controversial propositions and, thus, know that a relevant counterargument is misleading.
>
> (The linking premise) If you know, in a standard situation, that a counterargument is misleading, you shouldn't engage open-mindedly with the counterargument.

Therefore,

> (C) There are standard situations in which you shouldn't engage open-mindedly with a relevant counterargument against a controversial proposition.

The conclusion is true because you shouldn't *be* open-minded toward relevant counterarguments you know are misleading. This is so however willingness is construed— as an endorsed disposition or as requiring an act of will, a striving, or a trying. You shouldn't have or endorse a disposition to reduce your confidence in the relevant situations. Nor should you strive or try through an act of will to reduce your confidence.

The argument for the knowledge premise occupied much of the first half of this book, in particular Chapter 2. To briefly sum up that argument, if you know a proposition, then you can know by a trivial inference that any argument against the proposition is an argument against a truth; it's an argument with a false conclusion. Only misleading arguments have false conclusions. Therefore, you can know to be misleading any argument against what you know. If you know controversial propositions, you can know that arguments against those propositions are misleading.

Controversial propositions are just propositions that people disagree about. There is nothing in that fact that prevents your epistemic support for controversial propositions from being quite strong— strong enough for knowledge.[6] The mere existence of controversy does not itself render justified belief and, hence, knowledge impossible. Nor is knowledge inevitably destroyed by the fact that, when it comes to controversial beliefs, there might well be relevant counterarguments out there in which, after engagement, you wouldn't be able to locate a flaw. If your knowledge that $1 \neq 0$ survives the actual existence of trick arguments to the contrary in which you cannot find a flaw, then knowledge of controversial propositions can survive the fact that there might well be such counterarguments against them.

---

[6] I do not assume any particular account of knowledge. In particular, I assume neither an externalist nor internalist account of the justification required for knowledge nor even that justification is required for knowledge. My sympathies are internalist, but what I conclude about knowledge of controversial propositions is a constraint on the proper account of knowledge, rather than supported by a prior account of knowledge.

Therefore, you can know controversial propositions. It follows that you can know that relevant counterarguments against your controversial beliefs are misleading and, what is more, you can know that they're misleading even if you haven't engaged with them. It's best to be clear about this: you can know that a relevant counterargument is misleading even if you have no idea what the content of the argument is (outside of its conclusion). In many central cases, you tend to have a least a minimal idea of what is involved in the counterargument and so, it might be thought, there has been at least minimal engagement already. But the claim here is stronger: you can know that an argument is misleading even if you are completely ignorant of its content (again, outside of its conclusion). You can know this because you can know that the conclusion of the argument is false. If someone tells you that they have an argument that the earth is less than an hour old, you know that the argument is misleading without knowing what the content of the argument is.

Similar considerations show that you can know that relevant counterarguments are misleading even if you have engaged with them and are unable to expose a flaw; again, you don't lose knowledge that things move simply by being unable to figure out what's wrong with Zeno's arguments. So the mere fact that you can't expose a flaw in some counterargument against one of your controversial beliefs does not by itself destroy knowledge. Because there are standard situations in which you know controversial propositions when confronted by relevant counterarguments, you can know in standard situations that relevant counterarguments against controversial propositions are misleading. That's the knowledge premise.

6.3 FROM KNOWLEDGE TO CLOSED-MINDEDNESS

6.3.1 Worries about the linking premise

The linking premise states in that in standard situations, if you know that a counterargument is misleading, you shouldn't engage open-mindedly with it. You shouldn't be willing to reduce your confidence conditional upon engaging with the argument, finding the steps compelling, and being unable to expose a flaw.[7] I call this "the linking premise" because it links your knowledge that a counterargument is misleading to the proper attitude toward that counterargument: closed-mindedness.

One immediate worry is that the advice the linking premise gives seems useless at best and dangerous at worst. On the one hand, you might not be able to tell which counterarguments you know are misleading. In that case, it's useless to be told to be closed-minded with respect to counterarguments you know are misleading. On the other hand, if you don't know that your controversial views are true, you might still tend to think that you know that your controversial views are true. Therefore, you might tend to incorrectly think you know various counterarguments are misleading. If so, it's dangerous to advise you to be closed-minded when you know a counterargument is misleading. A global-warming denier who wrongly thinks they know that some climate-scientist's argument is misleading will use the linking premise to be closed-minded toward the climate-scientist's argument.[8]

---

[7] A similar premise is true of arguments you know are fallacious or have false premises, but don't know are misleading. If a premise corresponding to the knowledge premise is true, then a further conclusion follows:

> (The knowledge premise*) There are standard situations in which you know some relevant counterargument against a controversial proposition is fallacious or has false premises, but don't know whether it's misleading.
>
> (The linking premise*) If you know, in a standard situation, that an argument is fallacious or has false premises, you shouldn't engage open-mindedly with it, even if you don't know whether it's misleading.

Therefore,

> (C*) There are standard situations in which you shouldn't engage open-mindedly with a relevant counterargument against a controversial proposition, even if you don't know whether it's misleading.

The linking premise* is true because the fact that a counterargument is fallacious or has false premises is a decisive reason, in standard situations, to be unwilling to reduce your confidence in response to the counterargument.

This might be what's behind various defenses of free speech according to which, if it becomes legally permissible to restrict obviously offensive speech, that permissibility will turn around and bite the less obviously offensive speech of those who pushed for the restrictions in the first place.[9] The linking premise doesn't concern when it is legitimate for the government to restrict speech, so the specific concern is not relevant. But an analogous concern is. The free speech concern is that if it is legally permissible to restrict obviously wrong or offensive views, then it will be too easy to permissibly restrict non-obviously-wrong, non-obviously-offensive views. The analogous concern is that if it is morally or epistemically permissible to be closed-minded toward obviously misleading arguments, it will be too easy to be permissibly closed-minded toward non-obviously-misleading arguments.

The objection fails. First, it is sometimes useful to be told that if you know a counterargument is misleading, then you don't have to engage open-mindedly. I distinguish between those of my controversial beliefs I take myself to know and those I don't. I take myself to know that sexual orientation is morally neutral. I do not take myself to know that raising the minimum wage won't increase unemployment in some specific instance. Of course, some people will be worse than others at distinguishing what they know and what they don't. So this advice won't be very useful for them. But the same goes for all advice that makes the right course of action depend on matters of fact that are sometimes difficult to figure out.

---

[8] This concern has been pressed to me in conversation by many, including Gordon Hawkes, audiences at the University of Lethbridge and Carleton University, and others. I discuss it here and discuss a related worry later in the chapter.

[9] See, for example, (Bernstein 2003 246).

Second, even if it's not clear when you know that counterarguments to your controversial beliefs are misleading, this doesn't render useless the conclusion that, when you do know a counterargument is misleading, you shouldn't engage open-mindedly with it. It just means that epistemologists have work to do. Progress has been made, because we've been told what we have to figure out in order to get some really substantial advice: we have to figure out what it takes to know that a counterargument is misleading. I tried to make some headway on that in Chapters 3 and 4.

Third, the premise— that knowers shouldn't engage open-mindedly with counterarguments— isn't advice. It's a metaphysical principle about what it standardly takes to be licensed in engaging open-mindedly with a counterargument; it takes a failure to know that the counterargument is misleading. That people might get wrong when some case meets the required condition doesn't undermine the metaphysical point. Compare this principle: there is no *pro tanto* obligation to negotiate in good faith with enemies who are invariably unreasonable in their demands. If this principle is true, it is a substantial moral claim that is worth arguing for. But it is also a principle that is particularly subject to misuse, since many who use the principle will tend to falsely think that their enemies are invariably unreasonable in their demands.

Likewise, the principle that you shouldn't engage open-mindedly with counterarguments you know are misleading is a principle particularly subject to misuse. Many of those who use the principle will tend to falsely think they know that counterarguments are misleading. This remains consistent with the truth and importance of the principle itself.

Any conditional with a normative consequent has this difficulty.[10] There will be those who wrongly think the antecedent of the conditional is true and so draw the conclusion that the normative consequent is true. This neither impugns the truth of those principles nor the value of discovering them. Indeed, it's true of any conditional whatsoever, normative or not. If the gun is loaded, you shouldn't shoot. If the yolk is still runny, the hard-boiled egg isn't finished. It can be hard to tell sometimes when the antecedents are true and people can wrongly think the antecedents are true (and known) when they're not. But the conditionals themselves remain true and useful. So too with the linking premise, if it's true; if you know that a counterargument is misleading, you often shouldn't engage with it open-mindedly.[11]

---

[10] Thanks to Mark Migotti for referring me to a similar response in Mill's *Utilitarianism*:

A second worry about the linking premise is that it licenses too much. For most propositions you know, there are arguments that would disprove them or at least rightly undermine your confidence if those arguments were presented to you. You know that George Washington was the first president of the United States. But if sufficient numbers of respected historians testify that he wasn't the first president, that the heretofore orthodox opinion that he was the first president was the result of a massive hoax— if there comes to be consensus among respected historians that this was all so— you should reduce your confidence. It might seem that the linking premise, however, tells us that you should not reduce your confidence after such an

---

> There is no difficulty in proving any ethical standard whatever to work ill, if we suppose universal idiocy to be conjoined with it. . . (1863, 34)

and

> We are told that an utilitarian will be apt to make his own particular case an exception to moral rules and, when under temptation, will see an utility in the breach of a rule, greater than he will see in its observance. But is utility the only creed which is able to furnish us with excuses for evil doing, and means of cheating our own conscience? They are afforded in abundance by all doctrines which recognize as a fact in morals the existence of conflicting considerations; which all doctrines do, that have been believed by sane persons. (1863, 36)

[11] What if you justifiedly, but falsely, believe that the counterargument is misleading? You can have otherwise knowledge-level overall support for your position but, by bad luck, are in a situation in which your support is misleading. If you are unwilling to engage open-mindedly, you will cut yourself off from the only way you have of remedying your error and putting yourself on the right track. Isn't that a bad consequence?

It is indeed an unfortunate consequence. It's also the right consequence. Sometimes what you should do will have the worst results. Take the extreme case: the victim of the evil demon. The victim of the evil demon falsely believes on the basis of knowledge-level evidence that they have a hand. If they are approached by someone claiming (correctly) to have a proof that they lack a hand, they will invoke the lesson of this chapter and ignore the argument (assuming they don't find such arguments particularly fun to wrestle with) and they will be right to do so. Such is the fate of victims of evil demons; what such beings should do can prevent them from arriving at the truth.

argument comes along. After all, you know that George Washington was the first president. Therefore you know that if a sufficient number of historians testify that he wasn't, then they are testifying to a falsehood. Therefore, by the linking premise, you should not be willing to lower your confidence in response to this sort of argument. The same applies to any proposition you know that allows for possible disproof.

This second objection also fails. The linking premise does not have the problematic consequence that, once the historians testify, you should remain unwilling to reduce your confidence. Once the historians testify, you no longer know that George Washington was the first president, so the antecedent of the linking premise isn't true of you. This is the lesson of the standard solution to Harman's dogmatism paradox discussed in Chapter 2. According to the standard solution, your knowledge that a counterargument against *p* is misleading survives only as long as your knowledge that *p* does. Since, when you are presented with certain counterarguments, your knowledge that *p* doesn't survive, you fail to retain knowledge that the counterargument is misleading. So though, prior to confrontation with the counterargument, you should be unwilling to reduce your confidence if the argument is presented to you, after confrontation with the argument, you should nonetheless reduce your confidence.[12]

What the linking premise precludes is a distinct possible solution to the dogmatism paradox. This second solution grants that, when you know that *p*, you also know that counterarguments against *p* are misleading. But, on the second solution, this doesn't license you in being unwilling to reduce your confidence if the counterarguments ever come along. You should be willing to reduce your confidence in response to counterarguments you know are misleading. So, knowing that *p* doesn't problematically license dogmatic response to counterarguments against *p*.

---

[12] If unwillingness to reduce your confidence is a disposition to maintain your confidence then, prior to exposure to the knowledge-destroying argument, you should be disposed to maintain your confidence conditional on confrontation with the argument. But your disposition should be "finkish"— a disposition, in David Lewis' words, "which would straightaway vanish if put to the test" (1997, 144).

The linking premise precludes this solution to the dogmatism paradox, because the linking premise states that if you know that a counterargument is misleading, you shouldn't be willing to reduce your confidence in response to the counterargument: you shouldn't engage with it open-mindedly. Is the linking premise true?

6.3.2 The argument for the linking premise

The argument for the linking premise has two key steps. First, if you know that a counterargument is misleading, then you should do whatever that fact is a decisive reason to do. Second, in standard situations, that a counterargument is misleading is a decisive reason to be unwilling to reduce confidence in your position in response to the counterargument (even if you engage with it, find each step compelling, and are unable to locate a flaw). I defend the steps in turn.

First, if you know that a relevant counterargument is misleading then you should do whatever that fact is a decisive reason to do. This is an instance of a more general "knowledge-action principle": for any *p*, if you know that *p* then you should do whatever *p* is a decisive reason to do.[13] To illustrate: suppose you are the hapless drinker in Bernard Williams' (1981a, 102) famed gin/petrol case. That the glass contains petrol is a decisive reason to refuse to drink (given your overriding desire not to drink petrol). But you don't know it contains petrol; all your evidence suggests it contains gin. Therefore, there's an important sense in which you shouldn't refuse to drink; you should drink. The decisive reason for you to refuse to drink does not ensure that you should— in the relevant, internal sense— refuse to drink. However, if you know that the glass contains petrol, then you should do the thing that fact is a decisive reason to do: refuse the drink.

---

[13] I spend some time defending this principle, but for more detailed arguments see (Fantl and McGrath 2009, esp. Ch. 3) and, for a more accessible discussion, (Fantl and McGrath 2014)), as well as (Hawthorne and Stanley 2008).

Whether a reason requires you to do something depends on two things— your epistemic support for the reason and the connection between the reason and what it's a reason to do. If the reason is a decisive one, then the connection between the reason and what it's a reason to do has everything it needs for the reason to require you to do that thing.[14] Perhaps, for *p* to be a decisive reason to do something, *p* must objectively and decisively favor doing it, and, furthermore, you must be aware that it's a reason to do that thing.[15] Or perhaps there is no requirement that you be aware of the favoring relationship. What's important is that, if *p* is a decisive reason to do something, then the connection between *p* and what it's a reason to do has everything it needs to make it that you should do that thing.

Therefore, if *p* is a decisive reason to do something, the only thing that could make it that you need not do that thing is that you don't have enough epistemic support for *p*. According to the knowledge-action principle, knowledge is always enough. If you know that the decisive reason is true then your epistemic support for the reason is all it needs to be for the reason to require of you whatever it's a decisive reason to do. Therefore, when you know a decisive reason to do something, you should do it.

---

[14] Some— e.g. Jon Kvanvig (2014, esp. Ch. 5)— deny the existence of obligations to believe. Kvanvig would presumably also object to the possibility of decisive reasons for belief. The linking premise is consistent with both denials, because the linking premise does not say anything about what you are obligated to believe nor about what you have decisive reasons to believe. The linking premise says that there are situations in which you are obligated to (because you have decisive reason to) be unwilling to reduce your confidence. This is a practical, not a theoretical conclusion. You can be obligated to be unwilling to have a specified (lower) confidence level without being obligated to have that confidence level, just as you can be obligated to be unwilling to harm someone even if you can't help but harm them and so, if ought implies can, you aren't obligated to refrain from harming them. See, again, Chapter 1 for details on the relationship between willingness to reduce your confidence and the ability to purposely reduce your confidence.

[15] This "awareness requirement"— that for *p* to be a decisive reason to do something, it not only must decisively favor doing it, but must be seen to favor doing it— is helpful for preventing simple truths from obligating you in believing really complicated logical and mathematical truths. You are not obligated to believe, to use Miriam Schoenfield's example, that the millionth digit of pi is even (Schoenfield 2012, 208) just because you know something that entails it— for example, that Ottawa is in Canada. Grant that the fact that Ottawa being in Canada entails that the millionth digit of pi is even means that Ottawa's being in Canada in some sense decisively favors believing that the millionth digit of pi is even. Still, it need not count as a decisive reason, in my sense, for so believing, if you are not aware of the entailment, and if the awareness requirement is true. Thanks to an anonymous referee for Oxford University Press for suggesting I discuss the issue of complicated logical truths in this context.

Some have argued that knowledge does not always provide enough epistemic support. They point to cases in which you know something but shouldn't act accordingly because the stakes are too high. In Jessica Brown's (2008, 176) example, a surgeon knows that a patient's left kidney is to be removed but shouldn't do what that fact is a reason to do: operate immediately without checking the patient's records. Instead, the surgeon should check the records before going to surgery, to make sure she doesn't remove the wrong kidney. When the stakes are high, it seems, knowledge sometimes isn't enough to act on.

There are two ways these sorts of cases can go. One allows that, when you know something, you can use it in your reasoning but that, when the stakes are high, the reason is defeated by the potential costs and chance of being wrong. If this is the way the case is construed, it's not a counterexample to the knowledge-action principle. The knowledge-action principle says that when you know a *decisive* reason to do something then you should do it. If you shouldn't do that thing only because the reason is defeated by a countervailing reason, then the reason isn't decisive. Once you use a decisive reason in reasoning, it has everything it needs to remain undefeated.

Counterexamples to the knowledge-action principle have to be cases in which knowledge of a decisive reason doesn't allow you to use that reason in reasoning. Again, once it's usable in reasoning, a decisive reason can't be defeated. If Brown's case is to be a counterexample, the surgeon has to know that the left kidney needs removal, but is not thereby permitted to use her knowledge in reasoning— defeated or otherwise. This is implausible. It would allow the surgeon to truly say, “Yes, I know it’s the left kidney that needs removal. And this fact is relevant to what I should do (in fact, it’s a decisive reason). But I mustn’t even take that fact into account in my reasoning.” If you know something that is relevant to your decision, you at least get to take it into account. Better to say that, in such cases, the reason is defeated and, so, is not decisive. I argue against this latter possibility, below. Regardless, Brown’s case should not be taken as a counterexample to the knowledge-action principle, above.

When you know a decisive reason, not only may you use it in your reasoning, you should. Suppose the fact that your car’s brakes don’t work is a decisive reason for you to walk, rather than drive, to work. Despite that fact, you drive to work, only to wrap your car around a tree at the bottom of the first hill. How might you defend yourself? One way is to argue that the world was at stake if you didn’t get to work as quickly as possible; the broken brakes were not a decisive reason for walking, because the nearly certain accident was worth risking. Another way is to claim ignorance of the fact that your brakes were broken. But if it’s permissible to fail to use known decisive reasons in reasoning, the following can be truly said in your defense:

> Yes, I knew the brakes were broken. And, yes, that they’re broken was a decisive reason for me to walk to work; so it was totally relevant to my decision (and I knew it was relevant).[16] But it was o.k. not to take the broken brakes into account.

Your defense, to put it mildly, wouldn’t hold up in a court of law or anywhere else.

If this argument is sound, then you should do whatever your knowledge provides decisive reasons to do, no matter how high the stakes get. The stakes are high when you’re deciding whether to send a ship full of people out to sea. But if you know the ship is seaworthy and that it’s seaworthy is a decisive reason for you to send it out to sea, then you should send it out to sea nonetheless. Increases in stakes may make it more difficult to acquire knowledge— may increase the strength of epistemic support required to know— but that doesn’t affect the truth of the principle. If you manage to have knowledge in high stakes situations, you should do what your knowledge is a decisive reason for you to do in that high stakes situation.

---

[16] If to be a decisive reason it is required that you be aware it’s a reason, the parenthetical clause (that you knew it was relevant) is guaranteed by the fact that the reason is decisive. If there can be decisive reasons that you aren’t aware are reasons, then the parenthetical clause can be false. But, in that case, the parenthetical clause isn’t necessary for your defense; if the awareness requirement is false, then you need not be aware that decisive reasons are reasons for them to require you— when you know them— to do the things they are reasons to do.

The second key step in the argument for the linking premise is that, in standard situations, that a counterargument is misleading is a decisive reason for you to be closed-minded toward the counterargument. It need not be a decisive reason you have, of course, until you know that the counterargument is misleading— much as there being petrol in the glass is not a reason you have for refusing the drink until you know that there is petrol in the glass. But, in standard situations, that a counterargument is misleading is a decisive reason to be closed-minded toward the counterargument. It defeats all potential reasons to be open-minded toward the counterargument.

One potential reason to be open-minded toward a counterargument is that the counterargument is a good inductive argument— a non-fallacious inductive argument with true premises. That a counterargument is a good inductive argument can be a reason to be willing to (dramatically) increase your confidence in the conclusion of the argument and to thereby reduce your confidence that the conclusion is false. Still, the inductive worth of a misleading counterargument is defeated by the fact that the counterargument is misleading. If it were otherwise, this reasoning should look sensible:

> On the one hand, the argument is non-fallacious and has true premises. On the other hand, the conclusion is false. The non-fallacious nature of the argument is more important than the fact that the conclusion is false. So I’ll have high confidence that the false conclusion is true.

This reasoning is absurd. In standard situations, the falsehood of a conclusion is a weightier reason than the inductive worth of an argument.[17] Again, if you don't know that the conclusion is false, the inductive worth of an argument might make it that you should have high confidence in the conclusion. But that's because what you should do is a matter of what reasons you know about; the reasons you have come apart from the reasons there are. The situation is even more clear if, as is often the case, you know not only that the counterargument is misleading, but that it is fallacious.

The fact that a counterargument is misleading can sometimes be defeated by other considerations. This can happen if you have deviant preferences— for example, a preference for raising your confidence in the conclusions of misleading arguments. These preferences might stem from fundamental iconoclasm, or because you're in a bizarre situation in which you're being threatened with death if you don't raise your confidence in response to misleading counterarguments. They might come about because of the requirements of friendship, if friendship requires that you believe your friend's declaration of innocence even if your total evidence does not support your friend.[18] They might also have their source in the demands of solidarity with marginalized groups, if those demands require trusting the testimony of a member of that group in some specific case more than the total evidence allows.[19]

---

[17] The falsehood of a conclusion is not a reason to have low confidence in it, nor is the truth of a conclusion a reason to have high confidence in it. For example, the truth of a conclusion is not a reason to be certain that the conclusion is true (though it is a reason to believe the conclusion). However, when you know the conclusion of an argument is false, you have a defeater that prevents the argument from conferring positive epistemic status on the conclusion. For example, if it's a good but misleading inductive argument, the defeater would be a further piece of evidence such that when you add it to the inductive base, the resulting inductive base does not raise the probability of the conclusion.

[18] See, for example, (Baker 1987).

[19] Avery Kolers argues that solidarity can demand acting with others, "even if one disagrees with the group's chosen ends or means" (2012, 365). Though this is a demand on action, there might be similar demands on belief: perhaps solidarity can demand believing with others, even if one (initially) disagrees with what the group believes.

In standard situations, though, you don't have deviant preferences.[20] Therefore, in all standard situations, that a counterargument is misleading remains undefeated by these considerations. Absent a further consideration that defeats the fact that a counterargument is misleading— a consideration you should take into account in standard situations— when you know that a counterargument is misleading, you shouldn't be willing to engage open-mindedly with the counterargument; you should decide to maintain your confidence even if, after engagement, you can't locate a flaw.[21]

Are there other defeaters that, in standard situations, would render the fact that a counterargument is misleading a non-decisive reason to be closed-minded toward the argument? One possible defeater I already discussed above: the possibility that you take yourself to know things that in fact you don't know. Does this fact— that it's possible that you wrongly take yourself to know that a counterargument is misleading— defeat the inference from the misleadingness of the argument to the decision to be unwilling to be open-minded toward the argument? It might seem so. For this reasoning can sound legitimate:

> Sure, the counterargument is misleading. But, I know that lots of people *think* that they know some counterargument is misleading even though they don't really know this. And while I don't think I'm one of those people— this counterargument really is misleading— those people probably feel the same way. I certainly wouldn't want a global warming denier to say, "I know that there is no global warming. So I know that even if I can't figure out what's going wrong with scientific arguments to the contrary, those arguments are misleading. So I don't have to engage open-mindedly with them." If the

---

[20] It is a further feature of standard situations that you see obvious connections between reasons and what they're reasons to do. As noted above, it may be that you shouldn't do what there is an objective and decisive favoring reason to do, if you don't know that the reason favors doing it. But it is obvious that a counterargument being misleading is a reason to be unwilling to reduce your confidence in response to the counterargument in standard situations. Therefore, in standard situations, you should be unwilling to reduce your confidence if you know that the counterargument is misleading.

[21] Again, it might be that, after engagement, you lose knowledge that the conclusion of the argument is false and, thus, lose knowledge that the argument is misleading. In that case, you should, after engagement, be open-minded toward the counterargument. But that doesn't change the fact that, prior to engagement, you should be unwilling to reduce your confidence if, after engagement, you can't locate a flaw in the counterargument. See, again, Chapters 2 and 3, as well as earlier in this chapter, for discussion of the dogmatism paradox and the related issues of when knowledge is and isn't destroyed over the course of engagement with a relevant counterargument.

> global warming denier doesn't get to reason that way, why do I get to reason from my knowledge that some counterargument is misleading to the conclusion that I should be unwilling to engage open-mindedly with it?

Does the possibility that you wrongly take yourself to know that a reason is true defeat the inference from that reason to what it's a reason to do? I've already noted a number of difficulties with a "yes" answer to this question. In particular a "yes" answer would problematize the use in reasoning of all conditionals with normative conclusions and, indeed, all conditionals whatsoever, since all useful conditionals have antecedents you can fool yourself into falsely believing. Here I'll note an additional problem more specific to the current context: you don't generally allow reasons to be defeated by the possibility that you don't know the reason. You might give up commitment to the reason, but you don't maintain commitment to the reason while allowing it to be defeated by the possibility that you don't know it.[22]

To see this, consider how practical reasoning goes in a genuine case of defeat (imagine you're Clifford's shipowner, though somewhat more reflective):

> On the one hand, the ship is seaworthy. So the inspection will just tell me it's seaworthy and I'll have wasted time and money getting that very answer. On the other hand, the law requires an inspection, so if I just send it out to sea without getting it inspected, I'll go to jail. Freedom is more important than the wasted money. I'd better get it inspected.

In genuine cases of defeat, there's a reason why the defeater defeats the inference. There are competing goods, each reflected by one of the competing reasons. One reason defeats the other when one of the goods weighs more heavily than the other. In this case, avoiding jail time outweighs the savings in time and money. But when it comes to the possibility that you don't know that the reason is true, it's hard to see what the competing good is that weighs so much more heavily than the good that figures in the defeating reason:

---

[22] Because, as I've argued, when you know something that's relevant to your decision, you at least get to use it in reasoning, insisting that you should give up commitment to the reason requires insisting that you don't know the reason.

> On the one hand, the ship is seaworthy. So the inspection will just tell me it's seaworthy and I'll have wasted time and money getting that very answer. On the other hand, I might not know that the ship is seaworthy, so if I just send it out to sea without getting it inspected, I'll. . . what?

It's tempting to say that, as long as it's possible you don't know that the ship is seaworthy, then by sending it out to sea without getting it inspected, you run the risk of sending a ship out to sea that isn't seaworthy. But in fact you do know that the ship is seaworthy. So in maintaining your commitment to the seaworthiness of the ship, you know that the possibility of sending an un-seaworthy ship out to sea is not one that will obtain. Yes, sending the ship out to sea while the ship isn't seaworthy has potentially disastrous results. Sending the ship out to sea while the ship is seaworthy but not known to be so has no greater potential for disastrous results than sending out the ship while the ship is seaworthy and known to be so. Storms have no greater chance of sinking seaworthy ships that are not known to be seaworthy than they do of sinking seaworthy ships that are known to be seaworthy.

What's required for defeat is not that you might wrongly take yourself to know that the ship is seaworthy, but that there's a significant chance that the ship is not seaworthy. It is presumably this defeater that is supposed to be kicking in when you worry that you might not know that the ship is seaworthy. Perhaps, even if you know that a reason is true— and, so, should use it in your reasoning— the reason can be defeated by the chance that the reason is false. Assuming fallibilism about knowledge, it's possible to know things even if there's a chance those things are false. And if there's a chance that those things are false, then shouldn't you reason like this?

> Sure, the counterargument is misleading. But I've been wrong before, so I should adopt a policy of being open-minded toward all counterarguments, even the ones I know are misleading. If worse comes to worst, in being open-minded, I'll just confirm my initial belief that the argument is misleading. But if the argument is sound, I'll prevent myself

> from recognizing this if I engage closed-mindedly or do not engage at all. I'd better engage open-mindedly.

Do reasons, when known, get defeated in standard situations by the chance that those reasons are false? Again, imagine you're the shipowner. The relevant bit of practical reasoning would look something like this:

> On the one hand, the ship is seaworthy. So if I send it out to sea without inspection I'll save money on an inspection that will only tell me it already was seaworthy. On the other hand, the ship might not be seaworthy. So if I get it inspected first, I increase the chance that I'll find out it's not seaworthy (if it isn't), which will allow me to do necessary repairs and reduce the risk that the ship won't sink. I'd better get it inspected.

It's clearer here what the competing goods are supposed to be: risk reduction vs financial savings. Unfortunately, the practical reasoning on offer is implausible. Reasons don't get weighed against and defeated by the chance the reason is false. Implicit in this reasoning is a kind of vacillation that is more apparent in such expressions as "The ship is seaworthy. Isn't it? Maybe it's not. I'd better err on the safe side." In deciding to err on the safe side, you retract the commitment to the seaworthiness of the ship.[23] The above reasoning hides the vacillation better, but it is no less a vacillation. For when we alter the above reasoning to make explicit that no vacillation is taking place, it turns from sensible to bizarre:

> Yes, the ship is seaworthy. So, if I send it out to sea, it will have better results than if I don't. But while that's true, it might not be true, and if it isn't (which it is), and I send it out to sea without inspection, it will have disastrous results (which it won't). Because there's a chance (that doesn't obtain) that it's not seaworthy, I'll waste some money getting it inspected first. It's seaworthy and I'd better check whether it's seaworthy.

Here is the same difficulty that plagued Adler's and Riggs' accounts of open-mindedness and full belief in Chapter 1. The line of reasoning is absurd and the resultant decision on the basis of the line of reasoning is irrational: you shouldn't make that decision. When you know that *p* and *p* is a reason to do something, that reason is not defeated— barring deviant preferences— by considerations having to do with the chance that *p* is false. Seaworthy ships that have a chance of being unseaworthy are benefitted no more by inspection than seaworthy ships that have no chance of being unseaworthy.

---

[23] For more on playing it safe and the vacillation response, see (Fantl and McGrath 2012, 35–37) and (Fantl and McGrath 2014, 93).

Parallel considerations apply to practical reasoning about the misleadingness of relevant counterarguments. To make sense, your reasoning would have to look like this:

> On the one hand, the counterargument is misleading. So if I engage with it open-mindedly and I can't figure out where it goes wrong, I'll end up reducing my confidence in response to a misleading argument. On the other hand, the counterargument might be sound. So if I engage with it open-mindedly, I increase the chance that I'll find out it's sound (if it is), which will allow me a better chance at getting truth. I'd better engage open-mindedly.

As in the shipowner's reasoning, there is vacillation hidden in this statement of the reasoning. When you decide to become willing to reduce your confidence in response to the counterargument, you pull back on your commitment to the misleadingness of the counterargument. This becomes obvious when the reasoning is stated in such a way that the possibility of vacillation is eliminated. When it is, as with the shipowner's reasoning, it looks bizarre:

> Yes, the counterargument is misleading. So if I engage with it open-mindedly and I can't figure out where it goes wrong, I'll end up reducing my confidence in response to a misleading argument. But while that's true, it might not be true, and if it isn't (which it is), and I engage closed-mindedly, it will have disastrous results (which it won't). Because there's a chance (that doesn't obtain) that the counterargument is sound, I'll (unnecessarily and dangerously) engage with it open-mindedly. It's misleading and I'd better check whether it's misleading.

This line of reasoning, as with the corresponding line of shipowner reasoning, is absurd and the conclusion drawn is similarly irrational. Therefore, when you know that a counterargument is misleading, that reason is not defeated— barring deviant preferences— by considerations having to do with the chance that the counterargument is sound.

Perhaps what's really motivating the objection is not the claim that, sometimes, the chance that *p* is false defeats what would otherwise be a decisive inference from *p* to what it's a reason for. Rather, the worry is that the fact that you could be wrong about *p* or the fact that people can wrongly think they know that *p* means that you don't know that *p* in the first place. Perhaps what this objector would like to say to the shipowner is not, "Sure, you know the ship is seaworthy, but you should still get the ship inspected in case you're wrong," but rather, "Sure, the ship looks good to the naked eye, but you can't know it's seaworthy until you get it inspected." Likewise, when it comes to those who claim to know that some counterargument is misleading, perhaps what the objector would like to say is not, "Sure, you know that the counterargument is misleading, but you still should be willing to reduce your confidence in response to the counterargument in case you're wrong," but rather, "Sure, you've got some good evidence that the counterargument is misleading, but you don't really know it yet (you haven't examined it, after all), so you'd better be open-minded toward the counterargument."

That is, perhaps what's really motivating the objection is not unhappiness with the linking premise, but with the knowledge premise. The knowledge premise says that there are standard situations in which you know controversial propositions and, thus, know that arguments against those propositions are misleading. But it might be thought that in order to draw a substantial conclusion from the argument, the knowledge premise needs to say more. It needs to say that you continue to know that the counterarguments are misleading *when you have the opportunity to engage with them*. The argument for the linking premise might seem to provide a way to deny that.

The first step in the argument for the linking premise is the knowledge-action principle: if you know that *p* then you should do whatever *p* is a decisive reason to do. Therefore, when you know that a counterargument is misleading, you should act accordingly. You should be unwilling to reduce your confidence in response to that counterargument. But, an objector might say, in situations in which you're in a position to engage with relevant counterarguments against controversial claims, the stakes are pretty high. You run the risk of fostering false beliefs about matters of moral, political, religious, and scientific importance. In such contexts the heightened stakes make it that the small chance that a relevant counterargument is sound makes a difference to what you should do. In particular, the heightened stakes mean that you should be willing to reduce your confidence in response to the counterargument. In public arenas where important matters of politics, religion, ethics, and science are debated, the heightened stakes require open-mindedness that might not be required in private, low-stakes contexts.[24]

This is not because the linking premise is false. As noted above, if the linking premise is true, it's true no matter what: if you know that *p* in a high stakes situation, you should do whatever *p* is a decisive reason for you to do in that high stakes situation. Open-mindedness is required, not because the linking premise is false, and not because you can never know that relevant counterarguments are misleading, but because knowledge of controversial matters is lost in the public arena where those controversial matters are debated. Knowledge is lost precisely because the linking premise is true: if you know a counterargument is misleading, you shouldn't engage with it open-mindedly. But in the high-stakes context of the public arena, you should engage with relevant counterarguments open-mindedly. Therefore, by the linking premise, you don't know that those counterarguments are misleading and, thus, don't know the controversial propositions that entail that the counterarguments are misleading.[25] Or so the worry goes.

---

[24] There are heightened stakes as well to failing to engage at least closed-mindedly, for the reasons discussed in Chapter 5. But if you stand a lose a lot not just by failing to engage at all, but failing to engage open-mindedly, then you might lose knowledge that the counterargument is misleading.

The worry can be overcome. Though the heightened stakes afforded by an opportunity to publicly engage with relevant counterarguments can destroy knowledge, it's not true that heightened stakes invariably destroy knowledge. Heightened stakes only destroy knowledge if your strength of epistemic support is barely high enough for knowledge in low-stakes contexts. There's no reason to think that your epistemic support for controversial propositions can't be strong enough to survive in high-stakes contexts as well.

First, it's clear that there are propositions you would continue to know even if they became controversial and you were placed in a context where the stakes were high and you had the opportunity to engage open-mindedly (or not) with those views. Suppose the views that that the earth is flat,[26] that the sun revolves around the earth,[27] and that Hillary Clinton is an actual demon,[28] were threatening to infect the received views of the masses; people were gradually becoming convinced by charismatic figures and sophistical arguments. This social phenomenon plausibly raises the stakes on public disagreements about such matters. But your epistemic support is not thereby insufficient for knowledge that the earth is not flat and that Hillary Clinton is not an actual demon.

---

[25] This conclusion fits with a conclusion often drawn from principles like the knowledge-action principle— that knowledge is unstable. It's harder to know things in high stakes contexts: there is so-called "pragmatic encroachment" on knowledge. For versions of this kind of view, see, for example, (Grimm 2015), (Weatherson 2012), (Fantl and McGrath 2007), (Stanley 2005), (Hawthorne 2004), and (Fantl and McGrath 2002). (Grimm's version might not allow opportunities for public engagement to raise the stakes in a knowledge-destroying way, because according to Grimm, the standards for knowledge are already fairly high as long as there are others who might appeal to your belief; the standards for knowledge are set not just by your interests, but by the interests others might have.) See (Anderson 2015) and (Reed 2010) for objections to principles like the knowledge-action principle, deriving from worries about its implications about the instability of knowledge.

[26] The Flat Earth Society boasts over 500 members, all of whom have joined since 2009 (Flat Earth Society 2016). Are they sincere? Flat Earth Society Vice President Michael Wilmore say that he is "fully convinced" of the belief of "many members" (Wolchover 2012). NBA star Kyrie Irving is among several NBA players to go on record as a flat-earther, though there is some evidence they're all trolling (see (Wilder 2017)).

[27] A 2014 NSF survey of 2200 Americans found 26% responding that the sun revolves around the earth (Neuman 2014). This is up from the 18% figure found by a 1999 Gallup poll, which is itself comparable to a 1996 German poll which found 16% responding that the sun revolves around the earth (Crabtree 1999).

[28] In a 2016 Public Policy Polling survey 40% of Donald Trump supporters reported believing that "Hillary Clinton is an actual demon." On an optimistic note, 19% of Trump supporters were "not sure." See (Brinlee 2016).

More generally, no part of the argument in Chapter 2 that you can know controversial propositions and thereby know that relevant counterarguments against those propositions are misleading relies on situations in which the stakes are low. The argument relies on the claim that mere disagreement does not present any principled obstacles to having otherwise knowledge-level strong epistemic support for controversial propositions. If that's true, then it's likewise true that mere disagreement does not present any principled obstacles to having epistemic support for a controversial proposition that is otherwise knowledge-level strong even in high-stakes contexts.

Nor does the fact that there could easily be an apparently flawless counterargument automatically reduce the possible strength of support below the level required for knowledge in high-stakes contexts. Even in low-stakes contexts it's not like the decision about whether to engage is a trivial matter. Failure to privately engage open-mindedly has potential costs both for yourself and your fellows. As Clifford notes, "no one man's belief is in any case a private matter which concerns himself alone" (1886, 342). Therefore, though the stakes may be raised somewhat when it comes to public engagement, they won't be raised as dramatically as they would be if you went from observing a ship in a museum to owning the ship outright and having to decide whether to send it out to sea. As long as the arguments in the first half of the book correctly conclude that there are standard situations in which you know that relevant counterarguments are misleading, they further conclude that you can know such counterarguments are misleading in some somewhat higher-stakes contexts in which you have the opportunity to publicly engage on controversial matters.

Finally, the principled conclusion here fits with the intuitive data. Intuitively, there are controversial beliefs each one of us have that, intuitively, we know. Speaking for myself, I take myself to know that global warming is real and that human activity is a principle cause; that sexual orientation is morally neutral; that a widespread program of vaccination reduces health risks; that approximately 6 million Jews were killed in the Holocaust; and that there is no omnipotent, omnibenevolent creator of the universe (not to mention that the earth is not flat and that Hillary Clinton is not an actual demon). My intuition that I know these things does not diminish when I participate in public discussion. Others might take themselves to know different controversial propositions and, obviously, if others disagree with me, at least one of us is wrong. But the fact that the principled arguments (rehearsed above) *confirm* our intuitions tells us that all of the evidence— intuitive and principled— is on the same side: in the high-stakes contexts afforded by the public arena, you can continue to know controversial propositions are true.

Thus, even in the public arena, you can know that relevant counterarguments against controversial propositions are misleading. By the linking premise, you should not engage with those arguments open-mindedly. Therefore, there are standard situations in which you have the opportunity to, but shouldn't, engage with relevant counterarguments open-mindedly. This is not because you are ill-placed to engage with them or because you have countervailing reasons that defeat what would otherwise be decisive reasons to engage open-mindedly with them. It's that being open-minded toward arguments you know to be misleading is intrinsically epistemically improper.[29]

In addition to this epistemic impropriety in being willing to engage open-mindedly with a counterargument you know to be misleading, there is sometimes a moral impropriety, when the counterargument is harmful in certain ways to those it is about. When you are discussing sexual harassment policy or policies having to do with the rights of gay job applicants or issues having to do with disability, open-minded engagement is an act that can harm those the arguments concern. I argue in Chapter 8 that psychological, physical, and intrinsic harm can be done by just the appearance of open-minded engagement. But here I want to point out that intrinsic harm can be done by actual open-minded engagement. There is something distasteful about taking certain problematic speakers and arguments seriously enough to be open-minded toward them. Open-minded engagement can constitute a failure to stand in solidarity with members of marginalized, oppressed, or victimized groups and individuals.

---

[29] Again, there might also be an epistemic impropriety in engaging open-mindedly with counterarguments you don't know are misleading, if you do know they're fallacious or have false premises. See note 7.

I think this is what some authors are getting at when they worry about the tendency of philosophers to treat all arguments as worth discussing in a clinical way, without regard for the fact that the people the arguments are about are not themselves variables. As Rima Basu (manuscript) puts the point,

> if you are going to ask a morally controversial question—should we try to cure homosexuality”—you cannot treat that question with the level of playful indifference you would when out drinking with friends and theorizing about where your lap goes when you stand up.

I do not claim that any of these authors are correct about the specific areas they are discussing. I just note that one charge against a certain kind of open-minded engagement is that it involves a failure to stand in solidarity with or support the right people. Whether that charge sticks in a specific case requires more detailed analysis of that case. Still, the charge is common.

Grace Jantzen (1999), for example, in her critique of the “disengaged intellectual gymnastics” (262) that she sees characterizing mainstream academic discussion of the problem of evil, argues that “Both sides proceed from a detached intellectual perspective, as privileged onlookers, rather than in solidarity with those in suffering” (263). The sentiment is repeated in

Mark Lance's contributions to a debate on the Leiter Reports blog on whether the American Philosophical Association should sanction universities that require employees to sign statements pledging to abstain from sexual relationships with members of the same sex.

> I still have the feeling that a huge part of this discussion is going on in a very unhealthy abstraction from this context, as if we are discussing modal realism rather than a contribution to systematic structural violence against some of our colleagues. (Lance 2009a)

> Gays generally, in most parts of the US, have no right to visit their partner in the hospital. Gays are routinely denied the chance to adopt, or even foster, children. . . . Gays are assaulted, murdered, dragged behind trucks, bound with barbed wire, denounced daily in the media in the most vile terms. That is the context in which we consider the question of whether the APA should embrace colleges that explicitly state their intention to fire someone for engaging in gay sex. No matter how closely we stick to our academic focus, that context is there, and it is real. It is the context in which we consider whether to "remain neutral" on a matter "about which reasonable philosophers disagree".
>
> So many here speak as if we were considering philosophical disagreements about meriological [sic.] essentialism, or deciding whether a professional organization should officially pronounce on relevance logic, rather than considering whether to engage in symbolic solidarity with an oppressed and often despised class of people. (Lance 2009b)

Finally, here's disability rights activist and lawyer Harriet McBryde Johnson's account of a public interaction with Peter Singer:

> He responds to each point with clear and lucid counterarguments. He proceeds with the assumption that I am one of the people who might rightly have been killed at birth. He sticks to his guns, conceding just enough to show himself open-minded and flexible. We go back and forth for 10 long minutes. Even as I am horrified by what he says, and by the fact that I have been sucked into a civil discussion of whether I ought to exist, I can't help being dazzled by his verbal facility. He is so respectful, so free of condescension, so focused on the argument, that by the time the show is over, I'm not exactly angry with him. Yes, I am shaking, furious, enraged— but it's for the big room, 200 of my fellow Charlestonians who have listened with polite interest, when in decency they should have run him out of town on a rail. (2003, 53)

Basu (manuscript) argues that philosophers have an obligation to adopt the "participant stance" when considering how to present their arguments and what arguments to present. In particular, philosophers need to consider the impact their arguments and conclusions have on the feelings of those they are arguing with and about. More generally, philosophical discourse is subject to moral norms no less than philosophical writing is subject to aesthetic norms. In Chapter 8 I argue in more depth that respectful engagement with problematic arguments and speakers constitutes something like a betrayal and, thus, is a violation of a moral obligation. Here I just flag the issue and let it stand on its initial plausibility. There seems to be something impermissible— stemming from the disrespect shown to those who are victimized by certain kinds of problematic arguments and speakers— about taking certain speakers seriously enough to be open-minded toward their arguments, especially if that open-minded attitude is reflected in respectful behavior. If that's right, knowers can participate in both epistemic and moral impropriety when they engage open-mindedly with arguments they know are misleading. Therefore, in many cases in which you know a counterargument is misleading, you shouldn't engage with the argument any further— at least, not open-mindedly.

It is consistent with this conclusion that you have a *pro tanto* obligation to engage open-mindedly that has been outweighed by a stronger obligation not to. Many of the arguments surveyed in Chapter 5 have only the weaker conclusion— that there is a *pro tanto* obligation to engage open-mindedly with relevant counterarguments you know are misleading. But those who argue for this weaker conclusion should not be happy with the conclusion in this chapter, regardless.

For one thing, even if a *pro tanto* obligation to engage were defeated in standard situations in which you know the counterargument is misleading, it would conflict with what many of the defenders of engagement advocate— that the obligation does not get defeated by issues involving certainty, frivolity, silliness, offensiveness, sexism, or racism. More importantly, there are good reasons to think that the relevant cases don't involve one *pro tanto* obligation— the obligation to engage— being defeated by stronger *pro tanto* obligations—

obligations to act in accord with what you know and to stand in solidarity with members of marginalized groups.

Suppose you witness someone on the street shouting explicitly racist insults at a member of some ethnic minority. You respond by intervening and telling the harasser that their behavior is inappropriate. Telling a stranger that their behavior is inappropriate is sometimes impermissible, even in many cases in which their behavior is inappropriate. For example, you should usually not intervene when you overhear a stranger dominating their group's conversation at a nearby table in a restaurant. When the stranger is shouting explicitly racist insults, however, intervention is called for. There are two ways to think about this situation. One is that you have two *pro tanto* obligations— an obligation of politeness and an obligation to prevent racist harassment, and in this case the second obligation is stronger. The second way to think about the situation is that, in this situation, you have no *pro tanto* obligation to be polite to the racist harasser.

The latter option is correct. In this case, there is no competition between *pro tanto* obligations because there is no *pro tanto* obligation to be polite to the harasser. Note the contrast between this case and a case of uncontroversial conflict between *pro tanto* obligations. Suppose that in order to get your sick child to the emergency room quickly enough, you're forced to be rude to an innocent supermarket cashier. You have a duty of politeness and a duty to help your sick child and, in this case, the duty to help your sick child is stronger.

In the second case, but not the first, you should regret that you were forced to neglect the duty of politeness. You might even be right to feel guilty that you neglected the duty, and have a duty to later apologize to the cashier for being rude. There are various moral responses you should have to your violation of the *pro tanto* duty of politeness in the second case. None of that applies in the first case: you need not apologize for being impolite to the harasser; you should not feel guilty; nor should you regret doing it. While a violation of a merely *pro tanto* obligation is not blameworthy, it still has to mean something. But in the case of the racist harasser, all the normal consequences of violating a *pro tanto* obligation are missing.

They are missing too in the case of the knower who does not engage open-mindedly with relevant counterarguments. The knower should not regret their failure to be willing to be vulnerable to the argument. They should not feel guilty. They should not apologize. They shouldn't think, wistfully, "If only I could have been willing to reduce my confidence in response to that misleading argument, but my knowledge that it was misleading prevented me. . ." If there were a *pro tanto* obligation to engage open-mindedly with relevant counterarguments you know to be misleading, then when you don't do so, you should respond in various ways. In fact, none of these responses are called for. I conclude that not only is there an all-things-considered obligation for knowers to be closed-minded toward with relevant counterarguments, there is no *pro tanto* obligation to engage open-mindedly, either.

6.4 DEALING WITH THE ARGUMENTS FOR OPEN-MINDED ENGAGEMENT

In the last chapter I discussed three types of arguments for an obligation to engage with relevant counterarguments. First, an obligation to engage might derive from essential features of a social role, in particular the role of participant in the practice of argumentation. Second, an obligation to engage might derive from the goods— especially epistemic goods— that engagement might get us. Third, an obligation to engage might derive from the rights of those offering the relevant counterargument.

The first— that the obligation is grounded in essential features of the practice of argumentation— is the least likely to generate an obligation to engage open-mindedly. Ralph Johnson's argument— that argumentation has the essential purpose of achieving rational assent— does not generate an obligation to engage open-mindedly, since achieving rational assent does not depend on your being willing to reduce confidence in your position if you can't figure out where your interlocutor's argument goes wrong. It may be contingently true that achieving rational assent is best achieved when you are not moved by affective factors or subject to wishful thinking, and when you adhere to procedural norms. But because— as we saw in Chapter 1— these factors aren't sufficient for open-mindedness, even if the purpose of argumentation is to achieve rational assent, that purpose won't thereby ground an obligation to engage open-mindedly.[30]

A rule requiring open-minded engagement is also not a constitutive rule of the practice of argumentation in Williamson's sense. As we saw, it's hard enough to motivate a constitutive rule of argumentation according to which it is essential to participation in argumentation that you are obligated to engage with relevant counterarguments. It is even trickier to successfully argue that not only must you engage with relevant counterarguments, but you must do so with the willingness to reduce your confidence conditional on spending significant time with the counterargument and not finding anything wrong with it. In any case, such an argument is refuted by a positive demonstration that there are situations in which you have no *pro tanto* obligation to engage open-mindedly with relevant counterarguments, which is what I provide in this chapter. Because, ultimately, the argument for a dialectical constitutive rule of argumentation rests on intuitive responses to cases, the existence of situations in which the proper response is that you shouldn't engage open-mindedly leaves the positive case for a constitutive dialectical rule of argumentation toothless.

---

[30] Nor do other social roles generate an obligation to engage open-mindedly. For example, if you have taken on the role of commenter at a conference, you might thereby have an obligation to present objections to the paper you are commenting on, and perhaps adhere to various procedural norms. But it is not a feature of your role in standard situations that you have the obligation to be willing to reduce your confidence; the obligations accruing to the role are entirely social.

The third argument-style, which seeks to source an obligation to engage open-mindedly in the rights of the person putting forth the counterargument, is not persuasive either. You may owe it to the person putting forth a counterargument that you respond rationally to them, that you adhere to procedural norms, or prevent affective factors from influencing your evaluation of the counterargument. But you don't owe it to them— simply because they are agents— to reduce confidence if you can't figure out where their argument goes wrong, as long as there are positive reasons why you shouldn't be willing to do it. Again, because as I've argued in this chapter you often shouldn't engage open-mindedly with relevant counterarguments, you don't owe it to those putting forth the counterarguments that you're willing to do it.

For this reason, an obligation to engage open-mindedly with relevant counterarguments in standard situations is not generated by your entering into a social contract with those you are arguing with. You need not implicitly agree, when you're arguing with me, to give up your knowledge that the claim you are defending is true. While you can enter into contracts that require you to do things you otherwise shouldn't do, it shouldn't be impossible to enter into argumentation with someone about something you know without implicitly agreeing to do something you otherwise shouldn't do.

You might explicitly contract with me that if you can't figure out where my counterargument goes wrong, you'll reduce your confidence. This contract might generate an obligation to reduce your confidence if you can't figure out where my counterargument goes wrong, though it might depend on whether you have obligations not to enter into that kind of contract. But even if an obligation to engage open-mindedly with me is generated by your entering into such a contract, this is not a standard situation. It is a situation in which you have deviant preferences: (conditional) preferences for reducing your confidence in response to misleading arguments. Therefore, it is not a worry for the claim here— that knowers shouldn't engage open-mindedly with relevant counterarguments in standard situations.

One might think that if it is only legitimate for me to be subject to democratically enforced laws and policies if my voice had a chance to influence those laws, then if no one is willing to reduce their confidence in response to my arguments, it violates my rights when those laws are enforced on me nonetheless. However, all that is required to respect my rights in this regard is that I have the opportunity to present arguments which people are moved by if my arguments rationally compel people to be moved by them. I don't have the right to force movement of your confidence if movement of your confidence in response to my argument is irrational. It is not illegitimate to democratically enforce upon me just laws I do not manage to offer rationally compelling arguments against.

I shouldn't be silenced simply because my listeners are swayed by affective factors or because my listeners violate procedural norms. But if my arguments don't sway anyone because it is rational for my listeners to maintain their confidence despite my best arguments, there has been no violation of my rights when just laws are democratically enforced upon me. As above, given the existence of an independent positive demonstration that you shouldn't be willing to reduce your confidence in response to my counterargument, then in failing to be so willing, you haven't violated my rights, as long as you've respected procedural norms and avoided influence from affective factors.

Finally, we have to ask whether an obligation to engage open-mindedly derives from the goods— epistemic and otherwise— that open-minded engagement gets for yourself or the larger society. The response to this must involve a demonstration that open-minded engagement, in many cases, does not produce epistemic benefits for yourself or society at large. This, again, is the case I presented in this chapter. It does not produce epistemic benefits for yourself because, in knowing that the counterargument is misleading, you only stand to distort your confidence by being willing to reduce your confidence in response to a misleading counterargument.

Does it produce epistemic benefits for the society at large to engage open-mindedly with relevant counterarguments? As noted, it's possible that open-minded engagement can help prevent group polarization. But group polarization is not something that knowers should want to avoid. You shouldn't be avoiding the creation of an echo chamber of those who think that the earth is round. Some have argued— rightly, in my view— that one problem with charismatic politicians who casually spout bigoted remarks is that they normalize the legitimacy of making those remarks. For example, Brittany Stalsburg (2016) argues that Donald Trump's U.S. presidential campaign normalized bigoted speech against a number of different marginalized groups. You should want to leave some views on the margins. You should *want* to be in echo chambers of certain sorts. What sorts? The sorts of chambers that echo views you know are true.[31]

---

[31] Nor should you, behind the veil of ignorance, prefer a norm that requires open-mindedness toward all counterarguments—even those you know are misleading. For knowing, as you do, that once the veil is lifted you might well be confronted by apparently flawless misleading counterarguments, you will hope that knowing that such counterarguments are misleading can act as a preventative against being misled by those arguments. Nor does self-interest obviously favor preferring that people are willing to reduce their confidence in response to your apparently flawless counterarguments, if they know your counterarguments are misleading. It is not always in your interest for your false views to propagate. Thus, Lynch's epistemic method game (discussed in Chapter 5) does not favor adopting epistemic principles that privilege open-mindedness toward counterarguments you know are misleading.

There are other problems with echo chambers: they can distort your impression of how pervasive your views are. This sociological function has nothing to do with analysis of whether your views are true, and so if you want to avoid false demographic information, you should seek to exit your echo chamber regardless. This goal, however, is compatible with closed-mindedness. You can achieve the goal of sociological accuracy about the pervasiveness of false views by allowing those views a hearing on your Facebook page or by looking at public opinion polls about the prevalence of those views (though seeing them on your Facebook page might make vivid their prevalence in ways that merely looking at public opinion polls might not). This doesn't require engagement with those views, let alone open-minded engagement.

None of this is to say that the gatekeepers of arguments and information— like Facebook— don't have some obligation to allow counterarguments and information from all perspectives to show up on newsfeeds. A general policy of allowing exposure to relevant counterarguments might be appropriate even if, in specific cases, knowers don't have obligations to engage with those relevant counterarguments, or if knowers shouldn't do so open-mindedly. My conclusion is only the latter, not the former.

Public engagement can deliver great epistemic benefits to society as a whole, as the public gets to both observe and participate in the presentation of arguments, counterarguments, objections, and replies on behalf of both sides of controversial issues. But these benefits can be achieved regardless of whether those advocating for either side are willing to reduce their confidence if they can't figure out what's wrong with the counterargument; the epistemic goods to the larger society can be achieved just as well, and perhaps better,[32] by closed-minded engagement as they can by open-minded engagement. Therefore, the epistemic goods engagement delivers to the public do not require— and are sometimes in tension with— open-minded engagement. The question remains, to what extent knowers have an obligation to engage at least closed-mindedly with relevant counterarguments. It is to this topic I turn in the next chapter.

---

[32] "Perhaps better," because if knowers open-mindedly engage with counterarguments they know to be misleading, they might well be swayed when they shouldn't— simply by being unable to figure out where the known-to-be-misleading counterargument goes wrong. If they're swayed— when, knowing that the counterargument is misleading, they shouldn't be— the public can witness them doing this and be incorrectly swayed themselves.

7

Against Closed-Minded Engagement (in Some Situations)

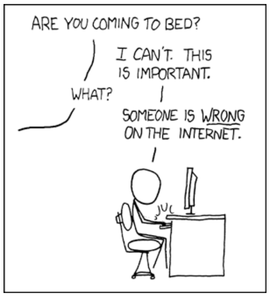
ARE YOU COMING TO BED?
I CAN'T. THIS IS IMPORTANT.
WHAT?
SOMEONE IS WRONG ON THE INTERNET.

(XKCD webcomic[1])

7.1 THE PITFALLS OF CLOSED-MINDED ENGAGEMENT

It is obvious that in some standard situations you should engage with relevant counterarguments you know are misleading. As I argued in the last chapter, you should not engage with them open-mindedly. The surprising conclusion is that, to the extent knowers have an obligation to engage with relevant counterarguments in standard situations, the obligation derives from an obligation to engage closed-mindedly.

There are a number of reasons you might engage with a relevant counterargument closed-mindedly. Consistent refusal to engage with certain views—even closed-mindedly—might have the result that proponents of the views are driven to extreme measures. The consistent silencing of even abhorrent positions can increase any antagonism proponents might have toward members of the silencing group. Even closed-minded engagement—if polite—might serve to lessen this antagonism in some cases. You also might engage closed-mindedly in order to correct easily correctable errors, to convince your interlocutor and your listeners that the arguments are misleading, to call out deeply offensive or victimizing speech, to stand in solidarity with members of marginalized groups, to resist your oppressor, to help your interlocutor overcome their prejudice or error, because your interlocutor is sincerely asking for help with their reasoning or ideas, or to resolve disputes in a manner that satisfies all parties. Closed-minded engagement, then, can be done in the service of yourself, society at large, groups of which the speaker is a member, the victims of the speech with which you closed-mindedly engage, or the speaker themselves.

---

[1] Permanent link: http://xkcd.com/386/.

Different techniques of closed-minded engagement might be appropriate for different goals. If you're trying to resist your oppressor or stand in solidarity with members of marginalized groups, your closed-minded engagement might rightly take the form of challenging and even aggressive speech. If you're trying to correct errors or make clear that some victimizing or offensive speech is unacceptable, your tone might rightly be unsympathetic and lecturing. Both of these kinds of methods fall under the general category of what is often referred to as *calling out*. On the other hand, if you're trying to convince the speaker or a hostile audience, trying to help the speaker overcome some prejudice or work through their arguments or ideas, or trying to resolve a dispute in a mutually satisfying way, then you might well try to emphasize compassion and openness—to hide or at least not advertise your closed-mindedness.

The various methods and goals come with different pitfalls. First, it is not clear to what extent closed-minded engagement actually achieves the goal of correcting errors or getting your interlocutor or audience to overcome prejudice. One familiar worry about closed-minded engagement—or engagement generally—is that it gives credibility to fringe arguments. This is what NASA thought in 2002 when they pulled a planned book debunking the moon-landing-hoax conspiracy theory. Sean O'Keefe, then head of NASA, remarked after the project was scrapped, "The issue of trying to do a targeted response to this is just lending credibility to something that is, on its face, asinine" (Dunn 2003, A11). Similarly, as the possibly apocryphal story of Lyndon Johnson's first statewide election campaign goes, after Johnson's campaign manager berated him for spreading rumors that his opponent, Ralph Yarborough, had sex with pigs, Johnson replied, "I just want to hear the son of a bitch deny it once" (Stevens 2001, 13).[2]

The danger of engagement is especially acute when the counterargument in question is so utterly on the fringe that its influence on the public hasn't been felt yet. In such cases, the likelihood that public engagement will succeed only in giving the view credibility it didn't have already is very high. Kristen Intemann and Inmaculada de Melo-Martín note that among the dangers of engaging with clearly fringe scientific views is that

> establishing public venues for all dissenters to be heard and have their criticisms considered may contribute to a false public perception that there is significant disagreement or that no scientific consensus exists. Indeed, some private companies and think-tanks have funded scientific research aimed at generating skepticism about climate change and environmental toxins, stalling the development of public policy, and creating doubt among the public and policymakers (Oreskes and Conway 2010; Michaels 2008). Although this may happen regardless of whether scientists seek and engage the participation of dissenters, there is fear that doing so risks bringing even more attention to dissenters and exacerbating the problem. (2014, 2752)

The risks are even higher when, though you know that the argument is misleading, you realize you can't figure out what's wrong with the argument. So, too, if you can figure out what's wrong with the argument but realize that the difficulties are highly technical, so that in engaging with the argument you will give the public the perception that it's one of those issues about which experts can disagree, and the details of which are beyond their understanding. In any of these cases, engagement easily backfires, if the goal is to convince interlocutors or prevent the spread of misleading arguments.

As countervailing data, Rose Solomon (manuscript) notes that subsequent to Barack Obama's release of his long form birth certificate, Gallup polling (2011) revealed that the percentage of adult U.S. citizens who believed that Obama was definitely born in the United States went up from 38% to 47% in the space of a month. And we've already seen (in Chapter 5) that engagement with counterarguments offered by group-members with diverse perspectives has been shown effective in limiting group polarization. Putting aside the issue (discussed in Chapter 6) that knowers may not want to limit group polarization surrounding propositions they know are true, there is some data that engagement with counterarguments can prevent people from entrenching in their views.

---

[2] Thanks to Richard Steinberg for first telling me this story.

Still, the data is at best equivocal, and it's not helped by recent work on the difficulties using reasoned arguments—or anything else—to change people's minds. Even Amiram Vinokur and Eugene Burnstein, who are among those who plump for the depolarizing effects of including diverse perspectives, note that depolarization effects are less noticeable the more popular and interesting are the subject matters:

> matters of social value (rated as most popular and most interesting) depolarized the least, matters of personal taste (rated as moderately popular and moderately interesting) depolarized a moderate amount, and obscure matters of fact (rated as least popular and least interesting) depolarized the most. (1978, 884)

Furthermore, as Magdalena Wojcieszak points out,

> disagreement may trigger stereotyping of out-group members (Sinclair & Kunda, 2000), leading people to dismiss the views expressed by outgroups. To the extent that extreme individuals see those who disagree with them as an out-group, polarization may increase. Studies on an anticipated opinion climate similarly find that people who expect to be in the majority prepare less for a debate and are more biased in thinking about their own position, findings which suggest that majority members may enter a debate ready to ignore opposing views (Levine & Russo, 1995; Zdaniuk & Levine, 1996). Because extreme people project their views onto others and believe themselves to be in a majority (Ross, Green, & House, 1977), they may be less open to dissimilar views and emerge from group discussions with polarized attitudes. (2011, 599)

The most well-known phenomenon in this context is the so-called "backfire effect." In Brendan Nyhan and Jason Reifler's influential (2010) study subjects were given fake newspaper stories and then the corrected versions. Subjects whose political leanings agreed with the fake stories did not merely dismiss the corrections; the corrections entrenched their false opinions deeper than they had been before the corrections. For example, subjects were given a fake news story indicating that weapons of mass destruction (WMDs) had been found in Iraq during the Iraq War. Those people who supported the war and described themselves as politically right of center, unsurprisingly, were more inclined to believe this than those who described themselves as politically left of center. Then came the retraction and the corrected information. Unsurprisingly, people who supported the war were less likely to believe the retraction than people who opposed it. What's surprising is that the corrected information made people who supported the war more confident that there were WMDs discovered in Iraq:

> The percentage of conservatives agreeing with the statement that Iraq had weapons of mass destruction before the US invasion increased from 32% in the control condition to 64% in the correction condition. By contrast, for non-conservatives, agreement went from 22% to 13%. (Nyhan and Reifler 2010, 315)

Similar results were observed in Nyhan and Reifler's (2014) study on the entrenchment of anti-vaccination beliefs and behavior in the face of various pro-vaccination interventions and in Kelly Garrett and Brian Weeks' (2013) study on beliefs about who has access to your electronic health records.[3]

While it might be that those antecedently disposed to believe some fringe view become more entrenched after receiving arguments to the contrary, it might also be that those on the fence are moved to believe the truth (this phenomenon is especially apparent in Garrett and Weeks' study). That's enough for closed-minded public engagement to serve an important good. Nor do the studies tell us what might happen to persistently corrected belief in the long run.

---

[3] For a discussion of the backfire effect and related issues in comic strip form, see http://theoatmeal.com/comics/believe. Thanks to Amanda Brown for the pointer.

Furthermore, as we will see, there might be methods other than those problematized by the various studies[4] that are more effective at changing people's attitudes and getting them to stick. For example, Nyhan and Reifler (manuscript), who are generally pessimistic about the possibility of using reasoned arguments to change minds, produce data indicating that if you want to convince someone, show them a *graph*.[5] As we will also see, though, the data is conflicted.

I flag the issue just because ineffectiveness and, worse, the backfire effect, is one of the pitfalls that closed-minded engagement needs to avoid in those cases in which you are trying to change people's minds or stop the spread of misleading arguments. One way it is thought intervention can be more effective is to treat your subject with a stance of compassion and openness. But to the extent you do this, a second pitfall looms. If you present yourself as open or at least fail to emphasize the degree to which your mind is closed, if you are in fact closed-minded you can easily be problematically insincere. If you are using the impression of open-mindedness to get your interlocutor to come to certain beliefs or attitudes, then you are exploiting their false belief to change their minds. This, in many cases, is a pernicious kind of manipulation.

---

[4] The methods problematized by the various studies include reasoned argument, empirical data, personal narratives, and disturbing images. See (Nyhan, Reifler, et al. 2014). In other contexts, however, disturbing images have been found to have measurable impact on behavior. Smokers who notice on their cigarette packages graphic images of the results of tobacco use show a significant tendency toward reduction in smoking. See (Hammond, Fong, et al. 2003).

[5] The difference when providing a graphic representation of the data is significant:

Figure 1

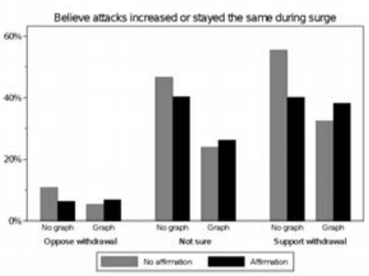

Figure 2

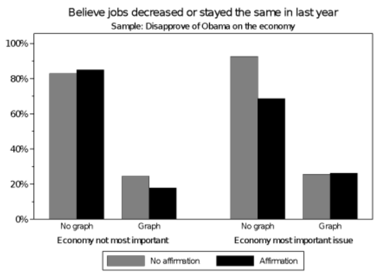

Figure 3

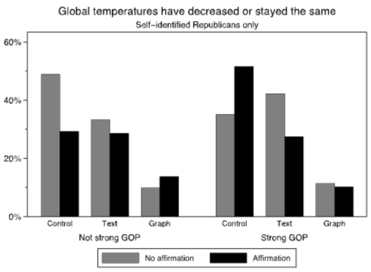

To avoid this, you might honestly represent yourself as closed-minded. You might challenge offensive speech in no uncertain terms or lecture to a speaker unsympathetically in order to call out offensive speech or behavior or to correct obvious errors. Aside from the difficulty doing this while also effectively persuading your audience, this kind of speech can easily become overly confrontational, patronizing, and problematically disrespectful.

What is more, some of the methods required to achieve one goal can be antithetical to achieving others. The challenging speech required to stand in appropriate solidarity with members of marginalized group is also the kind of speech that can alienate your audience and entrench them more in their offensive or erroneous views. Accurately representing yourself as closed-minded can backfire if you're trying to convince an audience. Calling out inappropriate behavior tends not to resolve disputes in a mutually satisfying way. So it might be that there are certain collections of goals and certain kinds of standard situations for which there is no permissible way to closed-mindedly engage with relevant counterarguments. Thus, in those situations, when you know the relevant counterargument is misleading, there is nothing for it but to refrain from engagement at all.

When should knowers engage with relevant counterarguments closed-mindedly in standard situations? In brief, you should engage closed-mindedly only if doing so is *non-manipulatively effective*. You should engage closed-mindedly only if you don't mislead people into taking you to be open-minded and thereby manipulate them into doing or believing what you want them to do or believe. The difficulty is that in some situations, the only effective way to closed-mindedly engage is by being deceptive. To be effective you sometimes have to allow your interlocutor to see you as open-minded. Therefore, there is no way for knowers to rightly and effectively engage in these situations: either you wrongly become open-minded, or you wrongly (so I will argue) deceive your interlocutor.

7.2 THE PITFALLS OF INSINCERITY

Are there ways to avoid these pitfalls? To sum up: if you are unwilling to reduce your confidence if you find a counterargument apparently flawless, then you have three options when hearing an argument out. Either you give the false impression that you are willing to reduce your confidence, you don't offer a representation of your attitudes, or you make clear what your attitudes are. The first two run the risk of problematic insincerity. The third runs the risk of becoming patronizing, insulting, or ineffective. In many situations, all the pitfalls can be avoided. But in some situations, they can't. I treat the first two options together, before turning to the third.

The first option is to actively give the false impression that you are open-minded toward the counterargument. You might do this explicitly, by prefacing your engagement with some statement of the form, "I'm willing to be convinced. What's your evidence?" It can also be accomplished in more subtle ways, through impressed smiles, pauses in which you look to be actively mulling over the details of the counterargument, approving nods of the head, exclaiming "nice point!", etc. Either way, in taking this first option, you positively and falsely represent yourself as willing to reduce your confidence in your opinion if you find each step compelling and can't figure out what's wrong with the counterargument.

The second option is that you don't falsely represent yourself as willing to change your attitudes in response to the argument, but only because you don't represent your attitudes at all. You can limit yourself to exposing where hypothetical opponents of the counterargument would be most tempted to object, without revealing whether you personally are one of those hypothetical opponents. On the other hand, you might, when taking this second option, accurately represent your beliefs on the matter without indicating whether you are open-minded or closed-minded toward the counterargument. In a formal debate, you can make plain where you stand on the issue without making plain whether you are open-minded or closed-minded toward the counterarguments you are engaging with.[6]

There are lots of situations in which these two options are permissible and perhaps obligatory. In addition to formal debates, educational settings stand out. Teachers especially might be expected to represent themselves as more open to being swayed by a student's argument than they really are. Teachers can likewise sometimes be expected to leave their opinions at the door and limit themselves, quite explicitly, to pointing out where readers might have the most trouble accepting the steps in a student's argument. (Perhaps situations in which a friend or colleague has requested help working through an argument that you know is misleading fit into this category as well.)

What's distinctive about formal debate and pedagogical settings is that the conventions surrounding them cancel any implicature from your participation in argumentation to conclusions about whether you personally are open-minded or closed-minded toward the arguments you are engaging with.[7] Formal debates can go so far as to arbitrarily assign argumentative roles to the various debaters, independent of the debaters' actual beliefs. But even when the role is assigned on the basis of the beliefs of the one occupying the role, the conventions don't vary. We expect any debater occupying the role of defending that particular position to behave pretty much the same way, regardless of their personal attitudes. So as long as their behavior fits within the confines we expect from someone occupying that role, we shouldn't infer from their behavior to what their personal attitudes are. They can act open-minded without our rightly inferring that they are open-minded.[8]

---

[6] One model of this kind of debate is John Corvino's and Mary Gallagher's (2012) book on the value of same-sex marriage.

[7] The same goes, I would argue, for academic papers and formal presentations. This is why it can be permissible not to represent yourself as closed-minded when presenting the case for the other side in a way that trains your audience appropriately (see the discussion in Chapter 3).

Similarly with pedagogical settings: teachers ought to exhibit a certain range of behaviors regardless of the personal beliefs of the person who occupies that role. For example, a teacher might behave more openly to a student's argument than they really feel. They might say things like, "This is really promising," while really feeling A) that the argument is awful as it stands and B) that though there is a germ of an idea that could, in the right hands, be developed into an excellent paper, the student in question does not have the right hands to do it.[9] Still, it is part of the conventional duties of the role that the teacher not say all the worst things on their mind, but instead say that, for example, the paper is really promising (at least when there is a germ of an idea that could in the right hands be developed into an excellent paper). It is part of the role to say this even when the teacher's personal attitudes are much more negative. Therefore, as long as a teacher adheres to the requirements dictated by their role, it doesn't follow from the fact that a teacher makes a certain set of comments that the teacher feels very positively toward the paper.

---

[8] In some formal settings, including formal debate settings, a façade of open-mindedness is demanded by the duties of the role you're in; it's a duty of politeness and respect to others in the formal setting. This duty arguably would not be permissibly discharged if there were a presumption that the appearance of open-mindedness reflected genuine open-mindedness. But because in the relevant formal settings, there is no such presumption, the duty to behave with at least the façade of open-mindedness is permissibly discharged. I argue in the next chapter that this formal duty precludes participation in formal settings with certain kinds of speakers.

[9] In Jennifer Lackey's classic (1999, 477) case of the creationist teacher who nonetheless follows her school's requirement to teach evolutionary theory, students can get knowledge of the truth of evolutionary theory from her lecture. The same conventions that license the teacher in attesting to the truth of evolutionary theory also prohibit the inference from the fact that a teacher sounds open-minded toward a student's argument to the conclusion that they really are.

So, in some formal settings, including but not limited to pedagogical and formal debate settings, it is a convention that occupants of the relevant roles behave in certain ways no matter what their personal attitudes. When you are an occupant of a role in such a formal setting, the insincerity involved in failing to accurately represent your attitudes is no worse than the insincerity involved when Robert Duvall said "I love the smell of napalm in the morning." In failing to accurately represent your attitudes, your behavior (as long as it stays in the range that is expected for occupants of the role) does not license inferences about what your personal attitudes really are.[10] Apparently open-minded behavior need not be problematically deceptive in such a setting, even if you're closed-minded.

It's possible to rationally take teachers, debate participants, and actors as having accurately represented themselves when they have made statements that don't reflect their views. If an audience member at a debate is rationally unaware of the norms of debate performance, if a student is rationally unaware that teachers might try to mask their true feelings toward a paper, if someone watching a play is rationally unaware that what they are viewing is a play, then they might rationally conclude that the speakers are conveying their true attitudes. Nonetheless, deception is sometimes required by the norms of the activity in question despite the fact that some might rationally draw the wrong conclusions from the speakers in that activity.

But there are other situations—ostensible peer-on-peer informal arguments, for example—in which, if you don't represent yourself accurately there are no relevant norms that prevent you from having wrongly deceived your audience.[11] In those cases, your interlocutor may be rational in taking you to be more sympathetic than you really are, and norms don't excuse or justify your behavior. In these situations, unbeknownst to your interlocutor, you are one of the very opponents you are warning your interlocutor about. Your interlocutor and the listening public have the right to take into account the fact that you are antecedently disposed against the view. In many settings, taking either the first or second option—failing to accurately represent yourself as closed-minded—makes you, in contemporary parlance, a *concern troll*.

---

[10] This is one (lesser) reason why sexually predatory behaviors of a teacher are especially egregious (other than the obvious and most important reasons involving betrayal and professional, psychological, and physical harm). Predatory behaviors allow teachers to exploit and misuse important conventions that prevent students from inferring to personal attitudes from public behavior. As long as the predator's behavior falls at least roughly within the normal range allowed by pedagogical conventions, the targets of those predators will doubt whether their suspicions about the predator's underlying attitudes are reasonable.

Of course, the good results of such deception might seem to outweigh the bad of the deception. Consider the Listening Project—a conflict resolution organization that tries to effect socially just outcomes in community disputes. As indicated by the title, the emphasis is on listening without judgment or correction—asking interested and curious questions when necessary to lead parties in disputes naturally to socially just and compassionate resolutions of their differences. The community interventions sponsored by the Listening Project—and its previous incarnation, the Rural Southern Voice for Peace—"help groups gauge and change attitudes toward homophobia, abortion, AIDS, defense spending, racial conflict, community development, breastfeeding, and the death penalty" (Imlay and Howard 2006, 2). Says Herb Walters, the founder of the Listening Project,

> We don't change people by clobbering them over the head. . . . We change people through a process of active, justice-seeking love. We teach empathy. Empathy isn't agreement. It's understanding where people are starting from and seeing the potential. It involves listening at a very deep level so that one builds a relationship of trust and respect. To listen at a deep level, you have to let go of your own strong beliefs. It's a spiritual process of seeking God in the other person. (Imlay and Howard 2006, 2)

---

[11] Whether a setting counts as formal in the relevant sense is highly context-sensitive. Many personal conversations, blog comments, social media posts, letters to the editor, questions in Q&A sessions, and political statements can be made in relevantly informal settings.

There is a seeming tension in this passage—between Walters' claim that their listening techniques don't require agreeing with your interlocutor and his claim that their listening techniques require letting go of your strong beliefs. If disagreement is allowed, then why can't you maintain your strong beliefs? But the tension is only apparent. The strong beliefs you are required to give up are not about the content of your interlocutor's speech, but rather about your interlocutor's worth, intelligence, or deservingness of condemnation. You can think they're wrong. But you can't think they're bad.

Again, the techniques emphasized by the Listening Project are mostly observational. You ask questions; they answer. Because the questions are entirely curiosity-driven, they do not put people on their guard. You are interested in them as people, and so they get led naturally, the hope is, to the truth—the truth you already know. Here's a description of a model listening session between members of the Listening Project—including Walters—and their subject, Jeff, in which the Jeff ends up coming around on some racist beliefs:

> Walters, a southerner himself, puts the man at ease with some casual questions, then poses an open-ended query: "What are your hopes and fears about race relations today?"
>
> Jeff's answer: "I think it's hopeless. I'm not a racist or anything, but Abe Lincoln should have sent them back to Africa," states he, pointing towards a nearby black neighborhood.
>
> Walters lets this statement sink in *without judgement*, then encourages Jeff to explain why he feels this way. "There's such a gap between us blacks and whites. We don't understand each other. They're not developed like we are," he says. "When you drive through there, it's like going back to Africa."
>
> "What has helped create such bad conditions in the black community?" Walters probes. "Is it just because they're black, or are there other reason [sic.]?"
>
> "Their environment affects them a lot," Jeff replies. "They don't want to work. They need to try to get off welfare. I worked for what I have, and they could work hard and make it, too."
>
> Walters steers the conversation towards specifics: "Do you know any blacks trying to get off welfare?" Jeff says he know a black woman living nearby who is struggling to find work. He's let her use his phone to set up job interviews. Talking about the woman's situation—no phone, no car, kids who need care—he realizes that his own experience with this is at odds with his beliefs. Without prompting, he begins to question his own sterotype [sic.] that "blacks are lazy."

> The interviewers' *nonthreatening questions and accepting manner* make Jeff comfortable. At their invitation, he explores his views about the Los Angeles race riots and the beating of Rodney King. In the give-and-takes that follows, he's surprised to learn that whites and Latinos were also arrested; that blacks had also risked their lives to save whites; that black unemployment is more than double that of whites; and that more than 44 percent of African-American children live in poverty.
>
> Jeff is asked why he thinks these differences exist. His answer is a breakthrough: "It's partially racial. I know guys doing the same work as me get three or four dollars less an hour. It's not fair. Blacks are almost still in slavery a little." (Imlay and Howard 2006, 1 ff., emphasis added)

Walters does not express his opinions or attitudes. He is not participating in argumentation. He is listening

> with an open mind and heart, and [asking] questions that help us get beyond fear and anger based beliefs; beyond labels like conservative and liberal. Deep listening takes us to a place where we are all human beings with the same basic need for understanding, mutual respect, fairness, and compassion. In that way, deep listening is both a highly effective process for social change, and a spiritual practice.[12]

There is no question that Walters is sincerely interested in the needs and ideas of those he listens to. But there is a kind of insincerity in Walters' "nonthreatening questions and accepting manner." First, because the questions have to be asked with an "open-mind and heart," Walters can't represent himself as disagreeing –which he presumably does—with Jeff's initial claims about the source of difficulties in the nearby black community.

Second, Walters does not give a complete representation of his purposes in asking Jeff the questions he's asking. Furthermore, he exploits Jeff's ignorance of these purposes in Walters' attempts to fulfill them. This makes the method a version of Le Fevre, et al.'s "pseudo inquiry" (discussed in the Preface). Again, the techniques of the listening project are supposed to "change people"—the people you are listening to. The techniques are designed to get people to give up stereotyping beliefs and start seeing things from a perspective more attuned to social justice concerns. This purpose is precisely the kind of purpose which, if it became known to Jeff, would put him on his guard, make him feel threatened, and thwart the attempt to revise his harmful attitudes.

---

[12] See the "About Us" section of the Listening Project's official webpage: http://www.listeningproject.info/aboutus.php.

The Listening Project requires that you be and represent yourself as sincerely interested in learning from the people you're talking to. But the way you're interested in learning from them is not the way you represent yourself as interested in learning from them. Presumably what gets the speaker to lower their defenses is not that they think you're interested in learning about them in some anthropological way, but that you're interested in learning from them something about the issues they're talking to you about. It would be disingenuous for Walters to insist he's really interested in getting insight from Jeff about whether the bad conditions in the black community are explained "because they're black or for some other reason." Walters, rightly, has his own opinions about that. Jack is not going to convince Walters that "it's because they're black."

This is also a difficulty with one common invocation among the political Left after (and before) the 2016 U.S. Presidential election that Clinton-supporters should treat Trump supporters with more respect by seeking to understand their concerns. For example, Michael Lerner implores us "to reach out to Trump voters in a spirit of empathy and contrition" (2016). And Charles Camosy "urge[s] us to look harder at the numbers and do the often difficult work of carefully listening to those with whom we disagree" (2016).

Regardless of the virtues of this compassionate listening, it easily comes with a certain amount of disingenuousness. Clinton-supporters presumably do not abandon opposition to Trump when they start compassionately listening to Trump voters. In particular, they may continue to believe that much of what Trump says is monstrous and that support for Trump, even if not motivated by sympathy for the more monstrous aspects of Trump's rhetoric, is at best culpably ignorant of the harm that is caused by encouraging that rhetoric. Nor do they come to doubt that Trump is disqualifyingly ignorant of many of the issues and demands of the presidency and that his inability to control his temperament makes him dangerously unfit for office. Compassionate listening is hard to do while accurately representing yourself as believing these things. The Left might seek to understand the motivations of Trump supporters by listening compassionately. But if they represent themselves not just as interested in the motivations of Trump supporters for purely anthropological or strategic reasons, but as interested in learning something from them about why voting for Trump was a good thing to do, their representations of their attitudes will be deceptive.[13]

There's a third way that compassionate listening of the relevant sort may be deceptive. Return to Walters' line of questioning of Jeff. Though Walters does not actively represent himself as open-minded, he gives the impression that he is. For it is essential to getting the speaker to lower their guard and keep talking themselves into changing their attitudes that they don't feel defensive. And it is impossible not to feel defensive when you are talking to someone about whom you know that, even if they can't figure out what's wrong with what you're saying, they are unwilling to reduce their confidence that you're wrong.

---

[13] If Clinton-supporters accurately represent themselves as believing in those ways, they risk undermining their appearance of compassion. This, as I discuss more extensively in the next section, risks undermining the efficacy of compassionate listening. It makes it harder to get Trump voters to open-up and thereby allow the Left to "rebuild its political base by helping Americans see that much of people's suffering is rooted in the hidden injuries of class and in the spiritual crisis that the global competitive market generates" (Lerner 2016). But if Clinton-supporters don't accurately represent themselves in these ways, they risk exploiting their deceived interlocutors in a way that manipulates Trump-voters into having the attitudes the Left wants them to have. The only alternative—the only way to be both non-manipulative and effective—is to give up the conviction that voting for Trump is problematic in all these ways. Assuming that Clinton-supporters know that voting for Trump is problematic in all these ways, the conclusions of Chapter 5 show us that they should not give up that conviction so easily.

Do Clinton-supporters know that voting for Trump is problematic in all these ways? Some think they do. Here is the closing of the *NY Times* columnist Charles Blow's Thanksgiving Day column addressing then President-Elect Trump:

> No, Mr. Trump, we will not all just get along. For as long as a threat to the state is the head of state, all citizens of good faith and national fidelity—and certainly this columnist—have an absolute obligation to meet you and your agenda with resistance at every turn.
>
> I *know this in my bones*, and for that I am thankful. (2016, emphasis added)

Walters claims that you need to ask questions "with an open mind and heart," but as I noted, what your mind and heart need to be open to is the fact that the one you're listening to is valuable, "with the same basic need for understanding, mutual respect, fairness and compassion" as any other human being. It is doubtful that by insisting that we need an open mind and heart that Walters means that we need to open our mind and heart to the possibility that blacks are ultimately at fault for their misfortune or, worse, that they are subhuman—that this is a possibility that deserves serious consideration. There is no hiding from the fact that the techniques of the Listening Project are deceptive in multiple ways, including in representing the listener as more open-minded than they really are, and that they exploit the deceptiveness in achieving their goals.

The kinds of methods used by the Listening Project are also employed in other contexts besides the search for social justice. Here's a passage from Jeff Giesea's column in the Harvard Business Review reflecting on the difference in corporate settings between calling someone out and what Giesea calls "calling forth":

> Pretend that you occasionally lose your temper in meetings, and my aim is to get you to change. The next time I see you lose your cool, I say one of two things.
>
> > *Hey, timeout. You just did it again — you lost your temper with Mario. This is the third time I've seen you do this in the last two days. C'mon, this behavior HAS to stop.*
>
> or:
>
> > *Hey, can we chat for a sec? I noticed you just lost your temper with Mario. Did you notice that too? You are so good at running these meetings, I can only imagine how much more effective you're going to be as you move past this behavior. What can I do to help?*
>
> In the first, I'm calling you out. In the second, I'm calling you forth. The content is similar. The messaging and tone are quite different. Which do you think is more effective? For most people and circumstances, it's the latter." (2014)

As with the Listening Project, Giesea's "aim is to get you to change." The recommended techniques mirror the Listening Project's as well: adopting a friendly, rather than a challenging tone; professing interest in the needs of the speaker; asking questions rather than issuing judgments and commands—all done to get the speaker to lower their defenses. Also, as with the Listening Project, calling forth easily fails to accurately reflect your attitudes. For, certainly, calling forth and calling out seem to express different attitudes. Calling forth expresses sympathy. Calling out expresses hostility.

It might be objected that in the relevant cases the representation is not really inaccurate. Walters might say, correctly, that he really is deeply interested in what the speakers have to say. In engaging with students, a teacher might protest, they aren't lying when they tell the students that their arguments are "promising" and that "there's a lot of interesting material in the paper." There are lots of ways to accurately represent a single attitude. If you hate my haircut, there are a number of things you might say, some more insulting and some less: "That's disgusting," "It's a bit shorter than I'd go with," "Do you like it?", "It's really different from your last one," "It's great!" The last is a lie, but is the first so obviously a more accurate representation of your attitudes than the others? Does the honest representation of your attitudes inevitably include the most insulting, most off-putting one?

Actually, yes. When you tone down the rhetoric in order not to offend, you depart, however permissibly, from the truth about your attitudes. This is as true in pedagogy as it is in social justice mediation as it is in the corporate world as it is in tonsorial evaluation. When you tell a student that their paper is promising, it's usually true that it contains a germ that could be developed by the right hands into an excellent paper. So it's not like you outright lie. But in selecting what you tell your student in order to avoid doing psychic damage, you misrepresent your attitudes. This is, of course, a good thing. Not all classes should be Faye Halpern's "Honesty Class"—a fictional class we can only fantasize about offering in which students sign up to hear their professors' completely unfiltered impressions about their work and in-class conduct. There is a difference between the Honesty Class and all other classes we offer. In those other contexts, you do not accurately represent your attitudes and, in situations that don't involve the kinds of formal roles discussed earlier, you allow and encourage a misleading inference to what your attitudes really are.

Is that o.k., as long as the deception is effective?[14] I argue not. Representing yourself as open-minded when you're really closed-minded—or even failing to represent yourself as closed-minded—violates potential obligations you have to your interlocutor and to your audience. It violates obligations to your interlocutor because your interlocutor might very well modify their style of argumentation depending on whether they take you as friendly or hostile, as open or closed-minded. They might be less careful, more revealing, less guarded if they take you as open-minded. It also is simply disrespectful. It allows you a feeling of power and superiority over your interlocutor and the listening public that is improper to indulge in.

Most importantly, it is problematically manipulative of your interlocutor and your audience. If you falsely represent yourself as open-minded for the purpose of preventing your audience from being swayed by your opponent's argument then, though you know the argument is misleading, you are falsely representing yourself in order to manipulate your audience into believing what you know to be true. This is treating your audience as a means *par excellence*. As Robert Noggle puts it,

---

[14] I don't mean to suggest that the deception is inevitably effective—that audiences and interlocutors are more easily convinced when they think you're open-minded than when they think you're closed-minded. Presumably whether open-mindedness or closed-mindedness is more convincing will depend on lots of independent factors: how likable or charismatic you and your interlocutor, respectively, are; how convincing your dismissal is; how defensive you sound in representing yourself as closed-minded and how insincere you sound in representing yourself as open-minded.

> Acting manipulatively toward someone, then, is an affront to her as a rational and moral being; for it is an attempt to thwart her moral and rational agency, which has as its goal the correct adjustment of her psychological levers. To attempt to thwart the goals someone has qua rational moral agent is to fail to respect her rational moral agency. And since a person's rational moral agency is crucial to her personhood, to fail to respect it is degrade her; it is to treat her as less than a person. And for that reason it is wrong. (1996, 52)

There are views of manipulation according to which the kind of deception on offer may not seem to count as manipulative or, at least, not wrongly so. Many philosophers view manipulation as, necessarily, an interference with a subject's rational capacities[15] or as a way of departing from what are in fact good reasons.[16] But, it might be thought, when you misrepresent yourself as open-minded in order to make your interlocutor and audience more receptive to you, you're not interfering with their rational capacities nor departing in any way from what are good reasons. Rather, the hope is, you're preventing your interlocutor and audience from being blocked by prejudice, bias, and other problematic affective factors from being responsive to you, your questions, and any good reasons you might offer. You're allowing them to respond to reasons and exercise their rational capacities by preventing them from being misled by distorting factors.[17]

---

[15] See, for example, (Baron 2003) and (Noggle 1996).

[16] See, for example, (Gorin 2014) and (Mills 1995).

[17] A similar contention might distinguish misrepresenting yourself as open-minded from a different case of problematic misrepresentation emphasized by Jennifer Lackey. Lackey points out that in some cases of isolated second-hand knowledge it is often wrong to outright assert that *p* even when you know that *p*. If Josie writes in a letter of recommendation that "Mitchell has very polished writing skills," but only because she knows it on the basis of "isolated, reliable testimony of her trustworthy colleague," then she has asserted improperly (Lackey 2011 259). She has asserted improperly because

> it is utterly reasonable for readers of her letter to presume that the asserter has the status of being a witness—which involves the possession of either firsthand or some other kind of privileged information—to the issue or event in question. For instance, given that Josie is offering a flat out assertion about the abilities of a student in a letter of recommendation, it is utterly reasonable for readers of her letter to presume that she experienced Mitchell's writing firsthand. (Lackey 2011 260)

Even if such a tactic counts as “manipulation” in some sense, it may not seem to be morally problematic. Is it any more problematic than, say, the manipulation involved in Marcia Baron’s example of “‘mood enhancers’ such as boiling some water infused with cinnamon or vanilla to release a delicious, comforting aroma before showing the house to potential buyers?” (2014, 114) On the contrary, it may seem less problematic. Unlike mood enhancers, your appearance of open-mindedness may not present any specific view or set of reasons in a particularly favorable light (in the way that the smell of cinnamon might make a house seem more cozy than it really is). It seems, rather, relevantly similar to the following case from Moti Gorin’s 2013 PhD dissertation:

> *Interview Anxiety*: Gabriel is in need of employment. His friend Simon has a good job and speaks to his boss about Gabriel, who has just the sort of experience Simon’s boss is looking for in a new employee. Simon convinces his boss to agree to meet with Gabriel for a job interview. Given his dire financial situation Gabriel has good reason to meet with Simon’s boss. The problem is that whenever Gabriel reflects on his finances he becomes paralyzed with anxiety. Consequently, if Gabriel were to associate his financial problems with his holding a meeting with Simon’s boss he would not make it to the interview. (2013, 104)

As Gorin goes on to say, “Simon might invite Gabriel to join him at his office for a holiday party and use this opportunity to introduce Gabriel to his boss, so that in preparing to go to Simon’s office Gabriel will be relaxed, not anticipating that he soon will be discussing financial matters” (104). (Assume that Gabriel will be emotionally fine once he’s actually in the discussion; it’s the anticipation that creates Gabriel’s paralyzing anxiety.) If Simon’s deception is problematic, it seems minimally so. One of the reasons it is only minimally problematic is that Gabriel plausibly

---

It might be objected that the distinctive feature of Lackey’s case is that Josie causes her readers to believe on a false basis. In falsely representing herself as a first-hand witness, Josie uses misleading evidence to get her readers to believe what she wants them to believe. But falsely representing yourself as open-minded doesn’t manipulate your audience into believing what you know to be true in the same way (it might be said). All the misrepresentation does is ensure your interlocutor and audience are appropriately receptive to you and your questions. They don’t base their resultant beliefs on the claim that you were open-minded toward them.

wishes there were some way to overcome his paralyzing anxiety and so wishes he could have a way to meet with the boss without being initially aware he is doing so.

Of course, if Gabriel likes that he suffers paralyzing anxiety when thinking about discussing financial matters, then it would more than minimally problematic for Simon to deceive Gabriel in the hopes of circumventing that paralyzing anxiety. Simon would be preventing Gabriel—without Gabriel's consent[18]—from doing something that Gabriel likes that he does. It's only because Gabriel doesn't want to suffer paralyzing anxiety when thinking about discussing financial matters that Simon's behavior stands a chance of being only minimally problematic.

When you falsely represent yourself as open-minded in order to ensure your interlocutor and audience is receptive to you, you do something more problematic than deceptively circumvent unwanted anxiety. It is rational to want to be unreceptive to those whose minds are closed against you. If your interlocutor and audience were fully aware of your closed-minded attitudes toward their arguments, it would make sense for them to want to be unreceptive to you. It's not just that, if they knew your real attitudes, they would in fact be unreceptive. It's that they would rationally endorse that unreceptiveness: "if you aren't going to give me the respect of serious engagement, I'm not going to give it to you." So, in deceiving them, you are preventing them from having attitudes they would otherwise both have and rationally endorse, and you are doing so without their consent.[19]

---

[18] As Allen Wood says, in discussing what's problematic about manipulation,

> when getting others to do what you want is morally problematic, this is not so much because you are making them worse off (less happy, less satisfied) but, instead, it is nearly always because you are messing with their *freedom*—whether by taking it away, limiting it, usurping it, or subverting it. By "freedom," I mean the capacity of a rational human adult to govern his or her life, rather than having it subject to the will of someone else. (2014, 17 ff.)

[19] It might be that, in the particular case, it is in your interlocutor and audience's interest to be receptive to you, even though they might deny it, were they to know about your closed-mindedness. But, arguably, being receptive to those

It's possible for there to be people who—like Gabriel—would wish to be just as receptive to you, whether you are open-minded or closed-minded. They might regret the fact that if you are closed-minded, they will be less receptive to your questions and arguments. To such people, deceptively representing yourself as open-minded might be permissible. But there are lots of standard cases in which your interlocutor and audience don't endorse being deceived in this way. In such cases, your deception is, at least *pro tanto*, impermissible.[20] It is a wrong to them.

The wrong to your interlocutor and audience can be outweighed or eliminated. If you are a member of a marginalized group, and there is no other way to escape persecution or reprisal than to make yourself out to be more open-minded toward someone's argument than you really (and rightly) are, then it can be too great a demand to place on you that you put yourself at risk.[21] Nor is this a mere excuse for a wrongdoing; you are not just excused, but can be justified in representing yourself falsely in such situations. The wrong might also be outweighed if there is an extraordinary amount at stake or if your interlocutor or audience doesn't deserve to be treated with honesty. For example, it might be the case that unrepentant Nazis have lost the right not to be deceptively manipulated into changing their minds.[22] But there are lots of situations in which the wrong of deception is not so easily outweighed. In many situations there is no imminent risk

---

who are closed-minded toward us is not typically in our interest. Therefore, that this is a case of problematic manipulation is a view that can be endorsed by Anne Barnhill (2014).

[20] This result can be allowed by Jennifer Blumenthal-Barby, who argues that the moral status of manipulation depends on, among other things, "the extent to which the instance of influence . . .threatens or promotes autonomy" (2014, 122). Autonomy is threatened or promoted, according to Blumenthal-Barby according to the "extent to which the person influenced is aware of and endorses, or were he aware of would endorse, the non-argumentative influence attempt as a process by which his desires, decisions, or actions were formed" (126).

[21] This point was suggested to me by a question asked by an audience member during Erik Anderson's presentation of his (2016). Anderson defends a "Principle of Sincerity" similar to the principle I defend here.

[22] This is the basis on which Christopher Nathan (forthcoming) argues that undercover police work can be morally permissible—when it is used against suspects whose criminal behavior has made them liable to such tactics. When it comes to public engagement, it seems to me, run-of-the-mill argumentation requires severe culpability of belief in order to make the believer liable to deceptive manipulation of the sort discussed in the body of the text. Thanks to Beige Lussier for suggesting the comparison between undercover police work and misrepresenting yourself as open-minded.

of harm to you, your interlocutor and audience haven't lost the right to be treated with honesty, and there isn't an extraordinary amount at stake (even if the issue about which you are arguing is of major import).

Giesea, in advocating for calling forth, suggests that, while you should call forth rather than call out, you also should have the attitude that calling forth would accurately represent. As Giesea says in justifying his position, calling people out

> comes from a place of ego or reaction. The intent, conscious or not, is to make the other person wrong. . . . Calling people forth, in contrast, comes from a place of service and an open heart. . . . Calling forth is a mindset, a way of showing up as a leader who fights for the greatness of within others. It starts with intention. (2014)

The virtue of calling forth as opposed to calling out is not just public; it's that it represents a better attitude. If you have the attitude that is accurately expressed by a calling out rather than a calling forth, according to Giesea, you should change your attitude to the one that is accurately represented by the language of calling forth.

This easy answer comes if the effective expression is the one that would accurately represent the attitude you should have. Perhaps this is the direction that members of the Listening Project would endorse as well: make your behavior accurate by changing your attitudes; become genuinely open. But, as I argued in the last chapter, knowers should be closed-minded in standard situations. Therefore, if effective expression is the expression that would accurately represent yourself as open-minded, the easy answer isn't available. What, then, should you do? If you're going to engage, there is no choice but the less effective option: sincere expression of your closed-mindedness.

7.3 THE PITFALLS OF SINCERITY

Sincere expressions of closed-mindedness can be ineffective if your goal is to convince those who disagree with you. But other goals might be thought to require honest closed-minded intervention. If your interlocutor has acted or spoken offensively, or is offering an argument that is demeaning to members of a marginalized group, or is making correctable errors, then you can respond in a way that makes clear you are not willing to reduce your confidence in response to what was said. Doing so might be the only way to demonstrate that the offensive actions or words are unacceptable, stand in solidarity with members of the marginalized group, stand up to the marginalizing speaker, or correct the errors.[23]

Donald Trump, for example, was clearly in the wrong when he declined to correct a questioner in a 2015 U.S. Republican Primary debate who commented, "We have a problem in this country, it's called Muslims. We know our current president is one." Trump's response? "Right." After being called out by a number of public figures and media personalities, Trump replied via a series of tweets: "This is the first time in my life that I have caused controversy by NOT saying something" and "If I would have challenged the man, the media would have accused me of interfering with that man's right of free speech. A no win situation!" Trump's protests aside, correction was the only proper response in this case.[24]

---

[23] This is why it is inappropriate to take public closed-minded declarations of moral principle as indicating that those making those declarations are guilty of what Justin Tosi and Brandon Warmke call "moral grandstanding," the practice of making "a contribution to public moral discourse that aims to convince others that one is 'morally respectable'" (2016, 199). There are surface features of speech that might be taken by some to indicate what Tosi and Warmke think are central cases of grandstanding—"(1) piling on; (2) ramping up; (3) trumping up; (4) excessive emotional displays or reports; and (5) claims of self-evidence" (203). But speech can have most of those surface features when it is aimed at ensuring that those victimized by the targets of the putatively grandstanding speech know that they are supported. And it can also have those surface features when the speech comes from marginalized speakers themselves, and when that speech is aimed at resisting their oppression.

[24] When John McCain made the better choice during a campaign rally in 2008—correcting a woman who claimed that Obama was "an Arab"—he received applause from the crowd (though McCain corrected the woman via an unfortunate contrast between being an "Arab" and being "a decent family man" and "citizen"). However, when he earlier attempted to correct a man who claimed to be scared of an Obama presidency—"each time [McCain] tried to cool the crowd, he was rewarded with a round of boos" (Martin and Parnes 2008). Media reports are routinely misleading about what McCain received boos for. Some have misleading headlines: "McCain: Obama not an Arab, crowd boos" (Martin and Parnes 2008). Some have misleading content: "Sen. John McCain drew boos at a town hall

The kind of honest closed-minded engagement that most naturally springs to mind in such cases is rather direct and might be described as rude. Honest closed-minded engagement is not rendered impermissible just because it's rude. For one thing, even rude engagement may be called for and obligatory in some cases, if the person you're engaging with is behaving despicably and their conduct needs to be denounced.[25] Sometimes people need calling out, especially when, as I've noted, this kind of direct approach is the only thing that can make marginalized groups assured that the non-marginalized stand in solidarity with them.

The requirement that you always respond toward obvious falsehood with polite, rational argumentation, even if honest about your closed-mindedness, doesn't do justice to the demeaning effect that satisfying this requirement can have on those who are harmed by certain kinds of falsehoods. To think that there is a requirement that women, blacks, or Jews have to respond rationally and politely to blatant sexists, racists, or anti-Semites—to think that sexists, racists, or anti-Semites must be engaged in rational argumentation—doesn't do justice to the experiences of the marginalized. Sometimes the right response to hateful speech is not rational discourse, but protest. Protest, as Bernard Boxill points out,

> is not an argument. . . . Typically, people protest when the time for argument and persuasion is past. They insist, as Dubois put it, that the claim they protest is 'an outrageous falsehood,' and that it would be demeaning to argue and cajole for what is so plain. Responding to a newspaper article that claimed 'The Negro' was 'Not a Man,' Frederick Douglass disdainfully declared, 'I cannot, however, argue, I must assert.' (1976, 64)[26]

---

when he corrected a woman who called then-Sen. Obama 'an Arab'" (Slattery and Greene 2015). Some misleadingly cut video: for example, changing the order in which McCain made various remarks (see http://www.cnn.com/videos/politics/2015/09/18/mccain–2008-presidential-campaign-audience-question-on-obama-as-arab.cnn).

[25] See (Kriss 2016).

[26] Thanks to Michael Randall Barnes (2016) for directing me to this passage.

Even when rational argumentation is required and rudeness is impermissible, honest closed-minded engagement might be called for if honest closed-minded engagement can be performed politely. One way to accurately represent yourself as closed-minded while not being challenging about it is to follow a procedure set out in Marshall Rosenberg's Nonviolent Communication process. Nonviolent Communication has features similar to the procedure developed by the Listening Project, but Rosenberg advocates a mutual honesty of expression that the Listening Project doesn't emphasize. While the Listening Project emphasizes passive techniques for drawing the speaker out until they change their mind, Rosenberg's "NVC Process" is designed to get the listener and the speaker to understand each other about their feelings and needs. Using his process, says Rosenberg, you can effectively resolve conflicts arising in "intimate relationships, families, schools, organizations and institutions, therapy and counseling, diplomatic and business negotiations, and disputes and conflicts of any nature" (2003, 8).

For example, if you are confronted by someone saying explicitly derogatory things about Jews, the process would require you to attend, nonjudgmentally, to what you are observing and feeling, to ascertain what needs drove you to feel that way, and honestly figure out what the person would need to do to satisfy those needs. Only then can you honestly impart these conclusions to your interlocutor and elicit their observations, feelings, needs, and requests. To illustrate, Rosenberg tells a story about sitting next to a man at an airport who said, "These kikes get up early in the morning so they can screw everybody out of their money." After sitting back for a moment to allow himself to enjoy "the image of actually grabbing his head and smashing it," Rosenberg decided to

> empathize with him, to hear his pain. Why? Because I wanted to see the beauty in him, and I wanted for him to fully apprehend what I had experienced when he made his remark. I know I wouldn't receive that kind of understanding if there were a storm brewing inside of him. My intention was to connect with him and show a respectful

> empathy for the life energy in him that was behind the comment. My experience told me that if I were able to empathize, then he would be able to hear me in return. It would not be easy, but he would be able to. (2003, 150)

After an extended discussion of the speaker's experiences and feelings that led to his comment, instead of saying, "That was really anti-Semitic," Rosenberg said,

> Let me repeat what I'm trying to say. I really want you to just hear the pain I felt when I heard your words. It's really important to me that you hear that. I was saying I felt a real sense of sadness because my experiences with Jewish people have been very different. I was just wishing that you had had some experiences that were different from the ones you were describing. Can you tell me what you heard me say? (2003, 152)

Rosenberg's goal is to get those he is intervening with to understand him. He does not say that his goal is to change their minds. This is in keeping with his contention that "the ultimate goal is to spend as many of my moments as I can in that world that the poet Rumi talks about, 'a place beyond rightness and wrongness'" (Rosenberg 2012, 82). But whether the goods we're supposed to get through Nonviolent Communication are epistemic or nonepistemic, the idea is that through total honesty with respect to your feelings, needs, and requests, it is more likely that others aren't blocked by their defenses from choosing the path you are hoping they will choose. Likewise, if this is right, then total honesty about your closed-mindedness and your reasons for being closed-minded, expressed with love and compassion, should not stand in the way of your interlocutor being convinced by you.[27]

There is not much data on the effectiveness of Nonviolent Communication. Elaine Fullerton notes that

> The main difficulty with the NVC model is that there is a lack of longitudinal analytical research. There are anecdotal examples of the outcomes of this process, and many of these are consistent and could confirm the plausibility of the hypothesis that: the NVC model does support conflict resolution. . . . There is however a lack of longitudinal research and a lack of research that specifically tests the theoretical basis of the model. (2008, 18)

---

[27] This is perhaps what those mentioned in the opening of the last chapter—who go out of their way to tell readers that they won't be convinced out of their views—are hoping to achieve by their disclaimer.

What is more, there is some evidence that techniques like those employed in Nonviolent Communication are singularly ineffective in eliminating prejudice. In Calvin Lai and Maddalena Marini, et al.'s (2014) study of 17 methods of intervention, only 8 of 17 methods were shown to be effective at reducing implicit racial preference. But techniques that aimed at exposing subjects to alternative perspectives and at training subjects in empathy—the techniques employed by Nonviolent Communication—were the least effective methods among the 9 methods judged *ineffective* (Lai, Marini, et al. 2014, 1779). Even the 8 effective methods were ineffective in reducing explicit racial preference, while in Lai, et al.'s follow-up (2016) study, none of the effects were shown to be lasting even after only days or, sometimes, hours. If Lai, et al., are right, Nonviolent Communication should be the least effective among the least effective methods.[28]

As usual, the empirical data is conflicted. In a study emphasizing the importance of presentation of alternative perspectives, David Broockman and Joshua Kalla—testing a method developed by the Los Angeles LGBT Center Leadership LAB[29]—found that a single 10 minute conversation had measurable and positive impact on subjects' attitudes toward transgender people: "After the intervention, the treatment group was considerably more accepting of transgender people than the placebo group" (2016, 223). The effect was still significant after 3 months (222) and was strong enough to withstand attack: both immediately after and 6 weeks after being shown attack ads, treatment subjects remained .3-.4 scale points more supportive of the non-discrimination law than placebo subjects (223).

---

[28] Some guides designed to help activists deal with bigoted attitudes and behavior do not have the kind of principled theoretical backing of the Listening Project or Nonviolent Communication. The Southern Poverty Law Center's (2005) *Speak Up!* manual at times recommends outright confrontation (when it's safe), at times non-judgmental listening, at times heartfelt attempts at reconciliation. What behavior is appropriate depends on context and audience. When confronting a stranger who has expressed a bigoted attitude, after you "assess your surroundings" to ensure that confrontation is safe, your intercession can be direct: "'I find that language very bigoted. It offends me.' Or, 'I think it's wrong to stereotype people.'" (70–71) When speaking to a family member, they recommend that you "speak up without 'talking back' (16), "ask clarifying questions: 'Why do you feel that way?'" (17) or "Appeal to parental values. . . .'Dad, when I was growing up, you taught me to treat others the way I wanted to be treated. And I just don't think that term is very nice'" (16–17). Always there is careful negotiation between being deceptive about your attitudes, alienating the person with whom you're intervening, and trying to achieve the best results.

[29] See their website, here: http://www.leadership-lab.org/. For a representative video of the method in action, see http://www.leadership-lab.org/2016-trans-equality-conversation-steven/.

Here's the method used by canvassers in Broockman and Kalla's study:

> Canvassers informed voters that they might face a decision about the issue (whether to vote to repeal the law protecting transgender people); canvassers asked voters to explain their views; and canvassers showed a video that presented arguments on both sides. Canvassers also defined the term "transgender" at this point and, if they were transgender themselves, noted this. The canvassers next attempted to encourage "analogic perspective-taking." . . . Canvassers first asked each voter to talk about a time when they themselves were judged negatively for being different. The canvassers then encouraged voters to see how their own experience offered a window into transgender people's experiences, hoping to facilitate voters' ability to take transgender people's perspectives. The intervention ended with another attempt to encourage active processing by asking voters to describe if and how the exercise changed their mind. (221)

Like the Listening Project, much of the method employed in this study emphasizes simply asking questions and listening: asking voters to explain their views, to talk about times they were judged negatively, to describe if and how the exercise changed their mind. However, there is also more emphasis than in the Listening Project on the presentation of argument—in videos showing arguments on both sides, and then again when "canvassers then encouraged voters to see how their experience offered a window into transgender people's experiences."

It is less clear to what extent canvassers accurately represented their attitudes.[30] It was presumably clear from context what beliefs the canvassers held about transgender rights; they introduced themselves to voters as volunteers from SAVE, an LGBT organization. It may not, however, have been clear from context whether the canvassers were open-minded or closed-minded. Videos of the technique in action reveal the canvassers engaged in a good deal of smiling-and-nodding as well as offering well-placed phrases like, “I think it’s really complicated”[31] when it seems clear to the viewer that the canvasser is nowhere near as conflicted about the truth of their subject’s claims as is indicated by the canvasser’s body language and vocalizations. When a subject says, “A man and a woman should be married and that’s what it says in the Bible,”[32] the canvasser does not presumably think the sentiment is the kind of food for thought their nodding suggests they think it is.

---

[30] Commenting on a video of the method eventually tested by Broockman and Kalla, LAB describes the context as one in which "the team was just beginning to zero in on a brand new same sex marriage conversation 'recipe'." Recipes for conversation are not the sort of thing that those who are accurately representing their attitudes need to zero in on. See http://www.leadership-lab.org/2010-marriage-conversation-laura/.

Still, canvassing is a relatively formal setting. As in formal debate or pedagogical settings, tone and mannerism can go a good distance toward shielding the attitudes of the canvasser from licensed inferences; there is a norm that requires that canvassers act in a relatively prescribed manner regardless of their actual attitudes. In not accurately representing themselves as closed-minded, if indeed they were, the canvassers need not have been problematically deceptive. But the methods tested and found effective by Broockman and Kalla are less the honest representation of attitudes involved in Nonviolent Communication and more the partially deceptive methods of the Listening Project.

So, if you’re trying to convince your interlocutor, there may be a way to steer between the twin dangers of closed-minded engagement. The main way to do it is by hiding your closed-mindedness in an appropriate formal setting. This, if certain empirical results are borne out, can be effective in convincing interlocutors and audiences. Honesty about your closed-mindedness can be effective, too, but primarily if your goals are to call a speaker out. The difficulty is that a lot of standard situations don’t satisfy either of these requirements. In those cases, you are trying to convince an interlocutor or an audience, and the only way to do it is to falsely represent yourself as open-minded. Because the settings are not of the appropriately formal sort that prevent licensed inferences from your behavior to your attitudes, the exploitation of misled beliefs about your attitudes to convince your interlocutors or audience is impermissible. Therefore, in many standard situations, you should not engage closed-mindedly with relevant counterarguments. Thus, if you know that the relevant counterarguments are misleading, you shouldn't engage with them at all.

---

[31] See https://leadership-lab.org/2016-trans-equality-conversation-steven/ @ 3:08.

[32] See https://leadership-lab.org/2010-marriage-conversation-laura/ @ 0:28.

When the conventions of the context do not allow deception, you can only permissibly engage closed-mindedly by being honest about your closed-mindedness. This can be a difficult trick to pull off effectively, if your goal is to convince others, since honesty about your closed-mindedness is just the sort of thing that can alienate your listener and put them on the defensive. If the goal of closed-minded engagement, however, is not simply to convince your interlocutor, but to show solidarity with groups marginalized or victimized by the interlocutor or show others that their speech is inappropriate, then this may be exactly what you should do. Depending on your goals, if what you have to say isn't nice, you should sometimes say it anyway.

8

On Inviting Problematic Speakers to Campus

> [Singer] proceeds to recount bits of Princeton history. He stops. "This will be of particular interest to you, I think. This is where your colleagues with Not Dead Yet set up their blockade." I'm grateful for the reminder. My brothers and sisters were here before me and behaved far more appropriately than I am doing. (Harriet McBryde Johnson (2003, 74))

> But why the hell were they inviting that hateful idiot to speak in the first place? (Rebecca Kukla)[1]

8.1 INTRODUCTION

One venue in which it might seem that there are special duties to engage with relevant counterarguments is academia. Both inside and outside the classroom, there is something particularly repugnant about stifling unorthodox opinion. Academia should be a forum for engaged and open-minded debate; if we refuse to invite speakers on the basis of the content of their speech or if we fail to expose students to a wide range of controversial views, we violate the ideal of the open academic community.

Call a "problematic speaker" a speaker whose invitation is viewed as demeaning or marginalizing to some individual or group. On this definition, subjective attitudes determine whether a speaker is problematic. All it takes to be a problematic speaker in this sense is to be a speaker whose invitation is *viewed* as demeaning or marginalizing. Because of this, it should be uncontroversial both that it is wrong to prohibit invitations to all problematic speakers and that it can be permissible to invite problematic speakers. A problematic speaker, on this definition, is just one whose invitation raises questions of permissibility. The questions, however, might be easily answered if the speaker's invitation is only irrationally viewed as marginalizing.

---

[1] Posted on Facebook in response to protests at Middlebury College of invited speaker Charles Murray. Quoted with permission.

A second feature of this definition is that problematic speakers need not be problematic only because their *speech* is (viewed as) marginalizing or demeaning. Problematic speakers are problematic because *invitations* to those speakers are (viewed as) marginalizing or demeaning. A speaker's invitation might be viewed as marginalizing or demeaning if they haven't made outright claims that are demeaning to any group or individual, but have committed acts that make their invitation marginalizing or demeaning to a group or individual. For example, if a speaker has assaulted a student or been found guilty of sexual harassment, their invitation will most likely be viewed as marginalizing or demeaning at least to those that they have harassed or assaulted. The speaker therefore counts as problematic on this definition.

Few universities have speech codes barring invitations to speakers whose history includes marginalizing or demeaning acts. But many universities have speech codes barring invitations to speakers whose speech is marginalizing or demeaning. The Foundation for Individual Rights in Education (FIRE), whose self-described mission "is to defend and sustain individual rights at America's colleges and universities" (FIRE 2017a) publishes an annual report summarizing their research into the degree to which American universities restrict speech on campus. According to their most recent report, 39.6% of 449 schools surveyed "maintain severely restricted, 'red light' speech codes that clearly and substantially prohibit constitutionally protected speech" (FIRE 2017b, 4).[2] In the United Kingdom, the National Union of Students maintains a "No Platform" policy which prohibits the invitation of speakers who advocate various racist, sexist, or otherwise marginalizing views. Such policies have resulted in the disinvitations of the biologist and public intellectual Richard Dawkins (for retweeting the video "Feminists Love Islamists"),[3] the journalists Julie Bindel (for advocating certain views about transgender women) and Milo

---

[2] For example, an explicit ban on "offensive speech" would count as a "red light" speech code (FIRE 2017b, 5).
[3] See (Blair 2016).

Yiannopoulos (for various tweets and comments about rape culture),[4] and Nick Lowles—director of the group Hope Not Hate—apparently for holding “Islamophobic” views.[5]

Much of the resistance to speech codes in the United States derives, when sincere,[6] from worries that they run afoul of the First Amendment. Legal challenges have prevented or reversed the implementation of some U.S. speech codes, for example at Stanford University and at the Universities of Michigan and Wisconsin.[7] I do not discuss the First Amendment issues except insofar as arguments regarding the First Amendment issues can be used to draw conclusions about the issue I’m interested in. Though I don’t think that a group’s refusal to invite someone or their decision to disinvite someone need be a violation of the First Amendment, this is not what I’m arguing for. Nor am I arguing that we should institute speech codes—that universities should bar invitations to certain problematic speakers. I don’t take a position on that, except to the extent that the premise that it’s wrong to invite such speakers can serve in an argument for the legitimacy of introducing such codes.[8]

---

[4] See (Chasmar 2015). More recently, UC Berkeley cancelled an invitation to Yiannopoulos in response to protests that were becoming progressively more violent, despite the fact that the university had reconfirmed Yiannopoulos’ right to speak. See (Wong 2017) for details. See (Cassell 2017) for evidence that Berkeley students were not responsible for the violence and for discussion of who might have been. Shortly thereafter, Yiannopoulos resigned from his position at Breitbart News after making comments that apparently endorsed sexual relationships between older men and boys as young as 13. See (Helmore 2017).

[5] See (Rawlinson 2016).

[6] Kate Manne and Jason Stanley argue that expressions of worry about freedom of speech are often motivated, not by sincere interest in maintaining freedom of speech, but by a need to maintain the status quo when those with less power object to those with more: “when people lower down in social and institutional hierarchies criticize the speech acts of those higher up, it often reads as insubordination, defiance, or insolence. When things go the other way, it tends to read as business as usual” (2015).

[7] For discussion, see (Gould 2010, 124–128).

[8] Some protesters attempt to silence invited speakers by shouting them down (for example, Charles Murray at Middlebury College (Kranz 2017) and Peter Singer at the University of Victoria (Fagan and Sauer 2017)). The argument marshalled in this chapter could—if sound—potentially be modified by those looking to justify certain kinds of silencing protests of wrongly invited speakers. I take no stand on whether this strategy would work.

My conclusion is not about the ethics or legality of speech codes, but about the ethics of invitations: that it's wrong for groups or individuals to invite certain problematic speakers to campus.[9]

8.2 LEARNING FROM PROBLEMATIC SPEAKERS

If you know that a speaker's words are false, then you can't learn from them the truth of what they're saying. But there are three other things you might be able to learn. First, you might learn critical reasoning skills and how best to marshal your arguments against problematic speech. This seems to be the view, as we've already seen, of Michael Scriven. According to Scriven, the primary goal of education should be training students in critical thinking. Critical thinking is the skill of being able to hear various arguments—new and old—and evaluate them according to the merits—weighing the pros and cons and deciding which arguments use the best evidence. Critical thinking skills should not be limited to those views that you already deny or are on the fence about. Unless they are trained to hone their critical thinking skills in defense even of their most cherished beliefs and values, students

> will have nothing but unthinking commitments which will evaporate under the first warm flush of emotion. Commitments with more backing than that are critical for survival. Hence a school environment that generates more than these superficial values is critical for survival. Thus training for critical thinking is critical for survival. (Scriven 1984, 11)

---

[9] Related issues arise concerning the ethics of accepting invitations from problematic institutions. For example, some performers invited to Donald Trump's inauguration ceremony felt duty bound to decline the invitation. For example, Broadway singer Jennifer Holliday pulled out of the inauguration, stating in an open letter,

> I did not take into consideration that my performing for the concert would actually instead be taken as a political act against my own personal beliefs and be mistaken for support of Donald Trump and Mike Pence. . . . [M]y only choice must now be to stand with the LGBT Community and to state unequivocally that I WILL NOT PERFORM FOR THE WELCOME CONCERT OR FOR ANY OF THE INAUGURATION FESTIVITIES! (Hod 2017)

If the arguments in this chapter are sound, there will be, in many cases, an obligation to decline such invitations when to accept them would require an invitee to knowingly betray some wrongly marginalized or targeted group or individual.

How do you train students in critical thinking? When you train students for deep sea diving, you'd better have them do some deep sea diving. When you train students for desert survival, you'd better have them do some surviving in deserts.[10] And when you train students for critical thinking, you'd better have them do some critical thinking.[11] Nor should you limit that critical thinking to the easy cases. The skills needed to think critically about highly charged, controversial, and deeply personal issues are different from the skills needed to think critically about purely theoretical issues, or issues that don't have much personal resonance. In the former class of cases students need to learn how not to be unduly influenced by ego involvement or emotional investment that they may not have in the latter class of cases. So says Scriven:

> My new requirement is that all the main examples should involve highly controversial issues of considerable personal, social, or intellectual importance that are not seriously addressed in the regular curriculum. Critical thinking is coping with controversy. That's the basis for its value and the nature of its approach. And you can't teach it if you can't freely discuss all sides of controversies. (12)

If Scriven is right, it's not merely permissible to expose students to such controversies; educators have an obligation to do so. His argument is anticipated by the National Association of Scholars, which argues that "Higher education should prepare students to grapple with contrary or unpleasant ideas, not shield them from their content" (National Association of Scholars 1991, 8) and repeated by Dennis Hayes, who claims, "Part of being educated used to mean you could hear any arguments and learn to answer them" (2009, 132).

---

[10] As Scriven says, "If one wants to teach survival skills to would-be astronauts . . . one first tries to set up training conditions as near as possible to the real thing" (9). Jonathan Haidt makes a similar point:

> Bones are anti-fragile. The more you use them, the stronger they get. This is one of the difficuloties [sic.] about going to Mars. We need to find enough challenges for astronauts to keep their bones strong so when they get to Mars they still can walk.
>
> Our immune system is anti-fragile. Keeping kids safe from germs will cripple their immune systems. We needs germs to make this system stronger. (Haidt 2015)

[11] See also (Bailin and Battersby 2016, esp 370–371).

It’s one thing to say that students should be considering controversial issues. It’s another thing to say that universities should invite problematic speakers to argue for one side of those issues. But the skills needed to reflect critically when faced with a flesh-and-blood advocate of controversial positions are different from those needed to reflect critically when considering a controversial position in the abstract. As Mill says, regarding one who advances a view without considering the case for the other side,

> Nor is it enough that he should hear the arguments of adversaries from his own teachers, presented as they state them, and accompanied by what they offer as refutations. That is not the way to do justice to the arguments, or bring them into real contact with his own mind. He must be able to *hear them from persons who actually believe them*; who defend them in earnest, and do their very utmost for them. He must know them in their most plausible and persuasive form; he must feel the whole force of the difficulty which the true view of the subject has to encounter and dispose of; else he will never really possess himself of the portion of truth which meets and removes that difficulty. (Mill 1859/2003, 104 ff. emphasis added)

Therefore, the National Association of Scholars and Hayes draw a stronger conclusion than Scriven’s: we should invite advocates of controversial positions to speak on behalf of those positions. This larger conclusion has been put into practice by Williams College’s Zachary Wood, whose student-run organization Uncomfortable Learning[12] invites problematic speakers to campus in order to hone critical reasoning skills that can only be developed through personal engagement with their views:

> being able to ask questions, make comments and put forth criticisms directly is only possible when there is a public platform for contentious discussion. When we say that a speaker should not come because of their views, we’re denying ourselves an opportunity to strengthen our own arguments. (Wood 2015)

On this argument, problematic speakers teach students how to hone their critical reasoning skills generally, as well as the best kinds of arguments to use in response to particular views they disagree with. But these argumentative virtues don't exhaust the benefits students might get from exposure to problematic speech. There are also more general moral virtues they can hope to develop.

---

[12] Among those speakers invited by Uncomfortable Learning but ultimately disinvited by Williams College were conservative thinker John Derbyshire (see (Kabbany 2016)) and author Suzanne Venker (see (Kabbany 2015)). A similar group—Open Campus Initiative—was formed at Harvard University in 2017 (see (Bauer-Wolf 2017)).

First, according to Lee Bollinger, the presence of problematic speakers on campus fosters in listeners the virtue of tolerance, especially when the speech is known to be false. Thus, Bollinger distances himself from the view that problematic speech should be allowed in order to serve the function of truth-discovery: "Even racist speech and advocacy of genocide are protected, despite the fact that anyone seriously interested in seeking truth would spend no time listening to such ideas" (Bollinger 1990, 980). Granting that there may be precious little truth to learn from certain kinds of hateful speech, Bollinger argues that "under the principle of free speech, we celebrate self-restraint, we create a social ethic of tolerance, and we pursue it to an extreme degree" (1986, 14). In a society whose failings are, if any, a propensity toward intolerance, this is an ethic sorely needed.

Of course, as with suffering generally,[13] the burden of learning to be tolerant falls heavier on some groups than others. The burden of hate speech on its targets, Richard Delgado and Jean Stefancic tell us, includes "psychosocial harms, including depression, repressed anger, [and] diminished self-concept" (2009, 362). In addition,

> The victim of hate speech, especially the one who fears more of the same, may behave circumspectly, avoiding the situations, places, and company where it could happen again. Needless to say, this "cultural mistrust," a mild form of healthy paranoia, has implications for both the mental health and professional chances of minority workers. Other victims will respond with anger, either internalized or acted-out (neither of which is calculated to make things better. (Delgado and Stefancic 2004, 14)

---

[13] Bollinger's view is strikingly similar to John Hick's influential "soul-making" response to the Problem of Evil (see, e.g. (Hick 1981)). According to Bollinger, evil speech is permitted so that we can develop our virtue of toleration in response to evil speech. According to the soul-making theodicy, evil is permitted so that we can develop our virtues in response to evil. A natural worry about this theodicy is that the cost of soul-building falls disproportionately and consistently on certain marginalized groups—the ones who suffer. Therefore, both the theodicy and Bollinger's position are subject to the same worry: the bulk of the burden falls on the victims of the evil or the evil speech, respectively.

Hence the aptness of Mari Matsuda's quip: "Tolerance of hate speech is not tolerance borne by the community at large. Rather, it is a psychic tax imposed on those least able to pay" (Matsuda 1989, 2323). The non-marginalized merely have to learn to tolerate falsehoods, many of which don't directly target them. The marginalized have to learn how to live with the fact that those contributing to their marginalization are getting polite and respectful treatment from their university's faculty and students. Similar remarks apply to Scriven's and Wood's contention that problematic speech allows us to develop critical reasoning skills and build convincing arguments against problematic speech. It's very nice for the non-marginalized that they get to bulk up their critical reasoning skills while playing the academic game of "Refute the Speaker." Meanwhile, the marginalized and victimized deal with the actual psychic costs of being surrounded by speakers who marginalize and victimize them, an establishment that provides a respectful platform for their speech, and students who respectfully listen to that speech in the name of open-mindedness.

Some who advocate inviting problematic speakers do not ignore this problem—do not ignore the special difficulties facing marginalized groups—but rather hold it up as a further reason to bring in problematic speakers. On their view, targets of problematic speech have especially valuable lessons to learn from that speech.[14] Nat Hentoff approvingly quotes future ACLU head Gwen Thomas as saying, "our young people have to learn to grow up on college campuses. We have to teach them how to deal with adversarial situations. They have to learn how to survive offensive speech they find wounding and hurtful" (Hentoff 1991, 549). David Bernstein adds that, "If Americans are going to preserve their civil liberties, then, they will need to develop thicker skin" (Bernstein 2003, 245). On this view, it's important to be exposed to problematic speech because it affords students the opportunity to toughen-up—to learn "resilience."[15]

---

[14] Manne and Stanley point out that

> the idea that oppressed and marginalized people should "grow up" has a long and ugly history. Women have frequently been dismissed using this stereotype, for instance. And the thing about children is that it is not always possible, nor even desirable, to reason with them. . . . Calling the student protesters "coddled" serves to excuse those touting freedom of speech as an ideal to spurn it in reality. (Manne and Stanley 2015)

This sort of response reflects the attitude that the reason why some students favor speech codes is that they find the expressed ideas too painful to hear and need protection from them: if the students need to toughen-up, then they are not already tough. When they interpret arguments in favor of speech codes, opponents of speech codes tend to take those arguments as demanding "protection" from "offense" as if the issue is that student ears are too delicate to hear those words. According to opponents of speech codes, advocates of speech codes want them in place for any of the following reasons:

- "to be protected from wounding language" (Hentoff 548)
- a supposed "right to be protected from offensive speech" (Lasson 1999, 44)
- "offense taken by the listener" (Bernstein 240)
- "some students, particularly those from ethnic minorities, might be so hurt by what they hear that they need to be protected from it" (Hayes 2009, 132)
- students need to be "shielded from words and ideas that make some uncomfortable" (Lukianoff and Haidt 2015, 44)
- to "create special protections for groups reductively defined as weak or vulnerable or to create buffers to spare sensitive young feelings from offense" (Paglia 2016).

> Responding to this kind of defense of speech codes, Hayes continues
>
> Is there not a whiff of class prejudice and racism in these arguments? Are the working class and ethnic minorities incapable of reasoning and need to be protected by middle-

---

[15] According to the burgeoning movement of "resilience science", "even serious disruptions in a child's life can offer unexpected opportunities for growth" (Brendtro and Larson 2004, 194).

> class student activist? . . . [T]he idea of a resilient student . . . has been entirely replaced by a more vulnerable picture of ethnic minorities and other groups. (2009, 132)

Greg Lukianoff and Jonathan Haidt join Hayes in arguing that the need for psychic protection stems from psychic fragility. The current movement to implement speech codes

> presumes an extraordinary fragility of the collegiate psyche, and therefore elevates the goal of protecting students from psychological harm. (Lukianoff and Haidt 2015, 44)

The progression of thought, then, leading to the conclusion that students should be exposed to problematic speakers in order to get them to toughen up has three stages:

A) students would only need to be protected from problematic speakers to avoid the psychological harm done by the content of "offensive" speech,

B) suffering psychological harm as the result of such speech is overly sensitive and to protect them from it would be "coddling," and therefore,

C) the proper response is to expose them to such speech in order to make them more "resilient."

8.3 THE CASE AGAINST INVITING (SOME) PROBLEMATIC SPEAKERS

8.3.1 Problematic speakers and psychological harm

This view of the matter gets everything wrong; against A, the problem with inviting problematic speakers is not solely or even primarily the psychological harm done to the targets of the speech by the speech itself. Against B, to the extent that there is psychological harm done to the targets of the speech, the targets of that speech are not being "offended" or "overly-sensitive" or require "coddling," and therefore (against C) the proper response is not to build resilience, since it wasn't resilience that was lacking in the first place.

The problems with B and C have their source in the problems with A, so the bulk of this chapter is concerned with exposing at least two ways A goes wrong. First, the speech itself and the ideas it expresses are not the only sources of psychological harm to the targets of the speech.

This is not to say that mere words can't cause psychological harms. Focusing primarily on slurs, Cassie Herbert (manuscript) argues persuasively that merely mentioning slurs can cause some of the same harms as actual uses of those slurs. Exposure to slurs increases feelings of dehumanization,[16] avoidance behavior,[17] and implicit bias[18] against members of the slurred group. Crucially, because in the relevant studies the slurs were merely mentioned—no one used them or directed them toward the participants in the studies, nor in some of the studies were any vignettes used in which the slurs were directed at anyone—it is arguable that the words themselves were the cause of the trouble.

Herbert notes that one explanation for these results is that exposure to slurs primes implicit and pernicious associations between group membership and other, negative stereotypical features. Because these associations can be implicit in a participant's psychology even while explicitly denied, the unfortunate effects of exposure to slurs will arguably be seen in anyone who has internalized the implicit associations. This includes members of the ingroup themselves. While the studies emphasize the effects of slurs on outgroup members—participants who are not members of the slurred group—Herbert notes that

> it's been well documented that ingroup members also internalize pernicious associations about the group of which they are a part (and sometimes do so even more strongly due to repeated exposure to the stereotype), and there's little reason to think that this doesn't hold with slurs—though the details of the associations called forth ought to be studied. Moreover, ingroup members are impacted by the slur via the stereotype threat it calls forth.

This means that, not only will the likely effects of exposure to slurs be seen in ingroup as well as outgroup members, the effects themselves need not be limited to avoidance, dehumanization, and implicit bias. They can include any harms that can be brought about by priming of pernicious implicit associations. In ingroup members, these harms include stereotype threat and, though Herbert does not explicitly discuss these possibilities, painful emotional consequences, feelings of inadequacy, and alienation from the group responsible for the appearance of the slurs. Such might be the explanation for the students' reactions in Herbert's motivating case—that of

---

[16] See, e.g., (Fasoli, Paladino, et al. 2016).
[17] See, e.g., (Carnaghi and Maass 2008).
[18] See, e.g., (Nicolas and Skinner 2012).

> a professor at Mount Holyoke who deployed a variety of slurs in his classroom in order to make a point about Robinson Crusoe. Going around the classroom, the professor reportedly identified slurs that could be directed against specific students on the basis of the students' race or ethnicity. Mentioning these contemporary slurs, he believed, would provide a valuable learning experience to better understand a historical slur used in the book. In some cases the professor required the students themselves to name the slur that could be directed against themselves. Now, despite the slurs being merely mentioned, students reported feeling "shocked" and "frozen." As a direct response to this experience, seven students transferred out of the course.

This case is less about implicit bias and stereotype threat and more about the kinds of emotional harms and feelings of alienation that are so often expressed as a response to offensive speech. If Herbert is right that merely mentioning slurs can prime implicit pernicious associations in ingroup members, then it is likely that the result will not just be the harms of stereotype threat, but more emotional harms as well.[19]

Herbert focuses on slurs, and the studies she emphasizes are studies on the effects of exposure to slurs. They don't deal with the effects of other kinds of speech: hate speech, speech that presents arguments for views demeaning to members of marginalized groups, or speech presented by a problematic speaker. Therefore, we need more evidence before we can conclude that the harm of all these kinds of speech lies in the speech itself—whether it is mentioned or used. Again, though, with respect to at least some of these kinds of speech, if what causes the relevant harms is the priming of implicit associations, it is natural to think that similar harms will result from exposure to the speech itself.

---

[19] One question left unanswered is what the source of the harm of mere mentionings is. *Why* do mere mentionings cause emotional harms? Why do they prime implicit associations? One explanation is that the mere mentioning of the relevant words brings certain people's thoughts to pernicious associations. If so, it would happen just as well if the words were encountered in a randomly strewn about collection of driftwood. A second explanation is that context matters. Perhaps it's essential that the problematic speech be put forth by a person: an experimenter or a professor. This might be why the students in the Holyoke professor's class had such a powerful reaction; if a person is so negligent with words of such power, it might seem, their attitudes must be problematic. Even in the attenuated studies emphasized by Herbert, in which isolated occurrences of slurs are shown on a screen, out of the context of a sentence, there is a person or group of people—the experimenters—that can be impugned by the appearance of the term.

If the second explanation is correct, then there is a quick argument to the conclusion I draw—that it's not only the speech itself that causes harm, but the fact that the speech receives a kind of endorsement that alienates the target of the speech. I leave open which explanation is correct. Still, even if the first explanation is correct and it's the speech itself that causes harm, it's implausible, for reasons invoked in the body of the paper, that the speech itself exhausts the causes of emotional harm.

I grant Herbert that mere exposure to problematic speech can be psychologically harmful in the ways she suggests (and more). Speech itself can be damaging. But it is not plausible—and nor need or does Herbert disagree—that the speech itself exhausts the source of psychological harm. For it is not plausible that the harm done by the speech itself is as damaging as the harm done by that same speech in other contexts. Exactly similar speech, shown on video as an exhibit of racist or sexist speech, or described in a fictional vignette won't always plausibly do the same sort of psychological harm as it would do when actually used. This is plain in various studies that have been performed to show that judgments of the harm resulting from hate speech vary from group to group.[20] To show this, experimenters present to subjects various scenarios in which hate speech is directed at a person or group within the scenario. Subjects are asked to judge how much the targeted person in the scenario is harmed. It is an implicit assumption of the experiment that the subjects of the experiment are not harmed by the presentation of the scenarios anywhere near the degree that it's possible that the fictional targets of the hate speech would be harmed. (If the subjects were harmed to that extent, the experiments would not get ethics approval.)

---

[20] See, for example, (Cowan and Hodge 1996).

So, while it’s true that being confronted—even as an exhibit—with certain kinds of racist or sexist speech can trigger some powerful psychological reactions if there is no warning given, once the warning is given, this kind of speech—while powerful—doesn’t dole out the lasting harms racist, sexist, and otherwise marginalizing speech can exert. On the contrary, there is evidence that with suitable warnings, exposure to the elements of past traumatic events can be helpful in relieving the effects of PTSD.[21]

It is partly for this reason that the recent backlash against trigger warnings and safe spaces is misguided. The University of Chicago, for example, in a 2016 letter to incoming students, wrote,

> Our commitment to academic freedom means that we do not support so-called trigger warnings, we do not cancel invited speakers because their topics might prove controversial and we do not condone the creation of intellectual safe spaces where individuals can retreat from ideas and perspectives at odds with their own. (2016)

Trigger warnings, of course, are not designed to allow students to avoid contact with views or ideas that they disagree with or make them uncomfortable. They are designed to give students a head’s up so that when they are confronted with such ideas they are able to respond to them in a productive way. While opponents of trigger warnings may be right that it can sometimes be psychologically helpful to confront those issues that are related to past trauma, evidence shows that it is only helpful when that confrontation is done is a safe space, where the victim of trauma can be relaxed and prepared for the difficult confrontation. Trigger warnings and safe spaces allow students to confront difficult material in a way that allows them to benefit the most from that material.[22] As Annaleigh Curtis (2016) says, in comparing content that triggers victims of PTSD to images that trigger her “weird phobia” of grasshoppers,

---

[21] See, for example, (Resick, Nishith, et al. 2002) and (Foa, Rothbaum, et al. 1991) on the merits of so-called “Prolonged Exposure Therapy.” It is essential to the therapy that the exposure be “controlled” in the sense that, first, the therapy “starts with education about the treatment,” second, that you learn to control your breathing “to manage immediate distress,” and third, that “you practice approaching situations that are safe but which you may have been avoiding because they are related to the trauma” (U.S. Department of Veterans Affairs 2015). I assume furthermore that the exposure needs to be *voluntary*—actively chosen. The implementation of all of these steps require that the subject have some warning about impending, potentially triggering experiences, as well as control over whether they’re exposed to the relevant stimuli.

> about the worst thing you could do to someone with a phobia is just spring triggers on them with no warning. I’m guessing it is similar for people with trauma. It will absolutely not help a survivor of rape to force them with no warning to read a narrative about rape or listen to a bumbling class discussion that reinforces rape myths. On the other hand, it might help a lot to know that there will be a thoughtful and respectful discussion of the topic so they can be ready.

Racist and sexist speech isn’t harmful merely because it exposes some listeners to the content of that speech. The context in which the speech is presented can magnify or reduce the harm.[23] The severity of psychological harms caused by bullying, hateful, or aggressive speech depends on whether it happens in front of others, whether authorities either are complicit or take a hands-off or tolerant policy, and how peers react to the speech.[24] The larger social situation matters.

Part of the larger social situation that makes the racist and sexist attitudes of invited speakers so damaging is *that they’ve been invited* and, therefore, as the norms of invitations require, that they are treated respectfully by the host: the university community. This explains why the most harmful consequences of abusive speech in a university setting include debilitating feelings of isolation and estrangement. Charles Lawrence approvingly cites the UT Austin’s Report of the President’s Ad Hoc Committee on Racial Harassment, which concluded that many students who feel that they experience racism on campus end up with “deep feelings of personal anger, distress and isolation in the academic community. Such experiences produce alienation from, rather than a supportive relationship with, the University" (Lawrence 1990, 465 n. 126). Delgado and Stefancic echo the sentiment in emphasizing that hate speech routinely results in "fear, nightmares, and withdrawal from society—what Joe Feagin and his collaborator call an impotent despair" (2004, 14).

---

[22] Similar considerations arise regarding confrontation with the arguments of problematic speakers. The argument in this chapter is not designed to show that certain views should be excluded from campus. It is designed to show only that there are different ways of presenting those views, and that inviting problematic speakers and showing them the respect required of a visiting speaker is morally problematic. There are better ways to present problematic and difficult views, that don’t require showing public solidarity with the wrong side.

[23] Nor is racist or sexist speech harmful simply because that speech makes its targets aware of racism and sexism. Members of marginalized groups are already quite aware of the existence of those attitudes; they are, after all, subject to them on a daily basis (see (Rankin and Reason 2005) and (Smith 1992)).

[24] See, for example, (Jay 2009, 88), (Geiger and Fischer 2006, 348 and 350–351), and (Delgado 1982, 143).

These feelings can only be brought about by a problematic speaker if the speaker is somehow associated with the group that the target of the speech feels isolated and estranged from. If you feel that the group will not support you, or will be ashamed of you, or embarrassed by you—if you feel that the group shares the speaker's attitudes or even just thinks of them as a live hypothesis—then it is very difficult to feel like you are a member of the group. In short, the psychological harm caused by the problematic speaker is not just "offense" at the speaker's words: it's a feeling of betrayal and isolation that is the result of a perception that the speaker is getting more respect from the larger group than are the targets of the speaker's words.[25] As Matsuda notes,

> Official tolerance of racist speech [in the university setting] is more harmful than generalized tolerance in the community-at-large. . . . It is harmful to targets, who perceive the university as taking sides through inaction, and who are left to their own resources in coping with the damage wrought. (1989, 2371)

This is because, generally, in the words of Uma Narayan,

> The disadvantaged cannot fail to realize that being hurt by the insensitivity of members of the advantaged groups they endeavor to work with and care about, is often more difficult to deal with emotionally than being hurt by the deliberate malice of members of advantaged groups they expect no better of. (1988, 35)

---

[25] Jeremy Waldron (2012), too, contrasts the feeling of offense at a speaker's words with what he regards as the primary harm of hate speech—that it undermines its target's dignity. He grants that targets will naturally find attacks on their dignity as "wounding and distressing" (108) but stops short of equating the powerful psychological harms that result from attacks on one's dignity with feelings of "offense" which is what one feels when one's "feelings are hurt." (108)

If Matsuda and Narayan are right, the harm done by universities that invite problematic speakers is akin to what Jennifer Freyd calls “betrayal trauma” (1998).[26] Betrayal trauma is a kind of trauma that “occurs when the people or institutions on which a person depends for survival significantly violate that person’s trust or well-being” (Freyd 2008, 76). Victims of betrayal trauma fare worse than victims of other kinds of trauma on several instruments (Freyd, Klest, et al. 2005, 97–98) measuring such symptoms as “headaches, chest pain, and abdominal pain” (Freyd, Klest, et al. 2005, 90) and anxiety, depression, and sleep disturbances. (Freyd, Klest, et al. 2005, 91) This is in keeping with other findings about the consequences of betrayal—that betrayal induces “visceral, intense, and protracted negative feelings,” (Koehler and Gershoff 2003, 245) and that that those who are betrayed “commonly experience an ‘unrelenting anger’ and view the world as a disordered place” (Koehler and Gershoff 2003, 245).

Summing up: one thing that advocates of unfettered invitations to problematic speakers get wrong is the claim that, if there is a problem with inviting such speakers, the problem lies in the psychological harms done to targets of the speech by the speech itself. On the contrary, the psychological harm is done not only by the speech itself, but by the fact that the speech is rendered in a certain context—by a speaker who is invited to the university and given respectful attention by their hosts and audience. In this context, the harmful psychological consequences are not simply offense at the speaker’s words, but feelings of isolation and betrayal directed not primarily at the speaker but at the university community.

8.3.2 Problematic speakers and intrinsic harm

---

[26] Freyd also claims that victims of betrayal trauma have a tendency to forget and even be unable to recall without assistance the initial trauma—what Freyd calls “betrayal blindness” (Freyd 1998, 1–11). I take no stand on this view and see (McNally 2007) for evidence to the contrary.

Of course, feelings of isolation and betrayal can be irrational or misguided; they can fail to reflect genuine isolation and betrayal. But in the relevant cases, these feelings are not irrational or misguided. And that's because of the second difficulty with what I called step A—the assumption that students would only need to be protected from problematic speakers to avoid the psychological harm done by the content of "offensive" speech. The second difficulty with this stage is that the problem with the invitation of problematic speakers is not just the forward-looking, consequentialist psychological harm that the invitations to such speakers visit on their targets.

To be sure, these are problems: as Stanley Fish says, commenting on ex Yale President Benno Schmidt's opposition to campus speech codes,

> The trick is nicely illustrated when Schmidt urges protections for speech "no matter how obnoxious in content." "Obnoxious" at once acknowledges the reality of speech-related harms and trivializes them by suggesting that they are *surface* injuries that any large-minded ("liberated and humane") person should be able to bear. The possibility that speech-related injuries may be grievous and *deeply* wounding is carefully kept out of sight, and because it is kept out of sight, the fiction of a world of weightless verbal exchange can be maintained, at least within the confines of Schmidt's carefully denatured discourse. (Fish 1994, 109)

The psychological harms are real. But the psychological harms don't exhaust the ethical problems with inviting problematic speakers. As we've seen, part of what causes psychological harm is the knowledge that the problematic speaker was chosen and tolerated by administrators of and peers in the group one hopes to identify with. It is not just arbitrary that this is particularly psychological painful. It's painful because the feeling of betrayal reflects an actual betrayal. Students who feel isolated and betrayed because of the invitations to problematic speakers are often right to feel isolated and betrayed because, in inviting the problematic speakers they have been harmed intrinsically, and not just in virtue of the psychological harms they have suffered. Or so I argue.

I start with a mundane observation: sometimes, how it's appropriate to behave in a social situation depends on who's watching or who's involved in the situation. If the chef is watching, it isn't appropriate to rudely disparage the awful food even if it might be appropriate when the chef isn't watching. How you'd discuss the progress of a student in your program can depend on whether the student is present. But sometimes, how you should behave when a person isn't there depends on how you should behave if the person is there. You should not flirt with strangers in ways you wouldn't be comfortable talking with them if your monogamous romantic partner were watching. If your friend were horribly and unjustly mistreated by someone, you shouldn't be respectfully deferential to that person simply because your friend doesn't happen to be around.

The wrongness of behaving these ways isn't due solely to the potentially unfortunate consequences if the news of your behavior gets back to the relevant person. True, if your friend or romantic partner hears about the behavior, they might feel betrayed. But the point is that they would *rightly* feel betrayed. If your friend is dead—so there are no potential emotional consequences to your friend—you still should not be offering respectful deference to their unrepentant murderer.

In the above cases, the obligations to behave only in ways you would behave were the relevant third party observing stem from the special obligations incurred by being involved in special social relationships: monogamous romantic relationships or personal friendships. But sometimes your obligations to behave only in ways you would behave if the third party were observing stem not from any special relationship, but from distinctive features of the third party. If your child says disparaging things about a stranger in the supermarket, you'll tell your child that such comments are inappropriate. Your child's excuse that they weren't overheard is irrelevant: if they’re not comfortable saying it in front of the person, they shouldn’t be saying it at all.

The *explanation* for why the behavior is wrong is not, in these cases, that you can’t permissibly manifest the behavior in front of the person. It’s simply that such behavior is impermissible. You shouldn’t be saying demeaning things about strangers. If you are in personal friendships of a certain sort, then you shouldn’t be respectfully deferential to their abusers. If you are in a monogamous romantic relationship of a certain sort, then you shouldn’t be flirting with people on the bus. But the imaginative heuristic of asking how you’d be comfortable behaving in front of the person can still provide helpful information. It tells you what kind of behavior counts as problematic. It tells you what behaviors count as inappropriate flirting, what interactions constitute problematically respectful deference, and what comments about a stranger count as demeaning.

It’s clear that there are ways of behaving toward problematic invited speakers that would be inappropriate were their targets involved in the social situation. Consider, first, a small discussion group consisting of you and a few students. If you invite a speaker to speak to your group, your duty as a host requires giving certain kinds of respect: mentioning their qualifications, presenting them as an authority, respectfully introducing them, etc. With respect to some settings and some speakers, none of these behaviors are appropriate. If your group includes a known Holocaust survivor, it is impermissible to behave in these ways toward an unrepentant Holocaust denier or a Nazi if that speaker is going to argue for the virtues of Naziism or the truth of Holocaust denial. If your group consists exclusively of white students, but for a single student of African descent, it is likewise impermissible to behave in these ways toward a Ku Klux Klan member who you’ve brought in defend the historical practice of lynching: “I’d like us all to give a warm welcome to our guest speaker for today, who will be presenting a challenging argument for a controversial conclusion. Let’s hear what he’s got to say.”

The problem also arises if the speaker isn’t going to talk directly on issues that target members of the group. You should also not behave respectfully toward speakers who have in the past acted in certain ways. If the speaker had, in the past, assaulted one member of the group, it is inappropriate to give them the kind of respect required of you as an inviting host. You should not treat the concentration camp guard formerly assigned to your Holocaust survivor with the deference due to an invited guest. This is so even if the assaulter or the concentration camp guard is going to speak on a completely unrelated subject. Some hands you just shouldn’t shake.

If you shouldn’t offer this respectfully deferential treatment to certain invited speakers, then you shouldn’t invite those speakers to speak to your group. In inviting them, you have an obligation to treat them in a certain way—with exactly the kind of respectful deference that you are not permitted to display when the targets of their speech and actions are around. If inviting them requires some action that you are not permitted to perform, then you are not permitted to invite them.

The respect-behavior and, thus, the invitation, don’t become permissible if the Holocaust survivor, the African-American student, or the assault victim happens to be absent that day or has coincidentally dropped out of the group the day before. You don’t get to give respectful deference to the speaker merely because the person who made the deference inappropriate isn’t there. The situation in which you have invited these problematic speakers is more like inappropriate flirting when your romantic partner isn’t present or giving respectful deference to your friend’s unrepentant murderer. If it’s not appropriate to give respectful deference to the

Holocaust denier when the Holocaust survivor is around, it's not appropriate to give respectful deference to the Holocaust denier when the survivor isn't around.

Finally, it's no response that inviting such speakers gives students the opportunity to develop resilience or tolerance or the critical reasoning skills necessary for responding even to really abhorrent views. There are limits to such defenses, even in the quest for resilience, greater tolerance, and heightened critical reasoning ability and this is one place where the limits have been reached. When the Holocaust survivor, appalled, asks you why you felt it necessary to invite their old concentration camp guard to lecture on the virtues of Holocaust denial, it is no defense that you thought it would be a perfect "teachable moment": that it was a wonderful opportunity to really instill in the class—and the survivor—the virtue of resilience or tolerance or the nuances of critically appraising false views that strike really close to home.[27]

Like inviting a speaker into your discussion group, inviting a speaker to give a public talk on campus comes with certain duties to the speaker: when you invite a speaker, you have a duty to treat them respectfully. When a speaker is invited as a joke, when after they are invited they are mocked or set up in a way that makes them obviously a fool, this violates an obligation to them in such a way that makes it clear they shouldn't have been invited at all.[28] In addition to the respect conferred by the invitation itself, the respectful treatment you are obligated to provide involves 1) a decent forum, 2) if not a guarantee, at least a likelihood of a decent audience, with

---

[27] The situation might be different if the students have knowingly chosen to sign up for a group whose purpose is to visit such surprises on its members. If a student wants the opportunity to debate their old unrepentant concentration camp guard, that may be their business. The considerations here are similar to those bearing on the issue of the ethics of members of marginalized groups inviting speakers who marginalize those very groups; see below for discussion.

[28] As commenter Jonathan says in a 2011 blog discussion in which a number of philosophers were mocking Carlin Romano's invited presentation at a conference on progress in philosophy,

> A lot of folks seem to be "shocked, shocked" to find that Carlin Romano was a shrill, anti-intellectual moron. But wasn't he invited for that reason? He has a long paper trail of ignorant hostility to academic philosophy. He showed up and said what everyone expected him to say and acted the way everyone expected him to act. One might even say he acted the way he was *supposed* to act. (Jonathan 2011)

3) advertising that doesn’t demean the speaker, but 4) makes the speaker out to be at least an authority and an expert, along with 5) an introduction and a setting that makes the speaker out the same way.

This sort of respectful and deferential behavior wouldn’t be appropriate in a social setting limited to you, the speaker, and a person who has or will be wrongly marginalized or victimized by the speaker or their speech. Of course, when you invite a problematic speaker to a university, often they haven’t assaulted any members of the potential audience, nor will their words single out individual members of the potential audience. But, as we’ve seen, in many cases these differences don’t make a difference to the ethical issues. If their words are not targeting individual members of the audience, but are rather targeting already marginalized groups that aren’t necessarily part of the audience, it doesn’t always eliminate the ethical worries. The absence from your discussion group of the target of the Holocaust denier’s speech doesn’t allow you to give respectful deference to the denier. The fact that the speaker isn’t singling out any member of the audience for ridicule is similarly irrelevant to whether respectful deference to the speaker is permissible. If they haven’t assaulted a member of your campus, but they have victimized others, it is often still not morally permissible to overlook these acts and invite them nonetheless, for the same reason it is not morally permissible to invite the concentration camp guard to speak to the discussion group just because the Holocaust survivor will be absent.

In all these cases, the facts that make invitation impermissible are known to you: you know that Holocaust denialism is false and that lynching is monstrous; you know that the Nazi was the survivor’s concentration camp guard and that the speaker assaulted your student. You know furthermore that the Nazi is unrepentant and that the speaker hasn’t made amends for assaulting your student. In short, in all these cases, you know that the speech or action was either false or wrong, respectively, and that the false speech or wrong action is indicative of the speaker’s current attitudes. All of this knowledge seems relevant to the fact that the absence of the speaker’s target from the social situation makes no difference to whether it is permissible to offer respectful deference to the speaker. It is impermissible to offer the respectful deference required of a host because you know 1) that their victimizing or marginalizing speech or behavior is false or wrong, respectively, and 2) that their behavior represents their current attitudes.

When you know that their victimizing or marginalizing speech or behavior isn’t wrong or false or that it doesn’t represents their current attitudes—or even fail to know whether any of this is true—then issues are trickier. For example, if a speaker is denigrating, not Jews or African-Americans, but anti-vaxxers, then it still might be impermissible to offer respectful deference to the speaker in a room with a committed anti-vaxxer. If the speaker, for example, is claiming that anti-vaxxers are essentially irrational, given the evidence widely available about the efficacy of vaccination, it might be impermissible to give respectful deference to that speaker while there is a sincere and well-meaning anti-vaxxer listening in; it’s generally rude, after all, to listen respectfully to someone saying insulting things about another well-meaning person present, even if the insulting things are true. But this doesn’t mean that it is impermissible to give respectful deference to that speaker in the larger context of an invited talk. That’s because you don’t know that the speaker’s words are false.

There are other tricky cases in which you don’t know whether the speaker has committed acts or spoken falsehoods that they are still accountable for. For example, if you have some evidence that a potential speaker has committed assault, but know neither that they did nor that they didn’t, then social situations involving both the speaker and the alleged victim would certainly be uncomfortable. Perhaps it is impermissible to behave respectfully toward the speaker in such settings. But I take no stand on the ethics of inviting speakers in situations of ignorance. There may be no substitute for expected utility calculations: the potential psychological and intrinsic harms to the alleged victim, the potential harms to a speaker wrongly accused of assault, the probability that the accusations are true, etc. But even expected utility calculations may not be up to the delicate task of negotiating these situations.

If you know that a speaker's prior false speech or wrong actions have marginalized or victimized some group or individual, then whether it is appropriate now to offer the person respectful deference can depend on the degree to which you know that their prior behavior represents their current attitudes. If the speaker has in the past said disparaging things about some marginalized group but has since repented—if they've sincerely recanted and sufficiently apologized for and rectified their past speech—respectful deference may be permissible. Such recanting and apology may not be sufficient for the speaker to have earned the right to be treated fully respectfully in an intimate social setting that includes members of the previously disparaged group. But it may be permissible to treat the person in the respectful ways required by an invitation to speak, though perhaps only with the endorsement of those previously marginalized by their speech. You might be able to permissibly invite and treat with due respect a reformed Klan member who now devotes their life to civil rights advocacy.

Similar remarks apply to the speaker who has committed assault, but has made it right, sincerely repented, sufficiently apologized, and made amends. In such cases, it still might not be appropriate to treat the speaker and their victim equally respectfully in a social situation involving the two of them. But that doesn't mean it's impermissible to be respectful to the speaker when the victim is not around. It can perhaps be permissible in some cases to invite such speakers, though, again, perhaps only with the endorsement of those they previously victimized; their previous actions do not represent their current attitudes.

But in other cases of previously victimizing speech or behavior, treating invited speakers with respectful deference is impermissible. If you know that their prior victimizing speech or behavior represents their current attitudes, then the reason why it would be wrong to treat them with equal respect in a social situation with both speaker and victim is that doing so in an important sense betrays the victim. If you know that a speaker has wrongly victimized or will victimize an individual or group—either by act or speech—in such a way that it reflects on their current attitudes, you should not offer them the kind of respectful deference a host is obligated to offer. Because you should not be offering the kind of respect you would be obligated to offer if you invited the speaker, you should not be inviting the speaker. This is the Rule of Non-Invitation:

> (Rule of Non-Invitation) You shouldn't invite a speaker if i) you know that the speaker has spoken or will speak falsehoods or has committed or will commit impermissible acts that accurately represent the speaker's current attitudes or moral status and ii) there is a victim of those acts or falsehoods such that in social situations consisting of you, the speaker, and the victim, it would be impermissible to behave toward the speaker in the way you'd be required to behave toward a visiting speaker.

The Rule of Non-Invitation permits inviting speakers whom you know are going to present misleading arguments—like an argument that Oswald was not involved in the assassination of JFK. Here there is no question of standing in solidarity with the speaker's victim. Nor need it be impermissible to invite a speaker expressing merely offensive speech that doesn't marginalize anyone, and who is such that being respectful to them doesn't require you not to stand in solidarity with any member of a marginalized group. If the speaker is simply generically insulting, or advocating an obviously false offensive position (like, "all people are jerks") inviting them need not satisfy the conditions. Nor need the invitation be impermissible if the speech is merely known to be potentially dangerous. Giving respect to a global-warming denier does not require you to fail to stand in solidarity with anyone marginalized by the denier's speech. There might be other principles in the vicinity of the Rule of Non-Invitation that preclude inviting the global-warming denier; if it is impermissible simply to give respectful deference to such a person, then it's impermissible on the public stage as well. But the Rule of Non-Invitation does not issue a verdict on that.

The Rule of Non-Invitation becomes complicated in cases in which the person or group inviting or hosting the problematic speaker is also marginalized or victimized by the speaker. This can happen any time there is a public debate between a problematic speaker and a member of a group marginalized by the speaker's speech, and organized by the member of the marginalized group. For example, if a gay advocate for same-sex marriage organizes a public debate in which they are up against a steadfast proponent of conversion therapy, it might seem unclear whether the Rule of Non-Invitation allows the invitation. Likewise if a black civil rights advocate organizes a respectful debate with the leader of the local chapter of the Ku Klux Klan. But what might seem clear is that such invitations are appropriate: members of marginalized groups can permissibly organize debates with speakers whose speech marginalizes them. If so, then if the Rule of Non-Invitation gives the contrary result, this is a counterexample to the Rule of Non-Invitation.[29]

---

[29] A particularly tricky case is the invitation extended by the historically black Dillard University to David Duke, the former Imperial Wizard of the Ku Klux Klan. Dillard had previously committed to host a debate between any candidates for Louisiana's Senate seat polling at least 5%. Duke unexpectedly met that threshold (at 5.1%). According to Dillard President Walter Kimbrough, "If we're trying to get out of it because one person is coming to campus, that's a problem for me in terms of what I value. . . . That's one of the criticisms of higher ed: we don't accept diverse opinions" (Seltzer 2016). A letter from a group of Dillard students disagreed:

> Dr. Kimbrough, respectfully, this response is specious. You are the president of a historically black college whose mere presence is anathema to EVERYTHING David Duke promotes. Instead of denying the presence of this terrorist onto our campus, you have ASSURED HIS SAFETY by Dillard University armed police, AGAINST US, your Dillard University student body. We write to you today not only to express our

The Rule of Non-Invitation doesn't give the contrary result—at least, not in any case in which the contrary result is incorrect. The issue hinges on whether it's permissible for members of marginalized groups to give respect to marginalizing speakers in social situations including the speaker and the member of marginalized group—namely, them. I think that the answer to this question is that it is permissible. For the most part, members of marginalized groups have authority over which of their marginalizers they are respectful toward and in what contexts. Therefore, the Rule of Non-Invitation allows minority invitations to problematic speakers.

But this is also true: there may be some cases in which members of marginalized groups cannot permissibly give respect to speakers that marginalize them, even in contexts in which they want to host a public debate. This can be so, for example, if in doing so, the inviter would end up speaking for the rest of the community in a problematic way. In this kind of case, the Rule of Non-Invitation would have the result that the invitation to the problematic speaker is inappropriate. But, in exactly this kind of case, that seems like the right result. Once it's inappropriate for the inviter to privately give respect and deference to the problematic speaker—whether the inviter is a member of the marginalized group or not—it also seems inappropriate for that respect and deference to be given on the public stage, as long as the first condition of the Rule of Non-Invitation is satisfied. Thus, when the U.S. Holocaust Museum invited Donald Trump to keynote their Days of Remembrance ceremony, an open letter signed by over ten thousand U.S. Jews and titled "American Jews to the U.S. Holocaust Museum: WTF?"[30] opened by saying,

---

> hurt and shame but also to fight for our ancestors and their struggles. How dare this administration stand for Duke's "safety" and not fight for our security and right to learn in a healthy space. (admin 2016)

Without the existence of the prior commitment, my position is that the students would be in the right. I take no stand on whether the prior commitment to host the debate generates a strong enough overriding obligation (though my sympathies are with the students there, as well).

[30] See http://www.weveseenthisbefore.org/not-this-president.

> President Trump’s administration has repeatedly insulted the memory of the Holocaust, and embraced the agenda and rhetoric of white nationalism and antisemitism. So how can the U.S. Holocaust Museum invite him to deliver the keynote remarks at the National Day of Remembrance?

Therefore, far from presenting a problem for the Rule of Non-Invitation, cases in which members of marginalized groups invite the relevant problematic speakers tend to confirm the Rule.

What of intra-campus groups who invite problematic speakers whose speech is consistent with the message of those groups? While some on campus might know that the speech of a white supremacist invitee will be false and marginalizing to a group on campus, if a group on campus is sympathetic to that speech, that group won’t satisfy the condition of the Rule of Non-Invitation. Is it, then, permissible for white supremacist groups to invite white supremacists to campus?

As far as the Rule of Non-Invitation goes, the answer may be yes. There may be other ethical problems with inviting such speakers that apply even to those campus groups that agree with them. But that doesn’t mean that the Rule of Non-Invitation has no application to these kinds of cases, for two reasons. First, while the group may fail to know that the white supremacist’s speech is false, they may be in a position to know that the speech is false. Their lack of knowledge may be solely due to an irrational lack of belief. If the Rule of Non-Invitation is true, an analogous rule is arguably true that requires only being in a position to know that a speaker’s words or acts are false or wrong, respectively.

Second, many who invite problematic speakers hope to make irrelevant their views on the content of the speaker’s speech or the worth of the speaker’s past behavior. The justification is supposed to be about increasing diversity of expressed viewpoints on campus and providing an opportunity to strengthen counterarguments, even if those inviting the speaker “personally disagree with [their] arguments” (Wood 2015). As University of Manchester’s Free Speech and

Secular Society put it, in protesting Manchester's decision to ban Milo Yiannopoulos and Julie Bindel from speaking, "We do not endorse the views of either of the speakers but merely wish to hear them speak and challenge them on what they say."[31] The Rule of Non-Invitation eliminates that weasel-room. If the argument for the rule is sound, to defend your invitation of the problematic speaker, you must admit to not knowing that their words are false and demeaning or their actions are harmful and not made right.

The Rule of Non-Invitation explains why one common response to arguments for disinviting problematic speakers is not to point. Some philosophers, in response to a consequentialist worry that inviting speakers with dangerous or false views might propagate those falsehoods or incite people to dangerous actions, quote Justice Brandeis' famed remark in Whitney v. California: "If there be time to expose through discussion the falsehood and fallacies, to avert the evil by the processes of education, the remedy to be applied is more speech, not enforced silence." So, it is suggested, we can eliminate the problem of inviting problematic speakers by following the speaker up with an alternative speaker on the other side.[32] But if the difficulty with inviting problematic speakers is not simply the spread of falsehood, but avoiding the intrinsic harm done to victims of the speaker by granting them the respect and deference due an invited guest, then allowing a rebuttal is insufficient. Respectful deference to the Klan member or the Holocaust survivor's concentration camp guard is not made right by allowing a historian to correct their errors after the fact.

---

[31] See http://www.free-speech-manchester.co.uk/blog/statement-on-the-ban-of-milo-yiannopoulos-and-julie-bindel/.

[32] See, e.g., (Lasson 1999, fn. 24). There are other responses to this possibility as well. Fish notes, in response to the possibility that "harmful speech should be answered not by regulation but by more speech," that this "would make sense only if the effects of speech could be canceled out by additional speech, only if the pain and humiliation caused by racial or religious epithets could be ameliorated by saying something like, 'So's your old man'" (Fish 1994, 109).

According to the argument for inviting problematic speakers, students only need to be protected from problematic speakers to avoid the psychological harm done by exposure to "offensive" speech. This, I've argued, is false: the problem with inviting problematic speakers is not solely or even primarily the psychological harm done to the targets of the speech by the speech itself. This is because, first, it's not the speech itself that is the primary cause of emotional and psychological harm; it's the fact that the speech was invited and presented in a respectful context in a way that made feelings of betrayal and isolation the natural result. Second, the psychological and emotional harm does not exhaust the wrongs done by the invitation: there is also the intrinsic harm done by the invitation itself and the inevitable respect and deference that hosting such a speaker requires.

8.3.3 Problematic speakers and the respect for truth

If the invitation to certain problematic speakers visits an intrinsic harm, then we can also see what's wrong with the second step in the argument for inviting problematic speakers—what I called step B, above, according to which suffering psychological harm as the result of problematic speech is overly sensitive and to protect students from it would be coddling. On the contrary, it's not over-sensitivity students exhibit when they are caused psychological harm by the invitation of problematic speakers. It would be over-sensitivity if what was causing harm was something not deserving of pained reaction. But in this case, it is eminently rational to be harmed—caused psychological pain—by the actions of the inviter.[33] For it is rational for students to feel betrayal and pain when their peers wrongly take stands with outsiders who glorify their intimidation and marginalization through words. Students are rational to feel like they have been left on the outside when the group they identify with has shown respect and deference to their victimizer.

---

[33] Again relevant, here, is Waldron's contention that more important than the psychological harms of inviting problematic speakers is the intrinsic harms such speakers do to their targets by attacking their dignity. If Waldron is right, then that's a further reason why the psychological harm caused is not simply a result of over-sensitivity; rather, targets of the speech are responding correctly to a grievous intrinsic harm. See (Waldron 2012, esp. ch. 5).

In short, students are rational to feel betrayed *when they are betrayed*. If, as argued here, inviting problematic speakers is to wrongly take a stand with outsiders who glorify the intimidation and marginalization of your ostensible colleagues, then it is perfectly rational for those colleagues to feel debilitating betrayal and isolation. (I don’t mean to say it’s ever irrational to feel betrayal and isolation. But in the case where feelings of betrayal and isolation are in response to actual wrongful betrayal, it is most obvious that the feelings of betrayal and isolation are not the result of over-sensitivity and a desire for “coddling.”) On the contrary, someone who did not feel betrayal and isolation in this case might be considered a doormat.

To emphasize the point: it is recognized on all sides that the targets of problematic speakers suffer feelings of betrayal as a result of the invitation of invited speakers. The question is whether those feelings constitute over-sensitivity, or whether they are rational and justified. I say that they are rational if those feelings are a response to an actual wronging—something like an actual betrayal. There is nothing oversensitive or irrational about having feelings of betrayal in response to something relevantly like an actual betrayal. So if, as I have argued here, the wrong done to targets of problematic speakers is often an actual betrayal, then their feelings of betrayal are not overly sensitive and irrational. And if they’re not overly sensitive and irrational, then to take steps to prevent them from suffering the relevant harms is not “coddling.”

Lastly, this conclusion tells us why we should reject the final step in the argument for inviting problematic speakers. According to that final step, the proper response to the suffering visited on the targets of problematic speakers is to expose them to such speech in order to make them more "resilient."[34] But, because the psychological harm often caused by inviting problematic speakers is perfectly appropriate, rational, expected, and even required, the proper response is not to cultivate resilience. Because the reactions are reactions to wrongful betrayal, we should not be trying to get students used to that betrayal so that. . . what? We can permissibly continue the practice? The right response to students who are rightly pained by our wrongful choices is to stop making those choices. While we do want to train students to handle problems thrown their way, we don't want to train students to think that actual sins against them are not problematic. We want students to be accurate monitors of when they are being mistreated and when they are being treated fairly. If training students to be "resilient" has the effect that they become deadened to the fact that they are being wronged, this is not the result educators should want.

This exposes a tension in recent defenses of "viewpoint diversity," defined by the "Heterodox Academy"[35] as "what happens when everyone in a field shares the same political orientation and certain ideas become orthodoxy."[36] On the one hand, according to members of the Heterodox Academy, departure from the ideal of viewpoint diversity elevates values other than truth—the value to which universities ultimately should be exclusively committed. As Haidt puts it in an interview with John Leo of *Minding the Campus*,

> It's really scary that values other than truth have become sacred. And what I keep trying to say—this comes right out of my book *The Righteous Mind*—is that you can't have two

---

[34] I won't discuss whether it's psychologically realistic in the first place to think that exposure to marginalizing speech actually will have the result of making students more "resilient." Some authors seem to take it for granted that it does. But since I've argued that fostering resilience in the relevant contexts isn't a proper goal, discussion of the best methods for fostering that resilience strikes me as a red herring.

[35] The Heterodox Academy, as stated on their website (http://heterodoxacademy.org/), is

> a politically diverse group of social scientists and other scholars who want to improve our academic disciplines. We have all written about a particular problem: the loss or lack of "viewpoint diversity." It's what happens when everyone in a field shares the same political orientation and certain ideas become orthodoxy. We have come together to advocate for a more intellectually diverse and heterodox academy.

[36] See http://heterodoxacademy.org/.

> sacred values. Because what do you do when they conflict? And in the academy now, if truth conflicts with social justice, truth gets thrown under the bus. (Leo 2016)

Haidt makes a similar point in a lecture on the choice between two universities: “Coddle U.” and “Strengthen U.”: as the spokesperson of Strengthen U. he says,

> We are like a cult. We worship truth. We will throw anything overboard if it conflicts with truth. We believe the greatest act of sacrilege is to betray the truth, to hide it or distort it for any reason. (Haidt 2015)

On the other hand, exclusion of diverse views denies students the opportunity to be trained in “coping skills” (Haidt 2015) and the virtues of resilience and tolerance. If truth is supposed to be the overriding value, it is then odd to criticize speech codes on the grounds that we should be toughening up our students or training them in tolerance, as we’ve seen a number of opponents of speech codes argue. For “toughening up” is a psychological good and tolerance is a moral good, both of which would stand as alternative goods to truth. And, as Haidt would have to say, “what do you do when they conflict?” This would not be a problem if “toughening up” or tolerance could never conflict with truth. But even some of their proponents don’t make this claim. We’ve already seen Bollinger state explicitly that “anyone seriously interested in seeking truth would spend no time listening to such ideas” (1990, 980). As I’ve argued in this book and especially in this chapter, the values of “toughening up” and tolerance can conflict with the value of truth in two ways.

First, if we know that the speakers brought in to toughen up the students or teach them tolerance are uttering falsehoods, then we are prioritizing those other values over the value of truth because we are allowing falsehoods an inroad to the university that they wouldn’t otherwise have. A cult-member who genuinely worships only truth is not going to allow onto their campus known falsehoods in the service of a soft-hearted need to develop their students’ resilience. Second, if my argument in this chapter is sound, psychic harms done to students by inviting problematic speakers are the results of the students' *accurate* perceptions of genuine betrayal. Therefore, in toughening students up or teaching them tolerance, we end up ensuring that they don't respond accurately to real harms. We ignore the value of truth because we teach students to inaccurately judge the intrinsic harms done to them by our coddling of false marginalizing speech. It is only by refusing to invite relevant problematic speakers to campus that we in fact give due respect to truth—the single ultimate value.

8.4 CONCLUSION

One worry about this argument for the impermissibility of inviting problematic speakers is that it marginalizes those who are the targets of the offensive speech. It makes the active players the inviters and the problematic speakers. Those targeted don't really play a role. It's not their pain and the harm done to them—their silencing—that is the operant reason why the speakers should not be invited. Rather, it is the negative status of the speaker and the requirement of the inviters to maintain their moral integrity that is doing all the moral work. Especially because this argument is targeted primarily at members of non-marginalized groups—as made clear in the discussion of cases in which targets of the speaker are also the inviters—the discussion seems to essentially place the marginalized once again on the outside.

This worry is well-taken. One consequence is that this argument should only supplement and not take the place of other decisive arguments for not inviting such speakers—namely that such speakers create a hostile environment for members of marginalized groups, that they silence members of marginalized groups, and cause targets of their speech, as well as victims of their behavior, emotional and, at times, physical pain. Second, to avoid this difficulty, failing to invite problematic speakers is not a substitute for members of marginalized groups occupying central and active positions in speaker series and campus life. This is one consequence of Narayan's thesis that "the oppressed should speak for themselves and represent their own interests." For,

> Historically, those in power have always spoken in ways that have suggested that their point of view is universal and represents the values, interests and experiences of everyone. Today, many critiques of political, moral and social theory are directed at showing how these allegedly universal points of view are partial and skewed and represent the view points of the powerful and the privileged. The oppressed will, therefore, be quite warranted in being skeptical about the possibility of 'outsiders' adequately speaking for them.
>
> Besides, the right and power to speak for oneself is closely tied to the oppressed group's sense of autonomy, identity and self-respect. That it will foster and safeguard this sense of autonomy and self-respect is a good enough reason to say that the oppressed should speak for themselves . . . (Narayan 1988, 38)

Members of dominant groups don't absolve themselves of culpability for excluding members of marginalized groups from central positions in speaker series simply by refusing to invite problematic speakers. These positions should be as both organizers and invitees. Nor is there any need to neglect members of marginalized groups simply to get certain kinds of respectable arguments out there. If there is some requirement to invite a speaker who will argue on behalf of fathers in custody battles, there is no need to invite a Men's Rights advocate. As Lindy West argues, the difficulties men have in custody battles, among other issues emphasized by Men's Rights advocates

> . . .are feminist issues. And they are being worked on—by men and women, feminist and not—in egalitarian, caring, open-air spaces, where every word and action isn't just misogyny in disguise. There are domestic violence advocates who care about men—they're called domestic violence advocates. There are people working on improving job safety for men—they're called workplace safety advocates. There are people working to integrate dangerous, traditionally male professions such as coal mining and military combat—they're called WOMEN. We still have a long way to go, especially in the widespread thinking about masculinity and gender roles (again, a feminist issue!), but drenching this conversation in misogynist hate is nothing but a hindrance . . . (2014)

The people who should be invited to speak on important issues are not those whose speech or previous actions wrongly marginalize or victimize anyone. The injunction does not require us to ignore important views or neglect the rights of large sections of the population. It just requires universities not to invite those specific speakers who would require the kind of respect or deference that prevents those universities from standing in solidarity with those who need solidarity. Such speakers provide a sensationalist spark that might garner occasionally useful publicity for a university or a speakers program. But they don’t advance human knowledge in the way that other speakers would who offer more nuanced and less obviously false views that don’t marginalize or victimize members of the university population. These are the speakers that universities should be focusing on, so let’s get on with it. As West says in closing her article, “We have work to do.”

References

Adler, Jonathan (2004a). "Reconciling open-mindedness and belief." *Theory and Research in Education* 2(2): 127–142.

Adler, Jonathan E. (2004b). "Shedding Dialectical Tiers: A Social-Epistemic View." *Argumentation* 18: 279–293.

admin (2016). "Dillard students ask president 'how dare you' allow, ensure David Duke's safety." *The Drum*. Retrieved November 24, 2016, from http://www.thedrumnewspaper.info/?p=5011.

Aikin, Scott F. (2011). *Epistemology and the Regress Problem*. New York, Routledge.

Alcock, James E. (2003). "Give the Null Hypothesis a Chance: Reasons to Remain Doubtful about the Existence of Psi." *Journal of Consciousness Studies* 10(6–7): 29–50.

Altemeyer, Bob (2002). "Dogmatic Behavior Among Students: Testing a New Measure of Dogmatism." *The Journal of Social Psychology* 142(6): 713–721.

Ames, Roger (2011). *Confucian Role Ethics*. Manoa, University of Hawai'i Press.

Anderson, Charity (2015). "On the Intimate Relationship of Knowledge and Action." *Episteme* 12(3): 343–353.

Anderson, Erik A. (2016). "Democracy and Disingenuousness: The Case for Public Sincerity." The North American Society for Social Philosophy, 33rd International Social Philosophy Conference: Power & Public Reason, Carleton University, Ottawa.

Anonymous (2016). "Did Swinburne Get Swindled?" *Rightly Considered*.

Arpaly, Nomy (2011). "Open-Mindedness as a Moral Virtue." *American Philosophical Quarterly* 48(1): 75–85.

Baehr, Jason (2011). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford, Oxford University Press.

Bailin, Sharon and Mark Battersby (2016). "Fostering the Virtues of Inquiry." *Topoi* 35(2): 367–374.

Baker, Judith (1987). "Trust and Rationality." *Pacific Philosophical Quarterly* 68: 1–13.

Ballantyne, Nathan (2015a). "Debunking Biased Thinkers (Including Ourselves)." *Journal of the American Philosophical Association* 1(1): 141–162.

Ballantyne, Nathan (2015b). "The Significance of Unpossessed Evidence." *The Philosophical Quarterly* 65(260): 315–335.

Baril, Anne (2016). "The Ethical Importance of Roles." *Journal of Value Inquiry* 50: 721–734.

Barnes, Michael Randall (2016). "Confronting Hierarchies: Second-personal address in egalitarian social movements and reactionary hate speech." The North American Society for Social Philosophy, 33rd International Social Philosophy Conference: Power & Public Reason, Carleton University, Ottawa.

Barnhill, Anne (2014). "What is Manipulation?" in *Manipulation: Theory and Practice*. C. Coons and M. Weber. Oxford, Oxford University Press: 51–73.

Baron, Marcia (2003). "Manipulativeness." *Proceedings and Addresses of the American Philosophical Association* 77(2): 37–54.

Baron, Marcia (2014). "The Mens Rea and Moral Status of Manipulation." in *Manipulation: Theory and Practice*. C. Coons and M. Weber. Oxford, Oxford University Press.

Basu, Rima (manuscript). "A Moral Constraint on Philosophical Inquiry: In Defense of Hurt Feelings." Retrieved May 9, 2017, from https://drive.google.com/file/d/0BzVXA-PKVzvvMFV2NVBGNUVSSlE/view.

Bauer-Wolf, Jeremy (2017). "In-Your-Face Free Speech." *Inside Higher Ed*. Retrieved May 14, 2017, from https://www.insidehighered.com/news/2017/04/10/harvard-student-group-promotes-free-expression-through-controversial-speakers.

Bem, Daryl J. (2011). "Feeling the Future: Experimental Evidence for Anomalous Retroactive Influences on Cognition and Affect." *Journal of Personality and Social Psychology* 100(3): 407–425.

Bem, Daryl J. and Charles Honorton (1994). "Does Psi Exist? Replicable Evidence for an Anomalous Process of Information Transfer." *Psychological Bulletin* 115(1): 4–18.

Benhabib, Seyla (1992). *Situating the Self: Gender, Community and Postmodernism in Contemporary Ethics*. Cambridge, Polity Press.

Benhabib, Seyla (1996). "Toward a Deliberative Model of Democratic Legitimacy." in *Democracy and Difference: Contesting the Boundaries of the Political*. S. Benhabib. Princeton, Princeton University Press: 67–94.

Bernstein, David E. (2003). "Defending the First Amendment from Antidiscrimination Laws." *North Carolina Law Review* 82: 223–248.

Betty, L. Stafford and Bruce Cordell (1987). "God and Modern Science: New Life for the Teleological Argument." *International Philosophical Quarterly* 27(4): 409–435.

Betz, Gregor (2013). *Debate Dynamics: How Controversy Improves Our Beliefs*. Dordrecht, Springer.

Blackmore, Susan (1992). "Psychic Experiences: Psychic Illusions." *Skeptical Inquirer* 16: 367–376.

Blair, Olivia (2016). "Richard Dawkins dropped from science event for tweeting video mocking feminists and Islamists." *Independent*. Retrieved February 22, 2016, from http://www.independent.co.uk/news/people/richard-dawkins-vdeo-twitter-necss-event-feminism-a6841161.html.

Blow, Charles M. (2016). "No, Trump, We Can't Just Get Along." *The New York Times*. Retrieved November 24, 2016, from https://nyti.ms/2jB1S3P.

Blumenthal-Barby, J. S. (2014). "A Framework for Assessing the Moral Status of 'Manipulation'." in *Manipulation: Theory and Practice*. C. Coons and M. Weber. Oxford, Oxford University Press: 121–134.

Bockover, Mary I. (2012). "Confucian Ritual as Body Language of Self, Society, and Spirit." *Sophia* 51(2): 177–194.

Bogardus, Tomas (2009). "A Vindication of the Equal-Weight Vew." *Episteme* 6: 324–335.

Boghossian, Paul Artin (1997). "Analyticity." in *A Companion to the Philosophy of Language*. B. Hale and C. Wright. Oxford, Blackwell Publishers, Ltd.: 331–368.

Bollinger, Lee C. (1986). *The Tolerant Society*. Oxford, Oxford University Press.

Bollinger, Lee C. (1990). "*The Tolerant Society*: A Response to Critics." *Columbia Law Review* 90(4): 979–1003.

Boxill, Bernard R. (1976). "Self-Respect and Protest." *Philosophy and Public Affairs* 6(1): 58–69.

Brendtro, Larry and Scott Larson (2004). "The Resilience Code: Finding Greatness in Youth." *Reclaiming Children and Youth* 12(4): 194–200.

Brinlee, Morgan (2016). "'Is Hillary Clinton A Demon' Is A Poll Question Voters Have Actually Been Asked." *Bustle*. Retrieved February 22, 2017, from https://www.bustle.com/articles/192296-is-hillary-clinton-a-demon-is-a-poll-question-voters-have-actually-been-asked.

Broad, C. D. (1949). “The Relevance of Psychical Research to Philosophy.” *Philosophy: The Journal of the Royal Institute of Philosophy* 24(91): 291–309.

Brogaard, Berit (2013). “Phenomenal Seemings and Sensible Dogmatism.” in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. C. Tucker. Oxford, Oxford University Press: 270–289.

Broockman, David and Joshusa Kalla (2016). “Durably reducing transphobia: A field experiment on door-to-door canvassing.” *Science* 352(6282): 220–224.

Brown, Jessica (2008). “Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning.” *Noûs* 42(2): 167–189.

Brown, Jessica (2013). “Immediate Justification, Perception, and Intuition.” in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. C. Tucker. Oxford, Oxford University Press: 71–88.

Brown, Roger (1974). “Further Comment on the Risky Shift.” *American Pscyhologist* 29(6): 468–470.

Burns, Jean E. (2003). “What is Beyond the Edge of the Known World?” *Journal of Consciousness Studies* 10(6–7): 7–28.

Camosy, Charles C. (2016). “Listening to Trump voters: 3 first steps.” *Dallas News*. Retrieved November 21, 2016, from http://www.dallasnews.com/opinion/commentary/2016/11/21/listening-trump-voters–3-first-steps.

Carnaghi, Andrea and Anne Maass (2008). “Derogatory Language in Intergroup Context: Are ‘Gay’ and ‘Fag’ Synonymous?” in *Stereotype Dynamics: Language-Based Approaches to the Formation, Maintenance, and Transformation of Stereotypes*. Y. Kashima, K. Fiedler and P. Freytag. New York, Lawrence Erlbaum Associates: 117–134.

Carr, B. J. and M. J. Rees (1979). “The anthropic principle and the structure of the physical world.” *Nature* 278: 605–612.

Carroll, Noël (2001). “Four Concepts of Aesthetic Experience.” in *Beyond Aesthetics*. Cambridge, Cambridge University Press: 41–62.

Carter, Stephen L. (1999). *Civility: Manners, Morals, and the Etiquette of Democracy*. New York, Harper Perennial.

Cassell, Paul (2017). “Did Yiannopoulos secretly send more than 100 thugs to Berkeley to break up his own speech?” *The Washington Post*. Retrieved February 25, 2017, from http://wapo.st/2kcjvo7.

Chasmar, Jessica (2015). “Journalists banned from U.K. university’s debate on censorship.” *The Washington Times*.

Christensen, David (2007). “Epistemology of Disagreement: The Good News.” *The Philosophical Review* 116(2): 187–217.

Christensen, David (2010). “Higher Order Evidence.” *Philosophy and Phenomenological Research* 81(1): 185–215.

Christensen, David (2011). “Disagreement, Question-Begging and Epistemic Self-Criticism.” *Philosophers' Imprint* 11(6): 1–22.

Clarke, Roger (2013). “Belief is Credence One (in Context).” *Philosophers’ Imprint* 13(11): 1–18.

Clifford, William (1886). “The Ethics of Belief.” in *Lectures and Essays by the Late William Kingdon Clifford, F.R.S.* L. Stephen and F. Pollock. London, Macmillan and Co.: 339–363.

Cohen, Daniel H. (2009). “Keeping an Open Mind and Having a Sense of Proportion as Virtues in Argumentation.” *Cogency* 1(2): 49–64.

Comstock, Gary David (2001). *A Whosoever Church: Welcoming Lesbians and Gay Men into African American Congregations*. Louisville, Westminster John Knox Press.

Conee, Earl (2001). “Heeding Misleading Evidence.” *Philosophical Studies* 103: 99–120.

Copp, David (2007). “The Collective Moral Autonomy Thesis.” *Journal of Social Philosophy* 38(3): 369–388.

Couser, G. Thomas (2001). *The Cases of Oliver Sacks: The Ethics of Neuroanthropology*. Bloomington, IN, The Indiana University Foundation.

Cowan, Gloria and Cyndi Hodge (1996). “Judgments of Hate Speech: The Effects of Target Group, Publicness, and Behavioral Responses of the Target.” *Journal of Applied Social Psychology* 26(4): 355–374.

Crabtree, Steve (1999). “New Poll Gauges Americans’ General Knowledge Levels.” *Gallup*. Retrieved March 22, 2017, from http://www.gallup.com/poll/3742/new-poll-gauges-americans-general-knowledge-levels.aspx.

Curtis, Annaleigh (2016). “Some Stray Thoughts on Trigger Warnings from Someone with a Weird Phobia.” *The Second Shift: Academic Feminism After Hours*. Retrieved October 24, 2016, from http://www.secondshiftblog.com/2016/09/some-stray-thoughts-on-trigger-warnings-from-someone-with-a-weird-phobia/.

Dare, Tim (2016). “Robust Role-Obligation: How Do Roles Make a Moral Difference?” *Journal of Value Inquiry* 50: 703–719.

Deleuze, Gilles (2008). *Kant’s Critical Philosophy*. New York, Continuum.

Delgado, Richard (1982). “Words that Wound: A Tort Action for Racial Insults, Epithets, and Name-Calling.” *Harvard Civil Rights-Civil Liberties Law Review* 17: 133–181.

Delgado, Richard and Jean Stefancic (2004). *Understanding Words that Wound*. Boulder, CO, Westview Press.

Delgado, Richard and Jean Stefancic (2009). “Four Observations About Hate Speech.” *Wake Forest Law Review* 44: 353–370.

Descartes, René (1641/1985). “Meditations on First Philosophy.” in *Philosophical Writings of Descartes*. J. Cottingham. Cambridge, Cambridge University Press. 2: 1–62.

Dilling, Todd (2017). “A Response to Fantl’s Account of Permissible Closedminded Dismissal.”

Dreher, Rod (2016). “‘Shut Up, Bigot!’ The Philosophers Argued.” *The American Conservative*. Retrieved Sept. 26, 2016, from https://www.theamericanconservative.com/dreher/society-christian-philosophers-richard-swinburne/.

Dunn, Marcia (2003). “Lunar lunacy: Even NASA’s unsure how to counter claims of faked moon landings.” *Calgary Herald* 5 January 2003: A11.

Dworkin, Ronald (1977). *Taking Rights Seriously*. Cambridge, MA, Harvard University Press.

Dworkin, Ronald (1993). “Women and Pornography.” *The New York Review of Books* 40(17): 117–164.

Edgington, Dorothy (1986). “Do Conditionals Have Truth Conditions?” *Crítica: Revista Hispanoamericana de Filosofía* 18(52): 3–39.

Edgington, Dorothy (1995). “On Conditionals.” *Mind* 104(414): 235–329.

Eger, Martin (1988). “A Tale of Two Controversies: Dissonance in the Theory and Practice of Rationality.” *Zygon* 23(3): 291–325.

Elga, Adam (2007). “Reflection and Disagreement.” *Noûs* 41(3): 478–502.

Elgin, Catherine (1988). “The Epistemic Efficacy of Stupidity.” *Synthese* 74: 297–311.

Elgin, Catherine Z. (1999). *Considered Judgment*. Princeton, Princeton University Press.
Engel, Pascal (2005). "Truth and the Aim of Belief." in *Laws and Models in Science*. D. Gillies. London, King's College Publishers: 77–97.
Everett, Hugh (1957). "'Relative State' Formulation of Quantum Mechanics." *Reviews of Modern Physics* 29(3): 454–462.
Fagan, Emily and Mylses Sauer (2017). "Protesters crash effective altruism debate." *The Martlet*. Retrieved March 10, 2017, from http://www.martlet.ca/protesters-crash-effective-altruism-debate/.
Fantl, Jeremy and Matthew McGrath (2002). "Evidence, Pragmatics, and Justification." *The Philosophical Review* 111(1): 67–94.
Fantl, Jeremy and Matthew McGrath (2007). "On Pragmatic Encroachment in Epistemology." *Philosophy and Phenomenological Research* 75(3): 558–589.
Fantl, Jeremy and Matthew McGrath (2009). *Knowledge in an Uncertain World*. Oxford, Oxford University Press.
Fantl, Jeremy and Matthew McGrath (2012). "Pragmatic Encroachment: It's Not Just About Knowledge." *Episteme* 9(1): 27–42.
Fantl, Jeremy and Matthew McGrath (2014). "Do Practical Matters Affect Whether You Know?" in *Contemporary Debates in Epistemology*. M. Steup, J. Turri and E. Sosa. Oxford, Wiley Blackwell: 84–95.
Fasoli, Fabio, Maria Paola Paladino, et al. (2016). "Not 'just words': Exposure to homophobic epithets leads to dehumanizing and physical distancing from gay men." *European Journal of Social Psychology* 46: 237–248.
Faust, Jennifer (2008). "Can religious arguments *persuade*?" *International Journal for Philosophy of Religion* 63: 71–86.
Feldman, Richard (2001). "Voluntary Belief and Epistemic Evaluation." in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. M. Steup. Oxford, Oxford University Press: 77–92.
Feldman, Richard (2006). "Epistemological Puzzles About Disagreement." in *Epistemology Futures*. S. C. Hetherington. Oxford, Oxford University Press: 216–236.
Feldman, Richard (2007). "Reasonable Religious Disagreements." in *Philosophers without Gods: Meditations on Atheism and the Secular Life*. L. M. Antony. Oxford, Oxford University Press: 194–214.
Field, Hartry (1996). "The A Prioricity of Logic." *Proceedings of the Aristotelian Society* 96: 359–379.
FIRE (2017a). "FIRE—Defending individual rights in higher education." Retrieved February 19, 2016, from https://www.thefire.org/.
FIRE (2017b). *Spotlight on Speech Codes 2017*. Philadelphi, FIRE.
Fish, Stanley (1994). "There's No Such Thing as Free Speech, and It's a Good Thing, Too." in *There's No Such Thing as Free Speech and It's a Good Thing, Too*. Oxford, Oxford University Press: 102–119.
FitzPatrick, William (2008). "Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge." *Ethics* 118: 589–613.
Flat Earth Society, The (2016). "Membership Register." Retrieved March 20, 2017, from https://theflatearthsociety.org/home/index.php/about-the-society/membership-register.

Foa, Edna B., Barbara O. Rothbaum, et al. (1991). “Treatment of Posttraumatic Stress Disorder in Rape Victims: a Comparison Between Cognitive-Behavioral Procedures and Counseling.” *Journal of consulting and clinical psychology* 59(5): 715–723.
Foley, Richard (2009). “Belief, Degrees of Belief, and the Lockean Thesis.” in *Degrees of Belief*. F. Huber and C. Schmidt-Petri: 37–47.
Fora_Fauna (2015). “No Time-gate to Explain.” [weblog comment]. Retrieved April 19, 2016, from https://www.reddit.com/r/DestinyTheGame/comments/3pnzz2/no_timegate_to_explain/.
Frege, Gottlob (1953). *The Foundations of Arithmetic: A logico-mathematical enquiry into the concept of number*. New York, Harper & Brothers.
Freyd, Jennifer J. (1998). *Betrayal Trauma: The Logic of Forgetting Childhood Abuse*. Cambridge, MA, Harvard University Press.
Freyd, Jennifer J. (2008). “Betrayal Trauma.” in *The Encyclopedia of Psychological Trauma*. G. Reyes, J. D. Elhai and J. D. Ford. Hoboken, NJ, John Wiley & Sons, Inc.: 76.
Freyd, Jennifer J., Bridget Klest and Carolyn B. Allard (2005). “Betrayal Trauma: Relationship to Physical Health, Psychological Distress, and a Written Disclosure Intervention.” *Journal of Trauma & Dissociation* 6(3): 83–104.
Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford, Oxford University Press.
Fritz, James (forthcoming). “Conciliationism and Moral Spinelessness.” *Episteme*.
Fullerton, Elaine (2008). “The Development of ‘Nonviolent Communication’ in an Early Years Setting.” 1–75. Retrieved May 4, 2016, from http://www.gtcs.org.uk/web/FILES/FormUploads/development-of-nonviolent-communication-in-an-early-years-setting1771_332.pdf.
Gandini, Sara, Edoardo Botteri, et al. (2008). “Tobacco smoking and cancer: A meta-analysis.” *International Journal of Cancer* 122: 155–164.
Ganson, Dorit (2008). “Evidentialism and Pragmatic Constraints on Outright Belief.” *Philosophical Studies* 139(3): 441–458.
Garrett, Kelly R. and Brian E. Weeks (2013). “The Promise and Peril of Real-Time Corrections to Political Misperceptions.” *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*: 1047–1058.
Geiger, Brenda and Michael Fischer (2006). “Will Words Ever Harm Me? Escalation From Verbal to Physical Abuse in Sixth-Grade Classrooms.” *Journal of Interpersonal Violence* 21(3): 337–357.
Gibbard, Allan (2009). *Evolving Moral Knowledge: The Lindley Lecture*. Lawrence, KS, University of Kansas.
Giesea, Jeff (2014). “Stop Calling People Out.” *Harvard Business Review*. Retrieved April 2, 2016, from https://hbr.org/2014/10/stop-calling-people-out.
Ginet, Carl (1980). “Knowing Less by Knowing More.” *Midwest Studies in Philosophy* 5: 151–161.
Goldman, Alvin (1979). “What Is Justified Belief?” in *Justification and Knowledge*. G. S. Pappas. Dordrecht, D. Reidel: 1–23.
Gorin, Moti (2013). *The Nature and Ethical Signficance of Manipulation*. PhD, Rice University.
Gorin, Moti (2014). “Towards a Theory of Interpersonal Manipulation.” in *Manipulation: Theory and Practice*. C. Coons and M. Weber. Oxford, Oxford University Press: 73–97.
Gould, Jon B. (2010). *Speak No Evil: The Triumph of Hate Speech Regulation*. Chicago, University of Chicago Press.

Govier, Trudy (1997). “Arguing Forever? Or: Two Tiers of Argument Appraisal.” Ontario Society for the Study of Argumentation 2, Argumentation and Rhetoric, Brock University.
Govier, Trudy (1999). *The Philosophy of Argument*. Newport News, VA, Vale Press.
Greenawalt, Kent (1989). “Free Speech Justifications.” *Columbia Law Review* 89(1): 119–155.
Grimm, Stephen (2015). “Knowledge, Practical Interests, and Rising Tides.” in *Epistemic Evaluation: Purposeful Epistemology*. D. K. Henderson and J. Greco. Oxford, Oxford University Press: 117–137.
Gutmann, Amy and Dennis Thompson (1990). “Moral Conflict and Political Consensus.” *Ethics* 101(1): 64–88.
Gutting, Gary (2014). “Why Take a Stance on God?” *The New York Times*. Retrieved May 9, 2016, from https://nyti.ms/2jSAKu2.
Hackett, J. Edward (2016). “Richard Swinburne's Toxic Lecture on Christian Morality at the Midwest Meeting of the Society of Christian Philosophers.” *Philosophical Percolations*. Retrieved September 26, 2016, from http://www.philpercs.com/2016/09/richard-swinburnes-toxic-lecture-on-christian-morality-at-the-midwest-meeting-of-the-society-of-chri.html.
Hackman, Michelle (2016). “Al Jazeera accused Peyton Manning of using the doping drug HGH. Here's what we know.” *Vox*. Retrieved March 22, 2017, from http://www.vox.com/2016/2/6/10925434/peyton-manning-hgh-superbowl.
Haidt, Jonathan (2015). “Coddle U. vs. Strengthen U.: Which is best for students in the era of diversity?” *Heterodox Academy.org*. Retrieved May 29, 2016, from http://heterodoxacademy.org/wp-content/uploads/2015/11/haidt.transcript-of-Coddle-U-lecture.pdf.
Hammond, David, Geoffrey T Fong, et al. (2003). “Impact of the graphic Canadian warning labels on adult smoking behaviour.” *Tobacco Control* 12(4): 391–395.
Hardimon, Michael O. (1994). “Role Obligations.” *Journal of Philosophy* 91(7): 333–363.
Hardwig, John (1985). “Epistemic Dependence.” *The Journal of Philosophy* 82(7): 335–349.
Hardwig, John (1994). “Toward an Ethics of Expertise.” in *Professional Ethics and Social Responsibility*. D. E. Wueste. Lanham, MD, Rowman & Littlefield Publishers, Inc.: 83–102.
Hare, R. M. (1955). “Reply to Flew.” in *New Essays in Philosophical Theology*. A. Flew and A. C. MacIntyre. Norwich, SCM Press: 99–103.
Hare, William (2003). “Is It Good to Be Open-Minded?” *International Journal of Applied Philosophy* 17(1): 73–87.
Hare, William (2009). “Socratic Open-mindedness.” *Paideusis* 18(1): 5–16.
Hare, William and T. H. McLaughlin (1994). “Open-mindedness, Commitment and Peter Gardner.” *Journal of Philosophy of Education* 28(2): 239–244.
Harman, Elizabeth (2011). “Does Moral Ignorance Exculpate?” *Ratio* 24(4): 443–468.
Harman, Gilbert (1973). *Thought*. Princeton, NJ, Princeton University Press.
Harvey, Gordon (1996). “Elements of the Academic Essay.” Retrieved March 21, 2016, from http://www.brandeis.edu/writingprogram/writing-intensive/instructors/files/harveyelements.doc.
Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford, Oxford University Press.
Hawthorne, John and Jason Stanley (2008). “Knowledge and Action.” *Journal of Philosophy* 105(10): 571–590.
Hayes, Dennis (2009). “Academic Freedom and the Diminished Subject.” *British Journal of Educational Studies* 57(2): 127–145.
Hazlett, Allan (2012). “Higher-Order Epistemic Attitudes and Intellectual Humility.” *Episteme* 9(3): 205–223.

Helmore, Edward (2017). “Milo Yiannopoulos resigns from Breitbart News over pedophilia remarks.” *The Guardian*. Retrieved February 25, 2017, from https://www.theguardian.com/us-news/2017/feb/21/milo-yiannopoulos-resigns-breitbart-pedophilia-comments.

Hentoff, Nat (1991). “‘Speech Codes’ On the Campus And Problems of Free Speech.” *Dissent* 38: 546–549.

Herbert, Cassie (manuscript). “Talking about slurs.”

Herrnstein, Richard J. and Charles A. Murray (1994). *The Bell Curve: Intelligence and Class Structure in American Life*. New York, The Free Press.

Hess, Beth B. (1990). “Beyond Dichotomy: Drawing Distinctions and Embracing Differences.” *Sociological Forum* 5(1): 75–93.

Hick, John (1981). “An Irenaean Theodicy.” in *Encountering Evil: Live Options in Theodicy*. S. T. Davis. Louisville, Westminster John Knox Press: 38–51.

Hitchcock, David (2002). “The Practice of Argumentative Discussion.” *Argumentation* 16: 287–298.

Hod, Itay (2017). “Jennifer Holliday Pulls Out of Trump Inauguration, Apologizes to LGBT Community.” *The Wrap*. Retrieved February 17, 2017, from http://www.thewrap.com/jennifer-holliday-trump-inauguration-lgbt/.

Holton, Richard (1999). “Intention and Weakness of Will.” *The Journal of Philosophy* 96(5): 241–262.

Honorton, Charles and Diane C. Ferrari (1989). “‘Future-Telling’: A Meta-Analysis of Forced-Choice Precognition Experiments, 1935–1987.” *Journal of Parapsychology* 53: 281–308.

Hornsby, Jennifer (1993). “Speech Acts and Pornography.” *Women’s Philosophy Review* 10: 38–45.

Hornsby, Jennifer and Rae Langton (1998). “Free Speech and Illocution.” *Legal Theory* 4: 21–37.

Huemer, Michael (2001). *Skepticism and the Veil of Perception*. New York, Rowman and Littlefield.

Huemer, Michael (2006). “Phenomenal Conservatism and the Internalist Intuition.” *American Philosophical Quarterly* 43(2): 147–158.

Huemer, Michael (2007). “Compassionate Phenomenal Conservatism.” *Philosophy and Phenomenological Research* 74(1): 30–55.

Hume, David (1748/2007). *An Enquiry Concerning Human Understanding*. Cambridge, Cambridge University Press.

Hyman, Ray (2010). “Meta-Analysis That Conceals More Than It Reveals: Comment on Storm et. al.” *Psychological Bulletin* 136(4): 486–490.

Imlay, Jo and Jerry Howard (2006). “Listening for a Change.” *New Age Journal* 1993(Nov/Dec): 1–4.

Intenmann, Kristen and Inmaculada de Melo-Martín (2014). “Are there limits to scientists’ obligations to seek and engage dissenters?” *Synthese* 191: 2751–2765.

Isenberg, Daniel J. (1986). “Group Polarization: A Critical Review and Meta-Analysis.” *Journal of Personality and Social Psychology* 50(6): 1141–1151.

James, William (1962). “The Will to Believe.” in *Philosophy in the Twentieth Century*. W. Barrett and H. D. Aiken. New York, Random House, Inc. 1: 241–258.

Jantzen, Grace M. (1999). *Becoming Divine: Towards a Feminist Philosophy of Religion*. Indianapolis, Indiana University Press.

Jaschik, Scott (2016). “U Chicago to Freshmen: Don’t Expect Safe Spaces.” *Inside Higher Ed*. Retrieved October 24, 2016, from http://ihenow.com/2nSfTtz.

Jastrow, Joseph (1902). “Belief and Credulity.” *Educational Review* 23: 22–49.

Jay, Timothy (2009). “Do Offensive Words Harm People?” *Psychology, Public Policy, and Law* 15(2): 81–101.

Johnson, Brian Earl (2014a). *The Role Ethics of Epictetus: Stoicism in Ordinary Life*. Plymouth, Lexington Books.

Johnson, Harriet McBryde (2003). “Unspeakable Conversations.” *The New York Times Magazine* 6(50): 50–55, 74–79.

Johnson, Ralph H (2014b). *The Rise of Informal Logic: Essays on Argumentation, Critical Thinking, Reasoning and Politics*. Windsor, ON, University of Windsor.

Johnson, Ralph H. (1999). “More on arguers and their dialectical obligations.” Ontario Society for the Study of Argumentation 3: Argumentation at the Century’s Turn, Brock University.

Johnson, Ralph H. (2000). *Manifest Rationality: A Pragmatic Theory of Argument*. Mahwah, NJ, Lawrence Erlbaum Associates.

Johnson, Ralph H. (2001). “More on Arguers and Their Dialectical Obligations.” Ontario Society for the Study of Argumentation 4.

Johnson, Ralph H. (2008). “Responding to objections.” in *Controversy and Confrontation: Relating controversy analysis with argumentation theory*. F. H. van Eemeren and B. Garssen. Amsterdam, John Benjamins Publishing Company: 149–162.

Jonathan (2011). “Stanley v. Romano on ‘Progress in Philosophy’.” [weblog comment]. Retrieved February 24, 2017, from http://leiterreports.typepad.com/blog/2011/09/stanley-v-romano-on-progress-in-philosophy.html.

Jones, Barry (2015). *Coup d'Etat: The Overthrow of an American President*. Atlanta, Anaphora Literary Press.

Kabbany, Jennifer (2015). “Conservative female author of ‘The War on Men’ disinvited to speak at Williams College after protests.” *The College Fix*. Retrieved February 22, 2016, from http://www.thecollegefix.com/post/24742/.

Kabbany, Jennifer (2016). “Paleoconservative accused of ‘hate speech’ disinvited to speak at Williams College.” *The College Fix*. Retrieved February 22, 2016, from http://www.thecollegefix.com/post/26302/.

Kaplan, Mark (1996). *Decision Theory as Philosophy*. Cambridge, Cambridge University Press.

Kaplan, Martin F. (1977). “Discussion Polarization Effects in a Modified Jury Decision Paradigm: Informational Influences.” *Sociometry* 40(3): 262–271.

Kelley, Elaine (2001). “Daniel Pipes’ Acrimonious Remarks Embarrass Organizers of Portland Panel on ‘Healing Words’.” *The Washington Report on Middle East Affairs* 20(5): 57.

Kelly, Thomas (2005). “The Epistemic Significance of Disagreement.” *Oxford Studies in Epistemology* 1: 167–195.

Kelly, Thomas (2010). “Peer Disagreement and Higher-Order Evidence.” in *Disagreement*. R. Feldman and T. A. Warfield. Oxford, Oxford University Press: 111–174.

Kelly, Thomas (2011). “Following the argument where it leads.” *Philosophical Studies* 154: 105–124.

Kock, Christian (2007). “Dialectical Obligations in Political Debate.” *Informal Logic* 27(3): 233–247.

Koehler, Jonathan J. and Andrew D. Gershoff (2003). “Betrayal aversion: When agents of protection become agents of harm.” *Organizational Behavior and Human Decision Processes* 90: 244–261.
Kolers, Avery H. (2012). “Dynamics of Solidarity.” *The Journal of Politial Philosophy* 20(4): 365–383.
Kranz, Laura (2017). “‘Bell Curve’ author attacked by protesters at Middlebury College.” *The Boston Globe*. Retrieved March 6, 2017, from https://www.bostonglobe.com/metro/2017/03/04/middlebury/hAfpA1Hquh7DIS1doiKbhJ/story.html.
Kriss, Sam (2016). “In defense of personal attacks.” *Idiot Joy Showland*. Retrieved May 28, 2016, from https://samkriss.wordpress.com/2016/05/26/in-defence-of-personal-attacks/.
Kvanvig, Jonathan L. (2014). *Rationality and Reflection: How to Think About What to Think*. Oxford, Oxford University Press.
Kvanvig, Jonathan L. and Christopher Menzel (1990). “The Basic Notion of Justification.” *Philosophical Studies* 59: 235–261.
Kwong, Jack M. C. (2016a). “Open-Mindedness as a Critical Virtue.” *Topoi* 35(2): 403–411.
Kwong, Jack M. C. (2016b). “Open-Mindedness as Engagement.” *The Southern Journal of Philosophy* 54(1): 70–86.
Lackey, Jennifer (1999). “Testimonial Knowledge and Transmission.” *The Philosophical Quarterly* 49(197): 471–490.
Lackey, Jennifer (2010). “A Justificationist View of Disagreement’s Epistemic Significance.” in *Social Epistemology*. A. Haddock, A. Millar and D. Pritchard. Oxford, Oxford University Press: 298–325.
Lai, Calvin K., Maddelena Marini, et al. (2014). “Reducing Implicit Racial Preferences:: I. A Comparative Investigation of 17 Interventions.” *Journal of Experimental Psychology: General* 143(4): 1765–1785.
Lai, Calvin K., Allison L. Skinner, et al. (2016). “Reducing Implicit Racial Preferences: II. Effectiveness Across Time.”
Lammenranta, Markus (2011). “Disagreement, Skepticism, and the Dialectical Conception of Justification.” *International Journal for the Study of Skepticism* 1: 3–17.
Lance, Mark (2009a). “A Counterpetition to the APA Emerges: Discrimination Against Gay Men and Women Really Isn’t Discrimination!” [weblog comment]. Retrieved June 25, 2015, from http://leiterreports.typepad.com/blog/2009/02/a-counterpetition-to-the-apa-emerges-discrimination-against-gay-men-and-women-really-isnt-discrimina/comments/page/2/#comments.
Lance, Mark (2009b). “Today’s Pacific APA Business Meeting in Vancouver and the Motion for the APA to Enforce Its Existing Anti-Discrimination Rules.” [weblog comment]. Retrieved June 25, 2015, from http://leiterreports.typepad.com/blog/2009/04/todays-pacific-apa-business-meeting-in-vancouver-and-the-motion-for-the-apa-to-enforce-its-existing-.html#comments.
Langton, Rae (1999). “Pornography: A Liberal’s Unfinished Business.” *Canadian Journal of Law and Jurisprudence* 12(1): 109–133.
Lasonen-Aarnio, Maria (2014). “Higher-Order Evidence and the Limits of Defeat.” *Philosophy and Phenomenological Research* 88(2): 314–345.
Lasson, Kenneth (1999). “Controversial Speakers on Campus: Liberties, Limitations, and Common-Sense Guidelines.” *St. Thomas Law Review* 12: 39–94.

Lawrence, Charles R., III (1990). “If He Hollers Let Him Go: Regulating Racist Speech on Campus.” *Duke Law Journal* 3: 431–483.

Le Fevre, Deidre M., Viviane M. J. Robinson and Claire E. L. Sinnema (2014). “Genuine Inquiry: Widely Espoused Yet Rarely Enacted.” *Educational Management Administration & Leadership*: 1–17.

Lehrer, Keith (1971). “How Reasons Give Us Knowledge, or the Case of the Gypsy Lawyer.” *The Journal of Philosophy* 68(10): 311–313.

Leo, John (2016). “A Conversation with Jonathan Haidt.” *Minding the Campus*. Retrieved May 29, 2016, from https://shar.es/1UMTfD.

Lerner, Michael (2016). “Stop Shaming Trump Supporters.” *The New York Times*. Retrieved November 21, 2016, from http://www.nytimes.com/interactive/projects/cp/opinion/election-night–2016/stop-shaming-trump-supporters.

Leslie, John (2002). *Universes*. London, Routledge.

Levine, John M. and Eileen Russo (1995). “Impact of Anticipated Interaction on Information Acquisition.” *Social Cognition* 13(3): 293–317.

Lewis, David (1996). “Elusive Knowledge.” *Australasian Journal of Philosophy* 74: 549–567.

Lewis, David (1997). “Finkish Dispositions.” *The Philosophical Quarterly* 47(187): 143–158.

Littlejohn, Clayton (2011). “Defeating Phenomenal Conservatism.” *Analytic Philosophy* 52(1): 35–48.

Littlejohn, Clayton (2016). “A response to Dick Swinburne.” *Think Tonk*. Retrieved Sept. 26, 2016, from http://claytonlittlejohn.blogspot.ca/2016/09/a-response-to-dick-swinburne.html.

Locke, John (1693/1995). *An Essay Concerning Human Understanding*. Amherst, NY, Prometheus Books.

Lockhart, Jeff (2017). “here we go again: biological and sociological responses to the latest sex difference op-ed.” *scatterplot*. Retrieved March 1, 2017, from https://scatter.wordpress.com/2017/01/21/here-we-go-again-biological-and-sociological-responses-to-the-latest-sex-difference-op-ed/.

Lucian (c. 150 AD/1925). “Alexander; or, the False Prophet.” in *Lucian*. Cambridge, MA, Harvard University Press, Loeb Classic Library. 4: 173–254.

Lukianoff, Greg and Jonathan Haidt (2015). “The Coddling of the American Mind.” *The Atlantic* 316(2): 42–52.

Lycan, William G. (2009). “Giving Dualism Its Due.” *Australasian Journal of Philosophy* 87(4): 551–563.

Lynch, Michael (2012). *In Praise of Reason: Why Rationality Matters for Democracy*. Cambridge, MA, MIT Press.

Lynch, Michael P. (2010). “Epistemic Circularity and Epistemic Incommensurability.” in *Social Epistemology*. A. Haddock, A. Millar and D. Pritchard. Oxford, Oxford University Press.

Lynch, Michael P. (2013). “Epistemic Commitments, Epistemic Agency and Practical Reasons.” *Philosophical Issues* 23: 343–362.

MacKinnon, Catherine A. (1996). *Only Words*. Cambridge, MA, Harvard University Press.

Manfred, Tony (2015). “Tom Brady suspended 4 games for Deflategate.” *Business Insider*. Retrieved March 22, 2017, from http://read.bi/1JCMO9v.

Manne, Kate and Jason Stanley (2015). “When Free Speech Becomes a Political Weapon.” *The Chronicle of Higher Education*. Retrieved February 17, 2016, from http://chronicle.com/article/When-Free-Speech-Becomes-a/234207.

Markie, Peter J. (2013). “Searching for True Dogmatism.” in *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. C. Tucker. Oxford, Oxford University Press: 248–269.

Markwick, Betty (1978). “The Soal-Goldney experiments with Basil Shackleton: New evidence of data manipulation.” *Proceedings of the Society for Psychical Research* 56(211): 250–277.

Martel, Charles, Conservatrarian, et al. (2016). “Did Swinburne get Swindled?” *Rightly Considered*. Retrieved November 21, 2016, from http://rightlyconsidered.org/2016/09/26/did-swinburne-get-swindled/.

Martin, Jonathan and Amie Parnes (2008). “McCain: Obama not an Arab, crowd boos.” *Politico* October 10, 2008. Retrieved November 12, 2015, from http://www.politico.com/story/2008/10/mccain-obama-not-an-arab-crowd-boos–014479.

Martin, Matthew M., Sydney M. Staggers and Carolyn M. Anderson (2011). “The Relationships Between Cognitive Flexibility With Dogmatism, Intellectual Flexibility, Preference for Consistency, and Self-Compassion.” *Communication Research Reports* 28(3): 275–280.

Matsuda, Mari J. (1989). “Public Response to Racist Speech: Considering the Victim’s Story.” *Michigan Law Review* 87(8): 2320–2381.

McGrath, Sarah (2008). “Moral Disagreement and Moral Expertise.” *Oxford Studies in Metaethics* 3: 87–108.

McGrew, Timothy, Lydia McGrew and Eric Vestrup (2001). “Probabilities and the Fine-Tuning Argument: A Sceptical View.” *Mind* 110(440): 1027–1037.

McNally, Richard J. (2007). “Betrayal trauma theory: A critical appraisal.” *Memory* 15(3): 280–294.

Menary, Richard (2007). “Writing as thinking.” *Language Sciences* 29: 621–632.

Michaels, David (2008). *Doubt is Their Product: How Industry’s Assault on Science Threatens Your Health*. Oxford, Oxford University Press.

Mill, John Stuart (1850). *A System of Logic, Ratiocinative and Inductive*. New York, Harper & Brothers.

Mill, John Stuart (1859/2003). *On Liberty*. New Haven, Yale University Press.

Mill, John Stuart (1863). *Utilitarianism*. London, Parker, Son, and Bourn.

Miller, Meg (2016). “Your Facebook News Feed Really is An Echo Chamber.” *Fast Company Design*. Retrieved October 5, 2016, from https://www.fastcodesign.com/3060061/infographic-of-the-day/your-facebook-news-feed-really-is-an-echo-chamber.

Mills, Claudia (1995). “Politics and Manipulation.” *Social Theory and Practice* 21(1): 97–112.

Milton, Julie and Richard Wiseman (1999). “Does Psi Exist? Lack of Replication of an Anomalous Process of Information Transfer.” *Psychological Bulletin* 125(4): 387–391.

Mizrahi, Moti (2015). “Why Simpler Arguments are Better.” *Argumentation*: 1–15.

Moody-Adams, Michele M. (1994). “Culture, Responsibility, and Affected Ignorance.” *Ethics* 104: 291–309.

Moore, George Edward (1959/1983). “Four Forms of Scepticism.” in *Philosophical Papers*. London, Allen & Unwin: 196–226.

Morales, Lymari (2011). “Obama’s Birth Certificate Convinces Some, but Not All, Skeptics.” *Gallup.com*. Retrieved August 21, 2015, from http://www.gallup.com/poll/147530/obama-birth-certificate-convinces-not-skeptics.aspx.

Myers, Dale K. (2007). “Italian Experts test JFK assassination gun.” *Secrets of a Homicide*. Retrieved March 21, 2017, from http://jfkfiles.blogspot.ca/2007/06/italian-experts-test-jfk-assassination.html.

Myers, David G. and Helmut Lamm (1976). “The Group Polarization Phenomenon.” *Psychological Bulletin* 83(4): 602–627.

Nagel, Thomas (2001). *The Last Word*. Oxford, Oxford: Oxford University Press.

Narayan, Uma (1988). “Working Together Across Difference: Some Considerations on Emotions and Political Practice.” *Hypatia* 3(2): 31–47.

Nathan, Christopher (forthcoming). “Liability to Deception and Manipulation: The Ethics of Undercover Policing.” *Journal of Applied Philosophy*.

National Association of Scholars (1991). “The Wrong Way to Reduce Campus Tensions.” in *Beyond PC*. P. Aufderheide. St. Paul, MN, Graywolf Press: 7–8.

Neta, Ram (2007). “Propositional Justification, Evidence, and the Cost of Error.” *Philosophical Issues* 17: 197–216.

Neuman, Scott (2014). “1 in 4 Americans Thinks The Sun Goes Around The Earth, Survey Says.” *The Two Way: Breaking News From NPR*. Retrieved March 22, 2017, from https://n.pr/1gBrOmi.

Nicolas, Gandalf and Allison Louise Skinner (2012). “‘That’s So Gay!’ Priming the General Negative Usage of the Word Gay Increases Implicit Anti-Gay Bias.” *Journal of Social Psychology* 152(5): 654–658.

Nisbett, Richard E. and Timothy DeCamp Wilson (1977). “Telling More Than We Can Know: Verbal Reports on Mental Processes.” *Psychological Review* 84(3): 231–259.

Noggle, Robert (1996). “Manipulative Actions: A Conceptual and Moral Analysis.” *American Philosophical Quarterly* 33(1): 43–55.

Nordmann, Alain J., Benjamin Kasenda and Matthias Briel (2012). “Meta-analyses: what they can and cannot do.” *Swiss Medical Weekly* 142(w13518): 1–11.

Nyhan, Brendan and Jason Reifler (2010). “When Corrections Fail: The Persistence of Political Misperceptions.” *Political Behavior* 32(2): 303–330.

Nyhan, Brendan and Jason Reifler (manuscript). “The roles of information deficits and identity threat in the prevalence of misperceptions.”

Nyhan, Brendan, Jason Reifler, et al. (2014). “Effective Messages in Vaccine Promotion.” *Pediatrics* 133(4): 1–8.

O’Neill, Onora (1996). *Towards Justice and Virtue: A Constructive Account of Practical Reasoning*. Cambridge, Cambridge University Press.

Oatley, Keith and Maja Djikic (2008). “Writing as Thinking.” *Review of General Psychology* 12(1): 9–27.

Office of the Press Secretary (2015). “Remarks by the President at Town Hall on College Access and Affordability.” *The White House* September 14, 2015. Retrieved November 12, 2015, from https://www.whitehouse.gov/the-press-office/2015/09/15/remarks-president-town-hall-college-access-and-affordability.

Oreskes, Naomi and Erik M. Conway (2010). *Merchants of Doubt*. New York, Bloomsbury Press.

Paglia, Camille (2016). “Free Speech & the Modern Campus.” *The Smart Set*. Retrieved May 30, 2016, from http://thesmartset.com/free-speech-the-modern-campus/.

Palmer, John (1986). “Progressive Skepticism: A Critical Approach to the Psi Controversy.” *Journal of Parapsychology* 50: 29–42.

Pipes, Daniel (1997). *Conpiracy: How the Paranoid Style Flourishes and Where it Comes From*. New York, Simon and Schuster.

Pipes, Daniel (2001). “Protecting Muslims while rooting out Islamists.” *The Daily Telegraph* 14 September 2001.
Pipes, Daniel (2014). “Islamism’s Trajectory.” *Daniel Pipes: Middle East Forum*. Retrieved May 19, 2016, from http://www.danielpipes.org/blog/2014/05/islamism-trajectory.
Planka, Delaney (2016). “The Struggles of Being A College Republican.” *The Odyssey Online*. Retrieved April 19, 2016, from http://theodysseyonline.com/unr/struggles-college-republican/337288.
Plato (1997). *Plato: Complete Works*. Indianapolis, Hackett Publishing Company.
Plato (2000). *The Republic*. Cambridge, Cambridge University Press.
Price, George R. (1955). “Science and the Supernatural.” *Science* 122(3165): 359–367.
Pryor, Jim (2000). “The Skeptic and the Dogmatist.” *Noûs* 34(4): 517–549.
Quine, W. V. O. (1980). “Two Dogmas of Empiricism.” in *From a Logical Point of View*. Cambridge, Harvard University Press: 20–46.
Radin, Dean (1997). “On Science and Psychic Phenomena.” *The New Times* 13. Retrieved May 11, 2016, from http://www.deanradin.com/97–10-scipsych.htm.
Radin, Dean (2009). *Entangled Minds: Extrasensory Experiences in a Quantum Reality*. New York, Simon and Schuster.
Ramsey, John (2016). “Confucian Role Ethics: A Critical Survey.” *Philosophy Compass* 11: 235–245.
Rankin, Susan and Robert Dean Reason (2005). “Differing Perceptions: How Students of Color and White Students Perceive Campus Climate for Underrepresented Groups.” *Journal of College Student Development* 46(1): 43–61.
Rawlinson, Kevin (2016). “Anti-racism campaigner ‘stopped from speaking at NUS event’ over ‘Islamophobic’ claims.” *The Guardian*. Retrieved February 22, 2016, from http://www.theguardian.com/education/2016/feb/18/anti-racism-campaigner-stopped-speaking-nus-islamophobia-claims.
Rawls, John (1993). *Political Liberalism*. New York, Columbia University Press.
Rawls, John (1997). “The Idea of Public Reason Revisited.” *The University of Chicago Law Review* 64(3): 765–807.
Reed, Baron (2010). “A Defense of Stable Invariantism.” *Noûs* 44(2): 224–244.
Resick, Patricia A., Pallavi Nishith, et al. (2002). “A Comparison of Cognitive-Processing Therapy With Prolonged Exposure and a Waiting Condition for the Treatment of Chronic Posttraumatic Stress Disorder in Female Rape Victims.” *Journal of Consulting and Clinical Psychology* 70(4): 867–879.
Riggs, Wayne (2010). “Open-Mindedness.” *Metaphilosophy* 41(1–2): 172–188.
Rokeach, Milton (1954). “The Nature and Meaning of Dogmatism.” *Psychological Review* 61: 194–204.
Rosenberg, Marshall (2012). *Living Nonviolent Communication: Practical Tools to Connect and Communicate Skillfully in Every Situation*. Boulder, CO, Sounds True.
Rosenberg, Marshall B. (2003). *Nonviolent Communication: A Language of Life*. Encinitas, CA, PuddleDancer Press.
Rosenthal, Robert (1984). *Meta-Analytic Procedures for Social Research*. Beverly Hills, CA, Sage.
Ross, Josephine (2002). “The Sexualization of Difference: A Comparison of Mixed-Race and Same-Gender Marriage.” *Harvard Civil Rights-Civil Liberties Law Review* 37: 255–288.

Ross, Lee, David Greene and Pamela House (1977). “The ‘False Consensus Effect’: An Egocentric Bias in Social Perception and Attribution Processes.” *Journal of Experimental Social Psychology* 13(3): 279–301.

Russell, Bertrand (1950). “Can We Afford to Keep Open Minds?” *The New York Times*.

Samuelson, Peter L., Ian M. Church, et al. (2013). “The Science of Intellectual Humility White Paper.” 1–98. Retrieved September 21, 2016, from http://thethrivecenter.org/wp-content/uploads/2015/08/IH-White-Paper.pdf.

Schefter, Adam (2016). “NFL concludes Peyton Manning didn't use HGH, any other PEDs.” *ESPN*. Retrieved March 22, 2017, from http://es.pn/2a8Tmng.

Schoenfield, Miriam (2012). “Chilling out on epistemic rationality: A defense of imprecise credences (and other imprecise doxastic attitudes).” *Philosophical Studies* 158: 197–219.

Schoenfield, Miriam (2014). “Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief.” *Noûs* 48(2): 193–218.

Schroeder, Mark (2008). *Being For*. Oxford, Oxford University Press.

Scott, C. and P. Haskell (1973). “‘Normal’ Explanation of the Soal-Goldney Experiments in Extrasensory Perception.” *Nature* 245(5419): 52–54.

Scott, Christopher and Philip Haskell (1974). “Fresh light on the Shackleton experiments?” *Proceedings of the Society for Psychical Research* 56(209): 43–72.

Scriven, Michael (1984). “Critical for Survival.” *National Forum* 65: 9–12.

Searle, John R. (1968). “Austin on Locutionary and Illocutionary Acts.” *The Philosophical Review* 77(4): 405–424.

Seltzer, Rick (2016). “Why Host David Duke?” *Inside Higher Ed*. Retrieved November 24, 2016, from https://www.insidehighered.com/news/2016/11/03/dillard-university-stands-decision-host-debate-even-when-david-duke-candidate.

Shearman, Sachiyo M. and Timothy R. Levine (2006). “Dogmatism Updated: A Scale Revision and Validation.” *Communication Quarterly* 54(3): 275–291.

Sinclair, Lisa and Ziva Kunda (2000). “Motivated Stereotyping of Women: She's Fine if She Praised Me but Incompetent if She Criticized Me.” *Personality and Social Psychology Bulletin* 26(11): 326–342.

Slattery, Denis and Leonard Greene (2015). “Donald Trump silent on accusation that Barack Obama is Muslim at New Hampshire rally.” *NY Daily News*. Retrieved March 10, 2017, from http://www.nydailynews.com/news/politics/trump-silent-questioner-calls-obama-muslic-rally-article–1.2365196.

Smith, Eleanor R. (1992). “Differential Perceptions of Sexism.” *Women and Therapy* 12(1–2): 187–200.

Smithies, Declan (2011). “What is the Role of Consciousness in Demonstrative Thought?” *Journal of Philosophy* 108(1): 5–34.

Smithies, Declan (2012). “Moore’s Paradox and the Accessibility of Justification.” *Philosophy and Phenomenological Research* 85(2): 273–300.

Sober, Elliott (2004). “The Design Argument.” in *The Blackwell Guide to Philosophy of Religion*. W. E. Mann. Oxford, Blackwell: 117–147.

Sober, Elliott (2009). “Absence of evidence and evidence of absence: evidential transitivity in connection with fossils, fishing, fine-tuning, and firing squads.” *Philosophical Studies* 143(1): 63–90.

Solomon, Rose (manuscript). “Dissent and Public Opinion: When Counter-evidence is Presented.”

Sorensen, Roy (1988). "Dogmatism, junk knowledge, and conditionals." *The Philosophical Quarterly* 38: 433–454.
Southern Poverty Law Center (2005). *Speak Up! Responding to Everyday Bigotry*. Montgomery, AL, Teaching Tolerance.
Spiegel, Jim (2016). "Swinburne, Homosexuality, and the Society of Christian Philosophers." *Wisdom & Folly*. Retrieved March 14, 2017, from http://wisdomandfollyblog.com/swinburne-homosexuality-society-christian-philosophers/.
Stalsburg, Brittany L. (2016). "Trump and the Normalization of Hate." *The Huffington Post*. Retrieved October 5, 2016, from http://www.huffingtonpost.com/brittany-l-stalsburg-phd/trump-and-the-normalizati_b_10101492.html.
Stanford History Education Group (2016). *Evaluating Information: the Cornerstone of Civic Online Reasoning, Executive Summary*. Stanford, Stanford History Education Group.
Stanley, Jason (2005). "Knowledge and Practical Interests." Oxford, Oxford University Press.
Stevens, Katharina (2016). "The Virtuous Arguer: One Person, Four Roles." *Topoi* 35(2): 375–383.
Stevens, Stuart (2001). *The Big Enchilada: Campaign Adventures with the Cockeyed Optimists from Texas Who Won the Biggest Prize in Politics*. New York, The Free Press.
Stoehr, Taylor (1967). "Writing as Thinking." *College English* 28(6): 411–421.
Storm, Lance, Patrizio E. Tressoldi and Lorenzo Di Risio (2010a). "Meta-Analysis of Free Response Studies, 1992–2008: Assessing the Noise Reduction Model of Parapsychology." *Psychological Bulletin* 136(4): 471–485.
Storm, Lance, Patrizio E. Tressoldi and Lorenzo Di Risio (2010b). "A Meta-Analysis With Nothing to Hide: Reply to Hyman (2010)." *Psychological Bulletin* 136(4): 491–494.
Stump, Eleonore (1994). "The Mirror of Evil." in *God and the Philosophers*. T. V. Morris. Oxford, Oxford University Press.
Sturgeon, Scott (2008). "Reason and the Grain of Belief." *Noûs* 42(1): 139–165.
Subtle Energy Sciences Blog (2015). "A List of 100+ Peer-Reviewed Papers that Offer Scientific Evidence for Psi Phenomena." *Subtle Energy Sciences Blog*. Retrieved August 20, 2015, from http://subtleenergysciences.net/blog-native/2015/4/29/a-list-of–100-peer-reviewed-papers-that-offer-scientific-evidence-of-psi-phenomena.
Sunstein, Cass (2002). "The Law of Group Polarization." *The Journal of Political Philosophy* 10(2): 175–195.
Swinburne, Richard (2016). "Christian Moral Teaching on Sex, Family and Life." Midwest Society of Christian Philosophers Conference, Evangel University.
Taylor, Luke E., Amy L. Swerdfeger and Guy D. Eslick (2014). "Vaccines are not associated with autism: An evidence-based meta-analysis of case-control and cohort studies." *Vaccine* 32(29): 3623–3629.
Thomson, Judith Jarvis (1994). "Goodness and Utilitarianism." *Proceedings and Addresses of the American Philosophical Association* 67(4): 7–21.
Tolkien, John Ronal Reuel (1966). *The Fellowship of the Ring*. London, George Allen & Unwin.
Tosi, Justin and Brandon Warmke (2016). "Moral Grandstanding." *Philosophy & Public Affairs* 44(3): 197–217.
Turri, John (2010). "On the Relationship between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80(2): 312–326.
Tyler, Christopher W. and Maureen B. Clarke (1990). "The Autostereogram." *Proceedings of the International Society for Optical Engineering (SPIE)* 1256: 182–197.

U.S. Department of Veterans Affairs (2015). “Prolonged Exposure Therapy.” Retrieved March 7, 2016, from http://www.ptsd.va.gov/PTSD/public/treatment/therapy-med/prolonged-exposure-therapy.asp.

Ubha, Ravi (2016). “Maria Sharapova banned for two years over meldonium drug use.” Retrieved March 22, 2017, from http://edition.cnn.com/2016/06/08/tennis/maria-sharapova-two-year-ban-tennis-meldonium/.

van Eemeren, Frans H. and Peter Houtlosser (2001). “Strategic Maneuvering with the Burden of Proof.” Ontario Society for the Study of Argumentation 4.

Vavova, Katia (2014). “Moral Disagreement and Moral Skepticism.” *Philosphical Perspectives* 28: 302–333.

Vazire, Simine and Erika Carlson (2011). “Others Sometimes Know Us Better Than We Know Ourselves.” *Current Directions in Psychological Science* 20(2): 104–108.

Velleman, David (2000). *The Possibility of Practical Reason*. Oxford, Oxford University Press.

Vinokur, Amiram and Eugene Burnstein (1978). “Depolarization of Attitudes in Groups.” *Journal of Personality and Social Psychology* 36(8): 872–885.

Wagenmakers, E. J. (2014). “Bem is Back ” *Open Science Collaboration*. Retrieved May 23, 2016, from http://osc.centerforopenscience.org/2014/06/25/a-skeptics-review/.

Waldron, Jeremy (2012). *The Harm in Hate Speech*. Cambridge, MA, Harvard University Press.

Wallis, Charles (2008). “Consciousness, context, and know-how.” *Synthese* 160: 123–153.

Walton, Douglas and Erik C. W. Krabbe (1995). *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany, NY, State University of New York Press.

Weatherson, Brian (2005). “Can We Do Without Pragmatic Encroachment?” *Philosophical Perspectives* 19, Epistemology: 417–443.

Weatherson, Brian (2012). “Knowledge, Bets, and Interests.” in *Knowledge Ascriptions*. J. Brown and M. Gerken. Oxford, Oxford University Press: 75–103.

Wedgwood, Ralph (2002). “The Aim of Belief.” *Philosophical Perspectives* 16: 267–297.

Weisberg, Jonathan (2005). “Firing Squads and Fine-Tuning: Sober on the Design Argument.” *The British Journal for the Philosophy of Science* 56(4): 809–821.

West, Caroline (2003). “The Free Speech Argument Against Pornography.” *Canadian Journal of Philosophy* 33(3): 391–422.

West, Lindy (2014). “No, I Will Not Take the Men’s Rights Movement Seriously.” *Jezebel*. Retrieved March 4, 2016, from http://jezebel.com/no-i-will-not-take-the-mens-rights-movement-seriously–1532799085.

Wheeler, John (1973). “From Relativity to Mutability.” in *The Physicist's Conception of Nature*. J. Mehra. Dordrecht, Reidel: 202–247.

Whitcomb, Dennis, Heather Battaly, et al. (2015). “Intellectual Humility: Owning Our Limitations.” *Philosophy and Phenomenological Research* early view: 1–31.

White, Roger (2010). “You Just Believe That Because. . .” *Philosophical Perspectives* 24: 573–615.

Wieland, Jan Willem (2013). “Is Justification Dialectical?” *International Journal for the Study of Skepticism* 3(3): 182–201.

Wilder, Charlotte (2017). “A conspiracy theory about Kyrie Irving's insistence that the earth is flat.” *SB Nation*. Retrieved March 20, 2017, from http://www.sbnation.com/2017/2/18/14658008/kyrie-irving-earth-flat-conspiracy-theory-trolling-us-hell-yeah.

Williams, Bernard (1981a). "Internal and external reasons." in *Moral Luck*. Cambridge, Cambridge University Press: 101–113.

Williams, Bernard (1981b). "Persons, character, and morality." in *Moral Luck*. Cambridge, Cambridge University Press: 1–19.

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

Wilson, Stephen (2016). "Maria Sharapova's doping ban cut from 2 years to 15 months." *The Star*. Retrieved March 22, 2017, from https://www.thestar.com/sports/tennis/2016/10/04/maria-sharapovas-doping-ban-cut-from–2-years-to–15-months.html.

Wojcieszak, Magdalena (2011). "Deliberation and Attitude Polarization." *Journal of Communication* 61: 596–617.

Wolchover, Natalie (2012). "Are Flat-Earthers Being Serious?" *Live Science*. Retrieved March 20, 2017, from http://www.livescience.com/24310-flat-earth-belief.html.

Wong, Julie Carrie (2017). "UC Berkeley cancels 'alt-right' speaker Milo Yiannopoulos as thousands protest." *The Guardian*. Retrieved February 25, 2017, from https://www.theguardian.com/world/2017/feb/01/milo-yiannopoulos-uc-berkeley-event-cancelled.

Wood, Allen W. (2014). "Coercion, Manipulation, Exploitation." in *Manipulation: Theory and Practice*. C. Coons and M. Weber. Oxford, Oxford University Press: 17–51.

Wood, Zachary (2015). "I ran a speaker series to expose Williams students to unpopular ideas. It was deemed 'too offensive.'" *The Washington Post*. Retrieved February 17, 2016, from https://www.washingtonpost.com/posteverything/wp/2015/10/27/i-ran-a-speaker-series-to-expose-williams-students-to-unpopular-ideas-it-was-deemed-too-offensive/.

Young, Iris Marion (2006). "Responsibility and Global Justice: a Social Connection Model." *Social Philosophy and Policy* 23(1): 102–130.

Zdaniuk, Bozena and John M. Levine (1996). "Anticipated interaction and thought generation: The role of faction size." *British Journal of Social Psychology* 35: 201–218.