

7-15-2020

Engaging Others

Teresa Allen

University of Connecticut - Storrs, teresa.allen@uconn.edu

Follow this and additional works at: <https://opencommons.uconn.edu/dissertations>

Recommended Citation

Allen, Teresa, "Engaging Others" (2020). *Doctoral Dissertations*. 2565.
<https://opencommons.uconn.edu/dissertations/2565>

Engaging Others

Teresa Allen, PhD

University of Connecticut, 2020

Abstract: We all have views on controversial issues: gun control, climate change, and abortion, to name a few. Due to the controversial nature of these issues, it can often be difficult to engage with those with whom we disagree. Precisely because of how much these issues matter, however, it is important to figure out whether and how we ought to engage. Aside from moral and political concerns, there are epistemic concerns. For instance, is open-minded engagement always epistemically responsible or required? The purpose of this dissertation is to clarify the role that open-mindedness ought to play in our lives. Given that providing a comprehensive answer to the question of when we ought to be open-minded is incredibly complicated—especially when it comes to controversial issues—a secondary aim is to provide a framework for properly addressing the question at all. I argue that a serious exploration of the question of whether to be open-minded must also include a detailed discussion of epistemic normativity more generally. On the view I develop, what we ought to do epistemically is largely determined by what we ought to do morally, politically, and socially. This pragmatic approach to epistemic normativity is a decisive break from traditional lines of thought, according to which epistemic normativity ought to both be treated as distinct from, as well as theorized separately from, non-epistemic varieties of normativity. One important implication of my view is that, while we ought to be open-minded in general, the importance of being open-minded typically increases with the non-epistemic significance of the matter in question.

Engaging Others

Teresa Allen

B.A., University of Illinois at Urbana-Champaign **2012**

M.A., Northern Illinois University, **2014**

M.A., University of Connecticut, Storrs, **2017**

A Dissertation

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

at the

University of Connecticut

2020

Copyright by
Teresa Allen

2020

ii

APPROVAL PAGE
Doctor of Philosophy Dissertation
Engaging Others

Presented by
Teresa Allen, B.A., M.A.

Major Advisor -----

Michael P. Lynch

Associate Advisor -----

Heather Battaly

Associate Advisor -----

Richard Wilson

University of Connecticut

2020

iii

Acknowledgments:

A core commitment of this dissertation is that the accomplishments we enjoy in life are not ours alone; rather, they are shared by those who helped make them possible. There are many people I have to thank for their help and support in completing this dissertation. I will do my best to name them here, and I apologize to anyone I forget to mention.

I'd like to begin by thanking Hyrum LaTurner for introducing me to philosophy and instilling in me a love of the discipline. Hyrum, if not for your encouragement and support, I would never have pursued a degree in philosophy in the first place. Thanks also to the philosophy community at Northern Illinois University. Geoff Pynn, Lenny Clapp, and David Buller, thank you for your mentorship and for encouraging me to join UConn's graduate program. Thomas Shea, Emma Duncan, Adrian Seeley, Drew Johnson, and Louie Gularte, thank you for your friendship and for all of the philosophical discussions we've had both at NIU and beyond. I'm grateful to have met such wonderful friends and peers, especially so early in my graduate career.

Major thanks to the UConn philosophy community. I am incredibly grateful for having had the opportunity to spend the past six years in such a rich and supportive academic environment. Thanks to the professors I've had at UConn who have helped to make me a better philosopher, including Keith Simmons, Bill Lycan, Paul Bloomfield, Marcus Rossberg, and Dave Ripley. Thanks especially to Dorit Bar-On, who has been both a mentor and a friend. Dorit, I am grateful for our conversations, my experience working for you, and for the opportunity to have taken several of your seminars over the years. I have learned, and continue to learn, so much from you.

Thanks to my committee members, to whom I am especially indebted. Richard Wilson, thank you for always encouraging me to appreciate the nuances of complex issues. I have learned a great deal from your seminars and from our conversations. Heather D. Battaly, thank you for being at once a tough critic and a major source of philosophical and personal support. I'm very fortunate that your time at UConn intersected with mine. Thanks especially to Michael P. Lynch for supervising this project. Michael, this dissertation simply would not be what it is without your guidance and support. You encouraged me when I doubted myself, and you called me out when you knew I wasn't putting forth my best effort. You frequently pushed me out of my comfort zone, and trusted that I would be stronger for it. I am a better philosopher, and a better person, because of your influence and the opportunities which you have provided me.

I must also sincerely thank my fellow UConn graduate students. First and foremost, thanks to Rasa Davidaviciute, without whom I almost certainly would not have made it through graduate school. Rasa, you are an incredible philosopher and friend, and I cannot emphasize enough how much I've benefited from our friendship and philosophical discussions. Thanks also to Heather Muraviov and Mandy Long. To say that graduate school is tough is an understatement. Your friendship and support have made the journey easier and much more enjoyable. Thanks to Colin McCullough-Benner, Nathan Kellen, Andrew Tedder, Jared Henderson, Robin Jenkins, Drew Johnson, Dana Miranda, Emily Pierce, Christopher Rawles, Taylor Tate, Jenelle Salisbury, Phillip Barron, Hanna Gunn, Junyeol Kim, Ken Ito, Darian Spearman, Tom Meagher, Mengyu Hu, Yuhan Liang, Michael Heagarty, Jordan

Ochs, Andrew Parisi, Nate Sheff, David Pruitt, Ryo Tanaka, Mary Gregg, Emma Bjorngaard, Aliyar Ozercan, and Kristin Culbertson. I'm grateful for the conversations, laughs, and drinks we've shared. I've learned so much from you all over the years, and I am lucky to have had you as my graduate peers.

Thanks also to those affiliated with UConn's Humanities Institute, especially the members and speakers of UConn's Social Epistemology Working Group. Being a part of this group has been such a rewarding experience, and this project has greatly benefited from both the talks and the lasting conversations that stemmed from them. Thanks to the philosophers outside of UConn with whom I've had the fortune of interacting, and whose thoughts have been an influence on this work. This includes: Jeremy Fantl, Jason Baehr, Walter Sinnott-Armstrong, Nathan Ballantyne, Alessandra Tanesini, and Casey Johnson, as well as the audiences at the 2017 Political Epistemology Workshop in Copenhagen and the 2019 Humility and Conviction in Public Life Capstone Workshop, at which portions of this dissertation were presented. I am especially indebted to Sandy Goldberg, whose work has been a major influence on this project.

To my friends back home in Illinois—especially Kayli White, Kayla Badon, Jax Adamiak, Joe Craig, Molly Bilderback-Ulrich, Abby Moldovan, and Christin Hoffner—thank you for love, your friendship, and for being there for me throughout this long process. Thanks especially to my partner, Nathan Rock, who has been with me nearly every step of the journey. Nate, I would not be where I am today without your love, patience, flexibility, and humor. I am honored to call you my partner and my best friend.

Thanks to my family, who has never ceased to be supportive and encouraging (even when I decided to pursue philosophy instead of physical therapy). Special thanks to my brothers, Mark, Scott, Eric, David, Joe, and Mike Allen. I am grateful for your love, humor, and constant reminders to not take life too seriously. To my sister, Katie Martin, thank you for your love, your positivity, and for always being there for me. I am deeply indebted to you for your encouragement and support in completing this dissertation, especially during these last few months. Above all, thanks to my parents, Mary and George Allen. I cannot adequately express my gratitude for all that you have done for me. You have taught me through example how disagreeing parties can have rich and productive conversations about divisive issues. You have taught me to focus on what matters in life. And you have shown me unconditional love and support, for which I will be forever grateful.

Contents

Introduction	1
1 Should We Be Open-Minded?	6
1.1 Against open-minded engagement	7
1.1.1 Premise 1	9
1.1.2 Premise 2	19
1.2 Against Fantl	24
1.2.1 Analogy between controversial and non-controversial propositions	24
1.2.2 Knowing when one knows	31
1.2.3 When, and the extent to which, bracketing is licensed	51
1.2.4 Fantl’s definition of open-minded engagement	58
1.3 Conclusion	64
2 A Moral and Social Account of Epistemic Normativity	66
2.1 Fantl’s proposed sources of the epistemic obligation to be open-minded	69
2.2 Stage-setting remarks	79

2.3	Epistemic obligations are agent-universal	82
2.3.1	The moral underpinnings of epistemic normativity	83
2.3.2	The social underpinnings of epistemic normativity	90
2.4	The belief-universality of epistemic obligations	117
3	Further Implications of the Moral-Social Account	124
3.1	Obligations to have or form beliefs	125
3.1.1	Topics relevant to one’s moral obligations	129
3.1.2	Topics relevant to one’s social obligations	136
3.1.3	Topics relevant to one’s political obligations	139
3.2	Proportional strength of obligations	141
3.3	Meeting our epistemic obligations	144
4	Objections and Replies	150
4.1	I make claims about the obligations we have. Who is the “we?” . . .	150
4.2	The account is too demanding	156
4.2.1	Peels’ account of epistemic obligation	157
4.2.2	Response to Peels	162
4.3	A problem of regress	172
4.4	These are not epistemic obligations	179
5	Open-mindedness	181
5.1	Preliminary remarks	182
5.2	The objects of open-mindedness	187
5.3	Relevant intellectual options	190

5.4	Critical Receptivity	202
5.4.1	The Socratic strand	203
5.4.2	The Russellian strand	210
5.5	Failing to be Critically Receptive	215
5.5.1	Failing to be receptive at all	215
5.5.2	Engaging in an unserious manner	226
5.6	Conclusion	227
6	Conclusion	228
6.1	On epistemic normativity	229
6.2	On the obligation to open-mindedly engage	231
6.3	On the context-sensitivity of the obligation to open-mindedly engage	233
6.4	My view compared to Fantl's	235
6.5	Open-mindedness and public discourse	238
6.6	Closing remarks	240
	Bibliography	241

Introduction

We all have views on, or opinions about, controversial issues: gun control, climate change, abortion, and the existence of God, to name just a few. Most of us would like to think that we at least sometimes do better than have mere opinions about controversial issues—more than that, we think we can sometimes know controversial propositions. For example, we may think we know whether or not God exists, whether or not abortion is morally wrong, or whether or not ramping up gun control in our current political context would be politically and/or morally good. Often because of the controversial nature of these issues, it can be quite difficult to engage with those with whom we disagree, especially when we think that we have knowledge that they lack.

And yet, precisely because of the nature of these issues—including how much they matter personally, politically, and morally—it is important to figure out whether and how we ought to engage with those with whom we disagree. Aside from moral and political concerns about how we ought to engage, there are also epistemic concerns. For instance, what should one do when one is made aware of, or directly confronted with, an argument that challenges a controversial proposition one takes oneself to

know? Should one engage open-mindedly? Should one engage, but do so closed-mindedly? Or should one not engage at all?

In a moment where political tensions are high, and ideological and political polarization seem to be threatening the very fabric of our democracy, a thorough and detailed examination of the role open-mindedness ought to play in our lives—particularly when it comes to morally or politically significant or controversial matters—is vital. This dissertation provides such an examination. As such, it is heavily practically motivated. It aims not simply to better understand open-mindedness as a concept or in purely theoretical terms. More than this, and more importantly, it aims to clarify the role that open-mindedness ought to play in our lives.

It perhaps should come as no surprise that providing a comprehensive answer to the question of when we ought to be open-minded—especially when it comes to morally or politically controversial issues—is incredibly complicated. On account of this, a primary aim of this dissertation is to provide a framework for how best to address the question at all. I argue that a serious exploration of the question of when we ought to be open-minded must also include a detailed discussion of epistemic normativity more generally. In light of this, I devote a large portion of the discussion to examining what gives epistemic norms their force. According to the view I develop, what we ought to do epistemically is largely determined by what we ought to do morally, politically, and socially.

The arguments contained within thus rely on a crucial insight that critical race theorists and feminist theorists have emphasized time and again, and, more recently, that social epistemologists have been highlighting in their own work. Namely, the

fact that we exist as socially situated subjects simply cannot be ignored in our epistemological theorizing. One important implication of my view is that, while we epistemically ought to be open-minded in general, the importance of being open-minded is roughly proportional to the moral, social, or political significance of the matter at hand. This distinctly pragmatic approach to epistemic normativity is a decisive break from traditional lines of thought in epistemology, according to which epistemic normativity ought to be both treated as distinct from, as well as theorized separately from, non-epistemic varieties of normativity.

Importantly, my account is intended to be general in the sense that I do not intend to discuss in any amount of significant detail what any particular individual's obligations are, epistemic or otherwise. Rather, my primary goals are the following. First, I aim to elucidate the ways in which what we ought to do epistemically is generally, albeit intimately, bound up with what we ought to do morally, politically, and socially—whatever that amounts to in each person's case. Second, I aim to draw some broad lessons about what this means for how we ought to conduct our epistemic lives, focusing on the topic of open-mindedness.

Chapter outlines

We begin in **Chapter 1** by considering Jeremy Fantl's comprehensive argument against open-mindedness. His view is especially relevant to this dissertation, given that he is also primarily concerned with the appropriateness of open-minded engagement with respect to morally and politically controversial issues. On his view,

laypersons regularly find themselves in situations in which they are confronted with arguments and evidence which challenge a controversial proposition they know. He argues that in such situations, it would be epistemically irresponsible for a layperson to open-mindedly engage. While Fantl's view enjoys a certain amount intuitive appeal, I argue that it is ultimately unpersuasive. More than that, I argue that it is not one to be taken lightly. For following his recommendations would likely worsen some very real democratic problems we're already facing—including, but not limited to, political polarization. The discussion in this chapter serves as a launching point for the rest of the dissertation. In the remaining chapters, I develop a positive view which recommends a different verdict on the appropriateness and value of open-minded engagement.

In **Chapter Two**, we shift focus from open-mindedness in particular to epistemic normativity more generally. As I contend, if we are interested in determining whether and when one ought to be open-minded, we need to understand what gives epistemic norms their normative force. There are at least two reasons to offer in support of this contention. First, insofar as claims such as “one ought to be open-minded” are truly normative, rather than merely evaluative, we should be able to explain why. Second, understanding the (or a primary) source of epistemic normativity puts us in a much better position to understand the extent to which epistemic norms ought to govern our behavior. Drawing heavily from William K. Clifford's *The Ethics of Belief*, as well as more recent work from Sanford C. Goldberg, I defend the view that much of the epistemic normative pressure we face in our day-to-day lives derives from moral and social normative pressure.

In **Chapter Three**, I argue for two claims, both of which follow from the arguments contained in Chapter Two. First, the importance of meeting our epistemic obligations increases with the non-epistemic significance of the content of the beliefs. Second, non-epistemic considerations make it such that we sometimes ought to *form* beliefs about matters that are morally, socially, or politically significant. I also explain what meeting these obligations looks like in practice.

Chapter Four is dedicated to addressing potential objections that could be leveled against the moral-social account of epistemic normativity. **Chapter Five** is the point at which I return to the topic of open-mindedness. First, I explain how everything that has been argued thus far concerning epistemic normativity in general is relevant to open-mindedness in particular. I then offer an original account of open-mindedness that is inspired by Heather Battaly's and William Hare's respective accounts. On my view, to be open-minded is to be critically receptive to relevant intellectual options.

Chapter Six is the conclusion of our discussion. In it, I unite the various themes of the dissertation. I begin by highlighting the three most important aspects of my overarching view. To bring the discussion full circle, I also explain how the recommendations of my view differ from the recommendations of Fantl's view. I conclude by advocating for a more nuanced and encompassing understanding of open-mindedness that better illustrates the role it can (and should) play in democratic life.

Chapter 1

Should We Be Open-Minded?

According to the view I develop in the subsequent chapters, one ought to be open-minded to evidence, arguments, and the like that are relevant to one's non-epistemic—i.e., moral, political, and social—obligations. Given that one's non-epistemic obligations often concern, or are relevant to, morally and politically controversial matters, it follows that on my view, one often ought to open-mindedly engage with such evidence and arguments. Before developing my view, however, it will be instructive to consider reasons to think that open-minded engagement, particularly when it comes to controversial issues, may often be epistemically irresponsible. This discussion will serve as a useful point of departure for the remainder of the dissertation, as it both motivates my view, and it reminds us of what concerns to keep in mind in our defense of open-minded engagement.

The discussion in this chapter focuses on Jeremy Fantl's novel and comprehensive argument against open-minded engagement. His view is especially relevant to

this dissertation, given that he is also primarily concerned with the appropriateness of open-minded engagement with respect to morally and politically controversial issues. While Fantl's view enjoys a certain amount intuitive appeal, I argue that it is ultimately unpersuasive. I begin by introducing Fantl's core argument, and then I summarize the sub-arguments he provides in favor of its two premises. This will set me up to challenge these sub-arguments in the following section.

1.1 Against open-minded engagement

On Fantl's view:

You are open-minded toward an argument iff i) affective factors do not prevent you from 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.

You are closed-minded toward an argument iff either i) affective factors prevent you from 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. (2018, 12)

A few brief notes. First, "affective factors" are meant to pick out "non-cognitive and vicious influences on beliefs," such as implicit biases, desires, and fears (2018, 7). Thus, one is closed-minded to an argument if these kinds of influences are what prevent one from being persuaded by the argument or prone to violate procedural norms.

Second, “procedural norms” are the norms governing good argument evaluation. To not violate procedural norms, a person must not be careless in her evaluation of an argument; she must take the time to appreciate the various moves in an argument and see how they work together to support the conclusion; and so forth. Third (and notably), Fantl doesn’t take a stand on how best to understand “(un)willingness” and “significantly.” When it comes to being (un)willing to be persuaded by an argument, he is neutral because he thinks his account is compatible with at least a few conceptions of (un)willingness. Fantl refrains from specifying what being willing to be *significantly* persuaded by an argument amounts to, given that he thinks the issue is orthogonal to his view: “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” (2018, 13).

As for a “standard situation,” Fantl explains that:

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. (2018, 129, ft. 3)

He also remarks elsewhere:

It is a further feature of standard situations that you see obvious connections between reasons and what they’re reasons to do. (2018, 140, ft. 21)

The work this term is meant to do, I take it, is to distinguish cases in which a person actually has knowledge and cases in which she falsely believes she has knowledge—where this false belief stems from some kind of irrationality, bias, fatigue, or lack of intellectual skill or virtue.¹

With these preliminary remarks in place, we can now turn to Fantl’s central argument:

1. (The knowledge premise) There are standard situations in which you know controversial propositions and, thus, know that a relevant counterargument is misleading.
2. (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,
3. (C) There are standard situations in which you shouldn’t engage open-mindedly with a relevant counterargument against a controversial proposition. (2018, 130-31)

1.1.1 Premise 1

Fantl’s defense of the argument’s first premise ultimately relies on an analogy between controversial and non-controversial propositions.

Knowledge of non-controversial propositions

He begins by pointing out that there are cases in which a person’s knowledge of a *non*-controversial proposition is not threatened by her inability to expose the flaws in

¹It’s possible that Fantl also means to exclude Gettier situations from the category of “standard situations,” though it is not entirely clear.

relevant counterarguments. Imagine that a student enrolled in a philosophy course is presented with trick arguments designed to show that $1=0$ or that motion does not exist. Despite her inability to expose the flaws in the arguments, she does not in any way reduce her confidence that $1\neq 0$ or that motion exists. Not only is this probably a typical response in such situations, but also, Fantl maintains, this is precisely the correct response to have. As he puts it, “There is something especially problematic about being so open to Zeno’s arguments that if each step is compelling, you are prepared to reconsider whether things move. The same goes for trick arguments that $1\neq 0$ and Sorites-style arguments that I’m not bald” (2018, 130). After all, anyone presented with such trick arguments antecedently *knows* that $1\neq 0$ and that objects move. Remaining steadfast in one’s beliefs is not epistemically improper, for one continues to know the target propositions even when one is presented with counter-evidence one is unqualified to evaluate—evidence which, at least to untrained eyes, seems “apparently flawless.” In addition to Zeno’s paradoxes and trick arguments that $1=0$, Fantl also discusses paradoxes about vagueness and magic tricks.² In all of these cases he maintains that an inability to expose a flaw in sophisticated and apparently flawless arguments does not always (indeed, often does not) destroy a person’s knowledge.

There are at least two explanations to offer for why a person’s knowledge is not undermined in cases such as these. In the first place, the evidence one has against the conclusions of such trick arguments “swamps” any evidence the arguments themselves provide.³ In other words, any person who encounters such trick arguments will

²Thus, “argument” is meant to be understood broadly enough to include visual illusions.

³See Fantl (2018, 39) for some brief discussion.

first have encountered mounds upon mounds of perceptual, testimonial, and theoretical evidence all establishing exactly the knowledge the trick arguments challenge. Thus, any evidence these counterarguments provide to the effect that motion does not exist, that $1=0$, or that magic is real, is vastly outweighed by a lifetime's worth of contrary evidence the average person possesses. Even if the evidence provided by these trick arguments should be taken into account in *some* way, the evidence does not hold much weight. If being unable to expose a flaw in a trick argument means that one ought to lower one's confidence to some degree that the conclusion of the argument is false, the degree to which one ought to lower one's confidence will be very small indeed. It certainly would not significantly threaten one's epistemic position such that being exposed to a few more trick arguments will push one over the threshold of knowledge, causing one to lose knowledge one previously had. On the contrary, one can walk away from such situations virtually epistemically unscathed.

While the swamping explanation is one Fantl briefly considers, it is not the explanation that is ultimately central to his view. Instead, his chief argument for why knowledge of propositions such as $\langle 1 \neq 0 \rangle$, $\langle \text{motion exists} \rangle$, and $\langle \text{magic does not exist} \rangle$ survives being unable to expose the flaws of relevant counterarguments relies on the fact that the kinds of people who tend to be unable to expose flaws in such arguments are laypeople. On the one hand, these laypeople have what he calls "sufficient facility" with their own beliefs, which grounds their knowledge. Namely, they have a lifetime's worth of perceptual, testimonial, and theoretical evidence all of which supports exactly what the trick arguments are designed to challenge. On the other hand, precisely *because* they are laypeople—and not, as it were, professional

philosophers, logicians, or magicians—it is unsurprising that they find themselves unable to expose the flaws in such sophisticated, but nevertheless misleading, arguments. They lack the expertise to properly evaluate the evidence types, methodologies, or techniques invoked in the trick arguments in question.⁴ Far from making it harder for them to retain knowledge, Fantl argues their lack of expertise actually licenses them to not take the evidence provided by such counterarguments into consideration. Although it’s generally true that evidence relevant to one’s beliefs ought to make a difference as to what one believes—and while it is true that new evidence can destroy prior knowledge—Fantl argues that this is a notable exception.

To build his case, he draws on relevant discussions in the literature on lower-order and higher-order evidence. Typically, lower-order evidence is described as ordinary, run-of-the-mill evidence. For example, my lower-order evidence that it is raining outside is my visual perception of the rain pouring down on my lawn. Higher-order evidence—though it can be spelled out in various ways—is broadly speaking evidence *about* evidence.⁵ Now, there are a number of ways in which lower-order evidence and higher-order evidence might strengthen or conflict with one another. To illustrate, suppose I have lower-order perceptual evidence that the ball in front of me is red. Suppose I also have recently updated higher-order evidence that my visual perception is reliable—e.g., I have confirmation that the lighting in the room

⁴Thus, his *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” (2018, 60).

⁵Fantl chooses to stick with this neutral characterization of higher-order evidence. It is worth noting, though, that others define or discuss it differently. For instance, Thomas Kelly (2010) uses “higher-order evidence” to refer to evidence about the quality of one’s lower-order evidence. David Christensen (2010a) discusses higher-order evidence specifically as evidence that one is not evaluating one’s lower-order evidence properly.

is normal, and I recently had an eye exam confirming that my color vision is perfectly normal. This higher-order evidence strengthens, or confers justification on, my lower-order evidence, for it is evidence that my lower-order evidence is reliable. Other times, lower-order evidence and higher-order evidence can pull in opposite directions. Consider the following example from David Christensen (2010a) (which, notably, Fantl relies on heavily in his own discussion):

Drugs

I'm asked to be a subject in an experiment. Subjects are given a drug, and then asked to draw conclusions about simple logical puzzles. The drug has been shown to degrade people's performance in just this type of task quite sharply. In fact, the 80% of people who are susceptible to the drug can understand the parameters of the puzzles clearly, but their logic-puzzle reasoning is so impaired that they almost invariably come up with the wrong answers. Interestingly, the drug leaves people feeling quite normal, and they don't notice any impairment. In fact, I'm shown videos of subjects expressing extreme confidence in the patently absurd claims they're making about puzzle questions. This sounds like fun, so I accept the offer, and, after sipping a coffee while reading the consent form, I tell them I'm ready to begin. Before giving me any pills, they give me a practice question:

Suppose that all bulls are fierce and Ferdinand is not a fierce bull. Which of the following must be true? (a) Ferdinand is fierce; (b) Ferdinand is not fierce; (c) Ferdinand is a bull; (d) Ferdinand is not a bull.

I become extremely confident that the answer is that only (d) must be true. But then I'm told that the coffee they gave me actually was laced

with the drug. My confidence that the answer is “only (d)” drops dramatically. (2010a, 187)

The way Christensen diagnoses this situation is by saying that I, the experimental subject, am presented with higher-order evidence (namely, that I have already been administered the drug) that calls into question the reliability of my relevant lower-order evidence (namely, the fact that the puzzle is simple and that the answer at which I have arrived appears to be correct). Because this higher-order evidence so drastically calls into question the reliability of my lower-order evidence, Christensen maintains that the correct response for me to have is to bracket—i.e., to “in some sense, and at least to some extent, *put aside*” (2010a, 195, original emphasis)—my lower-order evidence. In short, I should not believe I have solved the puzzle correctly.

Agreeing with this diagnosis, Fantl proceeds to argue that the same lesson can be applied to cases in which laypeople encounter sophisticated arguments which challenge what they know, but which they lack the expertise to properly evaluate. In more detail, given that they are laypeople who lack the expertise to properly evaluate trick arguments, it would be entirely unsurprising if they think unsound or misleading arguments are sound or not misleading. In other words, a layperson would be just as likely as not to be convinced by a misleading but apparently flawless argument, for she is entirely unreliable when it comes to judging the merits of the argument. So, Fantl argues, whenever a layperson encounters such an argument, she ought to bracket the evidence the counterargument provides, and instead only count as relevant the evidence which she is in a position to properly evaluate. Since this relevant evidence would be the evidence she had prior to being exposed to the

counterargument, and assuming this evidence originally grounded her knowledge, then her knowledge can survive being unable to expose a flaw with the argument (2018, 55-64).

Furthermore, when her knowledge survives, she can know that the argument is misleading, even if she spends significant time with the argument, finds each step compelling, and is unable to expose a flaw. After all, if her knowledge survives, and if the conclusion of the counterargument conflicts with what she knows to be true, then the counterargument must be misleading.⁶ Importantly, Fantl does not necessarily commit himself to the claim that a layperson’s epistemic position will be entirely unaffected when she does not open-mindedly engage with a sophisticated counterargument (though he does seem quite friendly to this possibility). Nevertheless, he thinks that her knowledge can definitely *survive* such an encounter.

Knowledge of controversial propositions

Fantl thinks the lesson to draw from discussions on bracketing lower-level evidence is applicable not just to knowledge of non-controversial propositions, but to knowledge of controversial propositions as well. To give his readers a sense of the kinds of arguments he has in mind, he writes:

Here’s an incomplete list of arguments I wouldn’t be surprised to find apparently flawless even if they’re misleading: Bayesian arguments for

⁶Fantl does admit that this license to bracket only works if one is a layperson considering “lay” propositions—i.e., “propositions about whose subject ... [one lacks] special training” (2018, 34). In other words, experts in the relevant areas do not enjoy this license. At the same time, these relevant counterarguments presumably won’t look apparently flawless to experts, as they possess the relevant expertise to detect the flaws in such arguments. If an expert can’t detect a flaw in a counterargument that challenges what she takes herself to know, Fantl agrees that she should lower her confidence that the argument’s conclusion is false.

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, [and] legal arguments that restrictions on same-sex marriages are constitutional. (2018, 58)

Of course, a natural worry about this proposal is that there is crucial difference between controversial and non-controversial propositions. The fact that the non-controversial propositions we've been discussing are known by the vast majority of people—i.e., the very fact that they're not controversial—is arguably what makes closed-mindedly dismissing relevant counterarguments epistemically permissible. The same cannot be said of controversial propositions, knowledge of which is sometimes much more difficult to come by, and knowledge claims about which are often contested.

The problem intensifies if we accept the (admittedly controversial) view of pragmatic encroachment. According to advocates of the view—Fantl notably being in the midst⁷—the more it matters *pragmatically* that one knows, the harder it is to have or retain knowledge. In other words, non-epistemic factors can affect one's epistemic standing. The following pair of popular examples from the literature will help to illustrate:

Low Stakes

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It

⁷See Fantl and McGrath (2002, 2007, 2012, and 2019).

is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it wasn't very important that their paychecks are deposited right away, Hannah says, "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning."

High Stakes

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, "I guess you're right. I don't know that the bank will be open tomorrow." (Stanley 2005, 3–4)

If pragmatic encroachment is true, this poses a potential problem for Fantl's view. After all, he is interested in morally and politically *controversial* propositions, and so much hangs on whether or not one has the right political or moral views. They determine (at least in part) how one votes, how one treats others, and more. If any propositions raise the stakes for knowledge, these propositions presumably do. This is especially the case when—as often happens with respect to politically or morally controversial issues—one's views are challenged in a more or less public setting.

Fantl anticipates both of these worries, but does not think either of them cause significant problems for his view. Concerning the nature of controversial propositions, he remarks:

[T]he 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. (2018, 40)

If pragmatic encroachment is true, Fantl argues that this would simply mean the threshold for knowledge is higher with respect to controversial propositions. This is compatible with both having and retaining knowledge of them. He writes:

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. (2018, 145)

Even if it is more difficult to know controversial propositions than it is to know non-controversial propositions, Fantl maintains that as long as one's knowledge of controversial propositions is sufficiently secure, it can survive being closed-minded to relevant counterarguments. As was the case with non-controversial propositions, when a layperson antecedently *knows* a controversial proposition,⁸ this means she also knows that counterarguments which challenge her knowledge must be misleading, and that their conclusions are false. Her knowledge can thus survive being exposed to an apparently flawless relevant counterargument she lacks the expertise to properly or reliably evaluate.

⁸If pragmatic encroachment is true, this would also mean that her knowledge would need to be sufficiently secure such that higher practical stakes would not cause her to lose knowledge.

1.1.2 Premise 2

(The linking premise) If you know, in a standard situation, that a counterargument is misleading, you shouldn't engage open-mindedly with the counterargument. (2018, 130)

In defense of this premise, Fantl advances the following sub-argument:

P) This counterargument is unsound (even though it might well be such that, after engagement, the steps seem compelling and I won't be able to locate a flaw).

C) So I won't reduce my confidence in my own position in response to the counterargument (even if I engage with it, find each step compelling, and am unable to locate a flaw). (2018, 190)

He argues that, assuming one knows that P, then P is a decisive reason for drawing the conclusion.⁹ In other words, if you *know* the relevant counterargument is misleading or unsound, “*then you should draw the conclusions that P is a decisive reason for you to draw*” (2018, 190, original emphasis). The conclusion in this case is that you should not in fact reduce, nor even be willing to reduce, your confidence in the falsity of the relevant counterargument's conclusion. In more detail: if a layperson does engage open-mindedly, she puts her epistemic position at significant risk, because she is willing to be persuaded by an argument she knows it to be misleading. If she does not engage open-mindedly, she can walk away from the situation relatively (if not entirely, Fantl seems to think) epistemically unscathed. So, she should not engage open-mindedly. Rather, she should bracket the lower-level

⁹Precisely put, “P is a decisive reason for drawing a conclusion just in case, if you fail to draw that conclusion, it's not because of any failure of the connection between P and the conclusion; it's because of some weakness in your epistemic position with respect to P” (2018, 190).

evidence the counterargument provides, and instead only give weight to the evidence she originally possessed—that is, the evidence which grounds her knowledge of the controversial proposition.

To back up this argument, Fantl draws on Thomas Kelly for support, who advances a similar claim:

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)

Since the topic was mentioned, it is worth spelling out why this kind of dogmatism is taken to be of a different variety than that which inspired the dogmatism paradox, which can be summarized as follows.¹⁰ Say I know some proposition p . If I know that p , I know that any evidence that not- p is misleading. Since misleading evidence should be dismissed, if I know that p , I should dismiss all evidence (including future evidence) that not- p . We can illustrate using Paxson and Lehrer's (1969) "Tom Grabit" example. Suppose I witness Tom stealing a book from the library. Because I know Tom stole the book, I know that any evidence that he didn't steal the book is misleading. Since misleading evidence should be dismissed, I should dismiss all evidence (including future evidence) that Tom did not steal the book.

¹⁰This paradox was first introduced by Saul Kripke during a lecture he delivered to Cambridge University's Moral Sciences Club in 1972, and later discussed in print by Gilbert Harman (1973, 147-49).

Here is the problem. The conclusion of the argument seems to follow from the premises. And yet, knowing things does not license one to simply dismiss all evidence—including evidence one might encounter in the future—that conflicts with what one knows. The way the paradox is typically resolved—indeed, the way Harman himself resolves the paradox—is to say that it “overlooks the way actually having evidence can make a difference” (1973, 148). In other words, there’s a difference between (1) knowing, here and now, that any evidence which conflicts with what one knows is misleading evidence; and (2) coming to possess conflicting evidence at some point down the road, and knowing that this evidence is misleading. For, as Harman points out, encountering conflicting evidence in the future can cause one to *lose* knowledge. And, once one has lost knowledge, one cannot infer: “I know that p , so I know that this evidence that not p is misleading.”

To see the point, let’s return to our example of Tom stealing the book. When I see Tom steal the book, I come to know that he stole the book, and thus that any evidence there might be that he did not steal the book is misleading. Suppose I know that Tom is always in his office at this time of day, which means it’s unusual for him to be in the library around the time I saw him. Since I know Tom stole the book, though, I know that evidence such as this is misleading. And yet, it is possible that I will encounter evidence that will actually *change* what I know. In other words, I might encounter evidence in the future that will undermine my knowledge, so I will no longer know that all evidence that Tom did not steal the book is misleading.

Suppose Tom’s mother tells me that it was not Tom who stole the book, but rather his identical twin brother, Tim (about whom I’d been previously unaware).

As it turns out, Tom’s mother is a pathological liar, and Tim does not even exist. Since she fabricated the story, her testimony is obviously misleading. Once it is in my possession, however, I cannot simply dismiss it, for the very act of encountering it undermines my knowledge that Tim stole the book.¹¹

How is this different from the type of dogmatism Fantl and Kelly are interested in? Fantl insists that the difference lies in the fact that, in cases in which I know a controversial proposition, it is entirely *unsurprising* that there would be evidence—in this case, in the form of a misleading, but sophisticated, relevant counterargument—which is in tension with what I know and which to me appears flawless. Here is why. There are often many people on both sides, as it were, of controversial issues. Furthermore, it is unsurprising that there are people who have false beliefs about controversial issues, yet who nevertheless are quite smart and capable of producing misleading but sophisticated arguments in favor of their views. There are thus two related factors worth highlighting. First, that there exists evidence which conflicts with what one knows is not surprising because of the nature of the propositions in question (they are, after all, controversial). Second, not only is the *existence* of such evidence unsurprising, but also, it would also be unsurprising that this evidence is of a type that the average person can’t reliably evaluate.

Contrast this with cases in which it would be surprising for one to encounter apparently flawless but nevertheless misleading evidence. In the library case, it would indeed be surprising to hear from Tom’s mother that Tom’s identical twin brother stole the book. Not only would it be surprising to hear this testimony, but also, since

¹¹Thus, once my knowledge is compromised or undermined, I can no longer infer: “I know Tom stole the book, so this contrary evidence is misleading.”

I don't know that Tom's mother is a pathological liar, I would presumably have no reason for doubting what she says. Since I wouldn't be able to dismiss this evidence on the grounds that it's misleading, my knowledge would be lost. In general, and unlike in cases in which controversial propositions are under consideration, there is often not a great deal of conflicting evidence for or against non-controversial propositions such as these. (This is, after all, what makes them non-controversial.) In such cases, then, when there is relevant counter-evidence that appears to be flawless, it is more likely that such evidence is actually flawless than that it is simply *apparently* flawless. This is not to say that misleading but apparently flawless evidence which conflicts with known non-controversial propositions doesn't exist. It is simply to say that it would be surprising to encounter such evidence. And that fact that it *would* be surprising to encounter such evidence makes a difference as to whether or not my knowledge would survive such an encounter. As Fantl remarks:

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... [For] if you already suspect, before exposure, that the evidence is there, then exposure to the evidence will often not change whether you know. (2018, 41)

He thus maintains that a layperson who knows a controversial proposition should often not engage open-mindedly with sophisticated counterarguments which challenge that knowledge. Her knowledge can survive closed-mindedly dismissing the arguments because (1) the existence of such arguments is unsurprising; (2) she antecedently knows the controversial proposition being challenged; and (3) being layperson who lacks the relevant expertise to competently evaluate the counterargu-

ment, she is thereby licensed to bracket that evidence. Not only *can* her knowledge survive such closed-minded engagement (or closed-minded dismissal), but also, she *should not* engage open-mindedly. For in doing so she would be willing to allow her beliefs to be affected by evidence she knows to be misleading.

1.2 Against Fantl

In this section, I apply pressure to various aspects of Fantl's arguments. Taken together, these criticisms significantly undermine the plausibility of his view.

1.2.1 Analogy between controversial and non-controversial propositions

My first criticism concerns the fact that there are plausible alternative explanations for why it is perfectly normal for a layperson to not open-mindedly consider the possibility, say, that $1=0$ or that motion does not exist. What these alternative explanations have in common is that they regard these kinds of propositions as somehow unique, and it is precisely because of this that doubting them under normal circumstances would be inappropriate—perhaps even impossible. These explanations are important to consider because they are as good as, if not better than, the explanation Fantl provides. Furthermore, they put pressure on the analogy Fantl draws between controversial and non-controversial propositions.

The first alternative explanation I'd like to consider is that the kind of commitment we have to propositions such as $\langle 1 \neq 0 \rangle$ and $\langle \text{objects move} \rangle$ is so fundamental that there is no real sense in which we could open-mindedly consider their negations

in ordinary contexts. In other words, such commitments arguably play such a crucial role in one's belief system that one literally could not call them into question absent some kind of a radical change in one's overall belief system. There are a number of ways of articulating this possibility. One is by way of Ludwig Wittgensteinian's *hinge commitments*. In *On Certainty*, Wittgenstein suggests that there are certain commitments we all have in virtue of being inquirers which we are quite literally incapable of doubting. Having these hinge commitments—so-called since they are the commitments upon which all of our beliefs and doubts hinge—is necessary for one to have beliefs or doubts at all. Consider one of Wittgenstein's examples:

My having two hands is, in normal circumstances, as certain as anything that I could produce in evidence for it. That is why I am not in a position to take the sight of my hand as evidence for it. (1969, §250)

Let's unpack this to see what Wittgenstein is getting at. The first point to make is that the qualifier "being in normal circumstances" is important. For there can be unusual circumstances in which it makes perfect sense to check whether I have hands by looking at them. Suppose I have just been seriously injured by a nearby explosion. If I am not sure what bodily harm I have sustained, it would make sense to check whether I have hands by looking to see whether they're still attached. Under normal circumstances, however—i.e., when I have no reason to think that I have lost my hands—it would make no sense to *check* whether I have them. And it would certainly make no sense to check whether I have hands by using my eyesight. For if I am at the point where I have no reason to think that I have *lost* my hands, and yet I begin to seriously doubt that I have them—for instance, if a radical skeptic

has really started to wear me down—I would also be at the point where I would seriously doubt what my eyes are telling me. Put differently, once I am at the point that I doubt the existence of my hands, I will also be at the point where I would not be trusting that my eyes are providing me with accurate information about reality. My commitment to the fact that I have hands is so fundamental that once I start doubting that I have hands under normal circumstances, I will also be doubting a host of other things about reality, including my ability to ascertain the nature of the external world. As Wittgenstein remarks elsewhere:

If a blind man were to ask me “Have you got two hands?” I should not make sure by looking. If I were to have any doubt of it, then I don’t know why I should trust my eyes. For why shouldn’t I test my *eyes* by looking to find out whether I see my two hands? *What* is to be tested by *what*? (Wittgenstein 1969, §125, original emphasis)

My commitment to the existence of my hands is equally as strong and fundamental as my commitment that my eyes provide me with reliable information about the external world. I can’t “check” or “justify” one of these commitments by way of the other, for they are both equally fundamental. Moreover, once I call one into question—assuming I can at all—the other commitment (as well as a host of other commitments) goes up in the air with it.

Of course, Wittgenstein’s position is not uncontroversial. That said, it is important to consider for it offers a rich alternative explanation for why it seems perfectly reasonable to not open-mindedly consider the possibilities that $1=0$, that objects do not move, and that magic exists. To make the case as clear as possible, we’ll begin with the proposition <objects move>. This seems like a prime candidate

for a proposition we are hinge-committed to, for it is one of the commitments we swallow down when we form beliefs (at all) about the world. If I somehow got to the point where I began to seriously doubt that objects move, I would not know *what* to think. Doubting that motion exists would not be possible absent a radical shift in my entire belief system—a shift which would include coming to doubt a host of other propositions—for it is the kind of commitment that holds a host of other commitments in place. It is a hinge upon which they all turn.

A similar case can also be made for $1 \neq 0$, for coming to take seriously the possibility that such a basic mathematical truth is false would presumably not be possible without also calling into question a host of other propositions about the nature of the world. Coming to doubt that $1 \neq 0$ would presumably also cause one to doubt the reliability of a great deal of others' testimony, including one's educators, one's parents, and anyone else who has ever maintained that $1 \neq 0$. For if one had somehow been misled about such a basic proposition as *that*, what else has one been misled about? For many people, <magic does not exist> might also be a personal hinge proposition. Consider what might happen, for instance, if a scientist comes to eventually believe that magic exists. It is plausible that she would then think that she had previously been radically and fundamentally deceived about the way the world is.

To sum up, Wittgensteinian hinge commitments offer an alternative explanation for why it is perfectly reasonable to not open-mindedly engage with propositions such as the ones being discussed. The explanation is not that we are epistemically permitted to be closed-minded toward them, but rather, that we simply cannot be

open-minded toward them—at least not absent a radical shift in our belief system.

Now, suppose for a moment that this the right way to think about propositions such as <motion exists> and <1≠0>. Two options remain for what we can say about the controversial propositions Fantl discusses, both in the event that the controversial propositions he wishes to draw conclusions about are also hinge propositions, and in the event that they are not. First, let's consider whether any of the controversial propositions Fantl discusses can plausibly be considered hinge propositions. It's possible that the existence of God is something to which some people are hinge-committed. For instance, for an extremely devout believer, coming to seriously doubt that God exists may not be possible without also calling into question a host of other beliefs—e.g., her beliefs about an afterlife, her beliefs about morality (including what grounds moral normativity), and perhaps much more. In other words, she could not seriously doubt whether God exists without a radical change in her belief system.

Alternatively, for many people, that the Holocaust occurred is plausibly something to which they are hinge-committed. This seems especially true of Holocaust survivors, for that the Holocaust occurred is likely more certain to them than any contrary evidence that could be given to them.¹² This is not to say that everyone is hinge-committed to the existence of God or to the Holocaust having happened. Indeed, it's quite plausible that even if some people are hinge-committed to these propositions, other people are not. The point is simply that if some people are hinge-committed to certain controversial propositions, Fantl's argument would not

¹²This particular case is especially complicated, for over and above any epistemic considerations about whether not survivors should or even can engage open-mindedly with such arguments, there are also extremely important moral considerations involved.

go through in the way he intends. For a person cannot be open-minded to a counterargument if she is incapable of doubting the propositions the counterargument calls into question. That said, even if some people are hinge-committed to some controversial propositions, it is also the case that a good number of the controversial arguments Fantl has in mind do not concern hinge propositions.

This brings us to the second horn of the dilemma. Let's assume now that people are not hinge-committed to some (or any) of the controversial propositions that Fantl mentions—for example, that x number of deaths are caused in the name of various religions every year, or that Israel is justified in retaliating in ways that risk civilian lives. In this case, Fantl would thus have failed to establish the requisite analogy between the non-controversial arguments which he uses to motivate his view and the controversial arguments about which he aims to draw conclusions. For in misrepresenting the nature of our commitments to propositions like $\langle \text{objects move} \rangle$ he disguises the fact that our dismissal of arguments that motion does not exist is different in kind from our dismissal of arguments that aim to establish conclusions we are in principle capable of believing. More to the point, the fact that we cannot doubt that motion exists seems to tell us nothing about whether we are epistemically permitted (or required) to be close-minded to arguments that can affect our degrees of belief.

Wittgenstein's analysis provides a compelling alternative explanation for why we do not doubt certain non-controversial propositions such as $\langle \text{motion exists} \rangle$ and $\langle 1 \neq 0 \rangle$, even when we are presented with apparently flawless counterarguments. However, this is not the only explanation to offer. Another is that at least some of

the non-controversial propositions with which Fantl opens his discussion are arguably self-evident. For example, it is sometimes held that basic mathematical truths (such as that $1 \neq 0$) are self-evident in the sense that one knows the truth of the proposition simply in virtue of understanding its meaning. One knows that $1 \neq 0$ simply in virtue of understanding the meaning of “1” and “0,” and one needs no further evidence or proof that $1 \neq 0$. Moreover, and importantly, entertaining any real doubt that “ $1 \neq 0$ ” would be tantamount to losing at least some degree of understanding of “1” and “0.”

In contrast, one would be hard pressed to explain how true controversial propositions are self-evident. First of all, one often needs further evidence or proof to be convinced of the truth of controversial propositions even when one understands the meanings of the propositions’ terms. Furthermore, if one could know a controversial proposition simply in virtue of understanding the meanings of its terms, the proposition would likely not be controversial in the first place. Finally, even when one knows a controversial proposition, entertaining a doubt about the truth of the proposition does not necessarily amount to losing some degree of understanding of the meanings of its terms.

Of course, not all of the non-controversial propositions that Fantl discusses are self-evident. For example, “magic does not exist” is not self-evidently true. However, the point I’m making is not that all of the cases he discusses require *uniform* treatment. (In fact, it would be surprising if multiple treatments are not required.) Rather, my point is that the reason laypersons should not doubt such propositions likely has more to do with the kind of proposition they are (including perhaps the kind of role these propositions play in one’s overall belief system) and less to do with

the fact that, as a layperson, one cannot reliably evaluate sophisticated counterarguments.

If either of these alternative explanations (or some similar alternative) is correct, Fantl faces a dilemma. Either the controversial propositions he discusses are not special in the same way the discussed non-controversial propositions are, and so his analogy breaks down, or at least some of the controversial propositions he discusses are special in this way, in which case the reason laypeople are epistemically permitted to not open-mindedly engage with relevant counterarguments is not due to their inability to reliably evaluate such arguments, but rather due to the nature of the target propositions. Either way, Fantl would fail to establish what he intends.

There is no obvious reason to prefer Fantl's explanation over the alternatives just outlined. What this suggests, at the very least, is that diagnosing why we do not expect people to open-mindedly engage with certain trick arguments is a complicated affair. Whether or not one is epistemically permitted to not open-mindedly engage with a trick argument may very well depend on the nature (or content) of the commitment targeted by the trick argument. On account of this, I submit that much more needs to be said before we draw any analogies—as Fantl suggests we do—between trick arguments concerning propositions like <motion exists> and trick arguments concerning true controversial propositions.

1.2.2 Knowing when one knows

Next, I'd like to address the natural worry that what Fantl is advocating can quite easily be abused. More specifically, those who lack knowledge of controversial propo-

sitions, but who falsely think they have it, may take themselves to be licensed in closed-mindedly dismissing sophisticated counterarguments which challenge what they take themselves to know. Far from being a minor problem for his view, this worry is actually quite significant, for many people have quite strong political and moral convictions. Moreover, these convictions are often in tension with, or directly opposed to, the convictions of others. Since they cannot all be true, it stands to reason that many people who take themselves to know controversial propositions will actually lack such knowledge. To make matters worse, it can be difficult to tell from one's own perspective whether one actually knows some proposition p or whether one merely thinks one knows.¹³

Not surprisingly—given that it is such an obvious problem for his view—Fantl does anticipate and address this worry. Here is what he has to say:¹⁴

First, it is sometimes useful to be told that if you know a counterargument is unsound, 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.

Second, even if it's not clear when you know that counterarguments to

¹³After all, those who merely think they know are often just as convinced that they know as those who actually know!

¹⁴I quote him at length because I want the reader to be able to see how he responds to this worry in his words, especially because I do not think he does nearly enough to address it.

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 standarily 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 reasonable 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; [sic] that is particularly subject to misuse, since many who use the principle will tend to falsely think that their enemies are invariably reasonable 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 that they know that the counterarguments are misleading. This remains consistent with the truth and importance of the principle itself.

Any conditional with a normative consequent has this difficulty. There will be those who wrongly think the antecedent of this conditional is true and so draw the conclusion that the normative consequent is true. This neither impugns the truth of the 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. (2018, 133-34)

Advice or not?

Fantl is right that the antecedent of any conditional can be falsely taken to hold when it in fact does not. In this way, any conditional, whether normative or not, is subject to potential misuse. He is also right that some conditionals will be particularly subject to misuse. Among other things, biases, motivated reasoning, and the like can make one prone to think that antecedents of certain conditionals hold when they do not. Presumably, this is (at least in part) what would lead one to falsely think that one knows a controversial proposition when one doesn't, or that one's enemies are invariably unreasonable in their demands when they are not.

That said, it is hard to take seriously the claim that the linking premise—i.e., if you know, in a standard situation, that a counterargument is misleading, you shouldn't engage open-mindedly with the counterargument—is not advice. There are perhaps some cases in which it is interesting enough to simply discover metaphysical principles, whether or not they can be of any use to us. However, this does not seem to be one such case. If it truly is a metaphysical principle, it would largely be interesting insofar as it *could* be used as a piece of advice. Why else would Fantl devote an entire book to developing a provocative view with so much potential for

practical application unless he thought it could be at least somewhat action-guiding? Notice, too, that whereas he insists that the linking premise isn't advice, he does refer to it several times as advice, and suggests that it remains useful, despite the fact that many people falsely take themselves to know controversial propositions—hence might falsely take themselves to know that relevant counterarguments are misleading. Perhaps what he meant to say, then, is that it isn't *first and foremost*, or isn't *only*, a piece of advice.

And yet, we should (and presumably do) want to know how to utilize it as a piece of advice, if it's true. The reason Fantl's book is so interesting is precisely because it tackles head-on the question of what one epistemically ought to do in certain situations involving disagreements about controversial issues. We are interested in what the epistemically responsible thing to do in these situations because these situations matter, and they matter *a great deal*. We don't just want to know whether it's possible for knowledge to survive being closed-minded to relevant counterarguments. Nor do we just want to know what an epistemic agent in ideal conditions should do—for instance, an agent who is not a victim of cognitive biases or motivated reasoning, and who thus would not think she knows a controversial proposition (and that relevant counterarguments are misleading) when she does not. We want to know what *we* should do when we find ourselves in situations like the one Fantl describes. For *we* are typically not in ideal epistemic conditions, and *we* are frequently victims of cognitive biases and motivated reasoning. This is not to say that Fantl's principle can't be useful, if true. It just means that if it is true, it needs to be utilized with caution, for it really is particularly subject to misuse. Regrettably, Fantl doesn't say

much about what is involved in using the principle responsibly. He merely remarks that epistemologists have more work to do.

Utilizing Fantl's principle responsibly

Since Fantl does not do this, let's take a moment to consider what's involved in responsibly making use of metaphysical principles, especially those that may be particularly subject to misuse. We can begin by considering the other metaphysical principle Fantl discusses, namely: if your enemies are invariably unreasonable in their demands, don't negotiate in good faith with them. Let's call this *the enemy principle*. The enemy principle is interesting, if true. However, as with Fantl's principle, it is largely interesting because it provides some practical guidance. That is, it is largely interesting insofar as it can shed light on what one *should do* in a difficult moral or political situation. Now, as Fantl points out, this principle does seem particularly subject to misuse.¹⁵ Thankfully, there are measures one can take to make it more likely that one does not misuse or abuse the principle. For instance, one can check with unbiased, third parties in an attempt to determine whether one's enemies are being unreasonable in their demands. Alternatively, one can take time to think before responding to demands; one can try to put oneself in the shoes of one's enemy to determine why they might be making their demands; one can reflect one what a reasonable demand amounts to; and so forth. The point is that there are concrete measures one can (and should) take in order to make sure—or, at least, make it more likely—that one is not abusing this principle. To not take such concrete

¹⁵Of all the people one is prone to think are unreasonable in their demands, one's enemies typically sit at the top of that list.

measures would be morally and/or politically irresponsible. Moreover, the fact that there are concrete measures one can take in an effort to not abuse it is what elevates it to something morally and practically useful, and not just metaphysical principle that will either be abused or remain practically useless.

Similarly, Fantl's principle, if true, is interesting. Again, however, it is largely interesting insofar as it could be used as a piece of advice. As before, there are measures one can (and should) take in order to make sure—or, at least, make it more likely—that one does not misuse or abuse Fantl's principle. And yet, I will argue both that (1) not abusing Fantl's principle is quite difficult (indeed, there is an important sense in which not abusing Fantl's principle is more difficult to do than not abusing the enemy principle); and (2) that it matters quite a bit that Fantl's principle, if true, not be abused. For this reason, even if Fantl's principle is true, it nevertheless should be used with extreme caution.

As we've discussed, determining whether the antecedent of the enemy principle holds involves judging the reasonableness of one's enemies demands. This is not always easy to do, even for un-involved third parties. When it comes to Fantl's principle, however, checking to make sure the antecedent holds comes with its own set of difficulties, a few of which I'll touch on now.

First, determining whether you know a controversial proposition p —and, by extension, that a relevant counterargument is misleading—is in effect to determine whether you know *that* you know that p . At the very least, it involves making sure that you justifiedly believe that you know that p . As William Alston (1980) rightly pointed out, knowing that p and knowing (or justifiedly believing) *that* you know

that p are different things.¹⁶ In order to responsibly make use of Fantl's principle in one's own case, it is not enough that one knows some controversial proposition p . For in order to make sure one is not misusing or abusing the principle—a principle which is, as we've already established, particularly prone to misuse given that many people think they know controversial propositions when they do not (and, by extension, that relevant counterarguments are misleading when they are not)—one ought to make sure the antecedent of the principle holds by determining that one is justified in believing one has knowledge. This involves, among other things: reflecting on the type of justification one has for p ; judging the strength of one's justification for p ; and trying to determine whether the type of justification one has, given the strength of the justification, would suffice for knowledge of p . All of this can be quite difficult, especially when it comes to determining whether one knows a given controversial proposition, as there is often a great deal of both confirming and disconfirming evidence with regard to such propositions.

Second, in order to utilize Fantl's principle responsibly, one's judgements arguably cannot be unduly influenced by biases or similar factors. Allow me to explain. Fantl suggests that in the good cases where knowledge does survive, part of what allows for this is that the layperson dismisses the counterargument for the right reasons. He thus leaves open the possibility that a layperson's knowledge may not survive closed-

¹⁶In brief, Alston argued that philosophers have a tendency to conflate different levels of justification and knowledge by talking as if what holds true of one level of justification or knowledge holds true of another level. For instance, epistemologists would sometimes assume that if one knows that p , one also know *that* one knows that p . Conversely, they would sometimes make the mistake of assuming that if one's second-level justification or knowledge were compromised, one's first-level justification or knowledge would be compromised as well. He strongly urged epistemologists to treat different levels of justification and knowledge as distinct.

mindful dismissal of a relevant counterargument if, say, the layperson dismissed it because of “problematic affective factors or because [the layperson] didn’t adhere to various procedural norms” (2018, 63). And yet there is good evidence that human judgments are regularly influenced by such affective factors, and that we often cannot tell when our judgments are so influenced. This point is important, so we will take some time to explore it.¹⁷

There is a general consensus amongst psychologists today that—at least in a loose or broad sense—human brains are divided into two systems. As Daniel Kahneman influentially put it, System 1 is the fast-processing system responsible for our unreflective, automatic judgments. System 2, in contrast, is the slow-processing, reflective system that makes it possible for us to engage in complex problem solving, plan for the long term, consciously weigh reasons, and more. While System 2 is what allows us to engage in more sophisticated cognitive activities, System 1 is unquestionably indispensable. Given that we encounter a superabundance of sensory information on a moment-to-moment basis, and given that System 2 processing is laborious, we simply could not get on with our lives if our brains did not quickly process the vast majority of information with which we are regularly confronted.¹⁸ Thus, our brains regularly take “shortcuts” in service of the aim of fast and efficient processing. For example, we have a tendency to automatically and unconsciously compare and contrast new objects we encounter and objects with which we are already familiar.¹⁹

¹⁷Much of the following discussion on biases also appeared in Section 2 of Allen and Lynch (Forthcoming).

¹⁸See Kahneman (2011).

¹⁹This tendency to take shortcuts in judgments starts from an early age. Infants compare and contrast new objects (e.g., a novel ball) with objects with which they have experience (e.g., familiar balls), and make assumptions about the new objects on this basis. See Baldwin et al.

This spares us the trouble of consciously and individually evaluating each new object we encounter.

Despite the practical and epistemic usefulness of System 1, judgments produced by it can also be epistemically problematic for the following reason: System 1 aims primarily at speed and efficiency, and can produce false beliefs in service of those aims.²⁰ In other words, while having efficiently-formed beliefs might often overlap with having true beliefs—which would help explain why this tendency to generalize evolved—efficiency and truth can come apart. Thus, judgments produced by System 1 are not always good from the epistemic point of view. Furthermore, the unreflective part of our brains responsible for making these judgments is not really in the business of determining when this is the case. Important to note for the purposes of our discussion is that this tendency to categorize, or associate, new with old is not limited to inanimate objects, but applies to people as well.²¹ We regularly take shortcuts when making judgments about people, which can be epistemically (not to mention morally and politically) problematic.

It would help to say a bit more about how this categorization works. First, because the purpose of System 1 shortcuts is to increase efficiency, the brain, when categorizing, highlights similarities amongst members of the same category as well as differences between members of different categories. In the context of social categorizing, this means that both in-group similarities and out-group differences are emphasized. Second, since our brains make sense of new people (as it does with all

(1993). For a relevant discussion, see Leslie (2017).

²⁰See Kahneman (2011, especially pp. 105-07).

²¹See Gendler (2011, 38-41).

new objects) by comparing and contrasting them to ones with which we already have experience, in an important sense, how we perceive the world is shaped by what we already believe and know.²² When we identify someone as belonging to a certain group, we expect them to be like other members of that group, to behave in similar ways, and so forth. To some degree, we see what we expect to see. Third, the more familiar we are with certain categories—including what to associate with, and expect from, members of those categories—the more automatic and involuntary our judgments concerning those categories become.²³

In addition, the ways in which we conceive of the groups with which we automatically associate people are likely socially informed. This means that racist, sexist, or otherwise discriminatory associations that have permeated the larger social context can distort how we conceive of the categories themselves. For example, the pervasive stereotype that Muslims are dangerous can distort the ways in which we come to think of, and treat, Muslim people; the pervasive stereotype that women are less self-assured or competent than men can distort the ways in which we come to think of, and treat, women; and so on.²⁴ And, the more pervasive the stereotypes are, the more likely they are to influence how people conceive of certain groups. To take a concrete example, we can consider Keith Payne's (2001) studies, in which subjects identified guns faster, as well misidentified tools as guns more often, when primed by non-white faces than they did when primed with white faces. The widely accepted explanation for this result is that the subjects were more likely to associate blackness

²²See Banaji (2002, especially p. 15102).

²³For a discussion of all three of these points, see Gendler (2011, 39-40).

²⁴Regrettably, the stereotypes that influence and infect the way we categorize people are numerous, and it is not difficult to think up other similar examples.

with danger than they were to associate whiteness with danger.²⁵ Arguably the worst part about all of this is that once these associations become rigid and automatized, we may unconsciously stereotype people by associating them with certain groups—or as having certain characteristics—even when we consciously disavow these same stereotypes.²⁶

To sum up so far: our tendency to both pigeonhole people into categories with which we are already familiar, and to expect certain things from them as a result, is worrisome. Not only are the categories themselves often influenced by broader social norms and attitudes which can be misinformed, but also these norms and attitudes are often morally and politically harmful. The relevance of this to the current discussion is that the ways in which we think *about* moral or political issues that concern stereotyped groups can be—indeed, often may be—influenced by these pervasive stereotypes. Again, this can happen whether or not we consciously endorse said stereotypes. For instance, the way I think about a women’s health or women’s rights issue can be influenced by pervasive stereotypes about women. If my beliefs about women are influenced by the stereotype that women tend to be more emotional and less rational, level-headed, and informed than men, I may perceive women’s

²⁵Notably, when the subjects had more time to think before producing a judgment, the subjects’ judgments were not as biased. However, Payne cautions us not to read too much into this shift in behavior, as there is also evidence that people self-impose correcting measures depending on how they want to present themselves. In other words, the decrease in biased judgments when subjects were given more time was not necessarily due to their attempts to correct their biases—for example, because they were explicitly disavowing the stereotypes that were automatically triggered. Instead, it may have been due to their attempts to save face (Payne 2001, 187). For more on attitudes and self-presentational strategies, see Fazio et al. 1995; Dunton and Fazio (1997); and Greenwald et al. (1998).

²⁶For a sampling of relevant literature, see Dovidio and Gaertner (2000 and 2004), Devine and Sharp (2009), and Gendler (2011, especially pp. 41-44).

rights movements as being born out of over-sensitivity and an illegitimate sense of unfairness, rather than actual discrimination, marginalization, and/or oppression.

A related point is that *who* I hear about certain moral or political issues from, or who I discuss them with, can make a difference as to how I perceive the issue. If I am influenced by the stereotype that women are generally more emotional and less rational or level-headed than men, then I will be more likely to take such movements seriously if I hear them defended by men than women. After all, given that I understand men to be more level-headed, informed, and rational, I will perceive them as being able to more accurately appraise the relevant issues. This very point is why members of marginalized groups seeking justice often ask that those in relative positions of privilege or power to assume the role of an ally.

Now, in addition to the worrisome evidence concerning System 1 reasoning, there is also evidence that System 2 is not as rational (or reason-driven) as we might have assumed, and that at least some types of judgments produced from this system are products of rationalization rather than proper reasoning. For instance, many people are strongly inclined to say that incest is wrong, even when it is consensual, it is done in secret (so nobody will find out and be appalled by it), protection is used, and no party to the incest will regret having done it in the future.²⁷ In other words, the reasons typically cited for the wrongness of incest are not apposite. And yet, people tend to insist that it is immoral, even as they are empty-handed when pressed to come up with an explanation. According to Haidt's (2001) social intuitionist model of moral reasoning, what explains such cases is that people make moral judgments on

²⁷See Haidt et al. (2000). This manuscript is unpublished, but it is cited in Haidt (2001).

the basis of intuition. In cases where someone is asked to explain why she made this or that moral judgment, “[she then] becomes a lawyer trying to build a case rather than a judge searching for the truth” (Haidt 2001, 814). His more general claim is that when people are presented with a moral problem, they respond intuitively and only later provide reasons, if at all. And, importantly, the reasons given are meant to accommodate a person’s intuitions.²⁸

On top of what has been cited, there is a wealth of other research corroborating these and similar points. For instance, another type of bias (or family of biases) worth mentioning is confirmation bias, which can be summed up as the following pair of tendencies: (1) we tend to seek out, or take more seriously, evidence and argumentation that confirms what we already believe, and (2) we tend to fail to seek out, or take seriously, evidence or argumentation that is in tension with what we believe.²⁹ These tendencies, coupled with the negative aspects of both System 1 and System 2 processing, should be worrying. It is bad enough that System 1 and System 2 can produce beliefs that are not well supported by the evidence, but are rather the effects of biased thinking. What makes matters worse is that we will often continue holding these beliefs, even if they are not evidentially supported, since we tend to prioritize and preferentially treat evidence that confirms what we already believe.

To make matters even worse, there is also evidence suggesting we do not have a reliable psychological capacity for telling the difference between beliefs formed on the basis of a given recognizable reason, and beliefs formed via implicit bias.

²⁸This, of course, is not to deny that people can determine that their moral intuitions are misguided. It is simply to suggest that, for many people and in general, our moral sensibilities are guided less by reason and more by gut reaction than we recognize.

²⁹See, for example, Nickerson (1998) and Kunda (1999).

In particular, many people exhibit what's known as a "bias blind spot"—i.e., an inability to recognize when their own judgments are informed by biases. Thus, a general recognition that judgments are often formed (at least partially) in response to non-rational factors only gets us so far if everyone considers themselves an exception to the rule. Raising this very point, Pronin and Schmidt (2013) draw from a host of psychological studies which demonstrate, across a number of domains, the ignorance people have of their own biases. Especially germane to this discussion are the blind spots people have to their ideological, prejudicial, and self-interested biases. For instance, as Pronin and Schmidt discuss, people are prone to deny that their political or other ideological beliefs are formed as a result of partisan or ideological alignment, insisting that their beliefs were reached instead on the basis of sound reasoning.^{30,31} Likewise, people are prone to deny that their judgments are race- or gender-biased,³² as well as that their judgments or decision have been influenced by monetary or social incentives.³³

Perhaps the existence of the bias blind spot should come as no surprise, as we have been aware for quite some time of the "Dunning-Kruger effect," whereby people tend to both generally overestimate their abilities, as well as rate their abilities as higher

³⁰This is true even though people are inclined to think their political or ideological opponent's views were largely shaped by their biases.

³¹See Cohen (2003); Pronin et al. (2007); and Robinson et al. (1995).

³²See Dovidio and Gaertner (1991, 2004); Uhlmann and Cohen (2007); and Vivian and Berkowitz (1992).

³³See Dana and Loewenstein (2003); Epley and Dunning (2000); Heath (1999); and Miller and Ratner (1998).

than those of their peers.^{34,35} Assuming this tendency carries over to the problem of bias-detection, not only is it the case that we regularly overestimate our ability to detect when our thinking is biased, but also, the worse we are at detecting this, the better we think we are.

To sum up this second point: we know that these biases exist, that they are pervasive, and that many people consider themselves to be exceptions to the rule. Moreover, while we know that our thinking is flawed in general, and that in general we think the grounds for our beliefs are stronger and less influenced by biases and other affective factors than they are, we should expect that matters will be worse when it comes to controversial topics. For one, such topics are often emotionally charged. Moreover, controversial topics are likely to bring to mind—whether consciously or unconsciously—the very stereotypes that biases embody, and these stereotypes can influence judgments in ways that are far less than apparent. For instance, discussions about the travel ban may bring to mind negative stereotypes about immigrants, and these stereotypes may negatively influence a person’s judgments in ways she is not even aware of. Moreover, this can happen even when the person issuing the judgments consciously disavows said stereotypes. Similarly, discussions about policies proposed by members of a certain political party may bring to mind stereotypes about members of that political party, and these stereotypes may negatively influence a person’s judgments in ways she is not even aware of. And so on.

³⁴Of course, this is not to say that everyone suffers from a tendency to overestimate their abilities. Some people have the opposite problem—namely, they regularly underestimate their abilities, rating themselves as less capable than they in fact are.

³⁵The interesting, and especially worrying, aspect of this effect is that oftentimes the less competency people have with respect to a certain skill, the more drastic their overestimation of their own capabilities. For more on this, see Kruger and Dunning (1999).

There is thus another way in which not abusing Fantl's principle is quite difficult—that is, if we assume that utilizing it responsibly requires one's judgements to not be unduly influenced by biases or other affective factors. For not only is it the case that our judgments are regularly influenced by such affective factors, but also there's a strong case to be made that our judgments are more likely to be negatively influenced by such factors when it comes to controversial issues than when it comes to non-controversial issues. What's more, it's very difficult to tell when one's own judgments are being heavily influenced by affective factors.

The third and final factor that makes utilizing Fantl's principle responsibly difficult is: when it comes to figuring out whether your epistemic position is strong enough to license being close-minded, it's not enough to figure out that your epistemic position is strong enough for knowledge. That is, it's not enough to justifiably believe that you know. You must also figure out whether your knowledge is secure enough to survive being unable to expose the flaw in a relevant counterargument. In other words, to determine that it's strong enough to license close-mindedness, you must determine (1) roughly how much being unable to expose such a flaw would weaken your epistemic position, and (2) that your epistemic position (prior to exposure to the counterargument) is sufficiently secure such that being unable to expose the flaw will not undermine your knowledge.

Related to this is the concern about pragmatic encroachment introduced in §1.1, according to which higher practical stakes make it harder to come by and/or retain knowledge. One reason that pragmatic encroachment is relevant is that, as we've discussed, our beliefs about controversial issues often make a moral or political dif-

ference. They influence how we vote, how we treat others, and much more. This alone raises the stakes. What can raise the stakes even further is that, in a good number of the cases Fantl has in mind, one's supposed knowledge of a controversial proposition is being directly challenged. In other words, when someone's presents you with a sophisticated counterargument for not- p , it's not just p that's challenged, it's also whether *you know* that p .

All of this is to say that not abusing Fantl's principle involves making sure the antecedent holds before one decides to not open-mindedly engage with a relevant counterargument. This, in turn, involves making sure that one's judgements are not unduly influenced by biases or affective factors, and that one's epistemic position is both strong enough for knowledge *and* strong enough to survive being unable to expose the flaw in a relevant counterargument. Furthermore, if pragmatic encroachment is true, coming by and retaining knowledge of controversial propositions is even more difficult since (1) the content of controversial propositions often matter morally and politically, and (2) in cases where one is presented with a relevant counterargument, often part of what's being challenged is *whether* one knows that p . Both (1) and (2) raise the contextual stakes. All of this taken together suggests that even if Fantl's principle is true, utilizing it responsibly is quite difficult to do.

Here is why it matters so much that Fantl's principle not be abused. (Much of this has already been touched on, but is worth re-emphasizing.) In the first place, you should not rashly make use of epistemic principles you know are particularly prone to misuse without at least taking measures to ensure that you yourself are not misusing or abusing the principle. Doing so would be epistemically irresponsible.

Second, the danger of misusing or abusing Fantl's principle is quite high. Although some of us actually have knowledge of controversial propositions, it's also true that many of us merely think we have knowledge.³⁶ If Fantl is right, most of us are also often faced with, or at least made aware of, sophisticated counterarguments which challenge what we take ourselves to know, and which we lack the ability to reliably evaluate. And yet, because many people think they know controversial propositions even though they lack such knowledge, the danger of a large number of people misusing or abusing the principle is quite high. Contrast this risk of abuse with the risk of abuse of the enemy principle, which is only relevant to a select number of people.³⁷

Third, when it comes to controversial matters such as those Fantl focuses on, we typically care about more than just lower-level knowledge or justification. That is, not only do we want to know (or justifiably believe), say, whether God exists, that discrimination is morally impermissible, and so forth. We also want to know (or at least justifiably believe) *that* we know. This is due, at least in part, to the fact that our beliefs about controversial issues often make a moral or political difference. They determine how we vote, how we treat others, what issues we fight for (or don't), and much more. This gives us plenty of reason to make sure that we have arrived at the right conclusions with respect to these and other such controversial issues. They *matter*. Furthermore, and importantly, we need to be able to articulate to ourselves the reasons we have for believing controversial propositions.³⁸

³⁶This, of course, is meant to be consistent with the fact that a single person can know some controversial propositions, and lack knowledge of others.

³⁷Most folks are not in the business of negotiating with their enemies.

³⁸This consideration should suffice to stave off potential externalist complaints. For instance, an

A fourth and related point is that, given that these issues are often the focal points in public discussions, it also matters that we are able to articulate to others why we believe what we do, and that we are able to defend our positions when they are challenged. In other words, given that controversial issues are often morally or politically laden, we not only have to answer to ourselves, so to speak; we also have to be able to answer to others.

A final note of caution about Fantl's principle: if we allow it to guide our behavior—that is, if we regularly take ourselves to be licensed in being closed-minded to complicated arguments which challenge what we (think we) know and which we lack the expertise to evaluate—this arguably will promote political polarization. Here is why. While it is possible to closed-mindedly yet seriously engage with relevant counterarguments, it is not necessarily difficult for others to determine when you have no intention of changing your mind as a result of your exchanges with them. When this happens, one's interlocutor would reasonably think: why should I engage with you in good faith if you are not willing to do the same (especially if I think you have the wrong views)? The problem of promoting polarization is obviously even higher if Fantl's principle is regularly misused or abused. This gives us all the more reason to not abuse the principle. Indeed, it gives us reason to be very wary of using the principle, even if it is true.

What this all boils down to is that, if Fantl's principle is true, a layperson should

externalist might argue that it is not necessary to know *that* you know the controversial proposition in question in order to responsibly utilize Fantl's principle. As long as you know the controversial proposition, and you follow Fantl's principle, then you've done the epistemically responsible thing (at least, assuming Fantl's principle is true). I think such externalist claims hold little weight in this context. For even most externalists will agree that internalist considerations matter sometimes, and I am providing reasons for why they matter here.

make sure that she really knows (or justifiably believes) whatever controversial proposition is being challenged before utilizing the principle. And this is by no means easy to do. I could say more at this point about what would be involved in responsibly making use of Fantl's principle, but that would be worth it only if there is good reason to think that it's true. As I will argue in the next section, Fantl's argument fails. The reason I spent so much time discussing what would be involved in responsibly making use of his principle (if it were true) is to demonstrate both that it would be quite difficult to do and that it would matter quite a bit that it would be utilized responsibly. In short, what he is advocating for is dangerous.

1.2.3 When, and the extent to which, bracketing is licensed

Fantl's argument relies fundamentally on the premise that when one encounters an argument she lacks the expertise to reliably evaluate, and when she antecedently knows whatever proposition the argument is meant to undermine, she is epistemically licensed in bracketing any evidence the argument provides, and instead only counting as relevant the evidence which originally grounded her knowledge. This premise in turn relies on the presumed similarity between (1) the cases in which Fantl is interested and (2) other kinds of cases in which bracketing evidence is widely accepted to be epistemically permitted or required. The third criticism I have of Fantl's argument is that there is an important dissimilarity between (1) and (2).

To see why, let's remind ourselves of the details of Christensen's *Drugs* case. The reason I (the experimental subject) ought to bracket my lower-level evidence is that I have been given higher-level evidence which indicates that my rational capacities

have likely been temporarily compromised. I know both that I have been drugged, and that those who are affected by the drug cannot tell that they were. Since those affected by the drug are just as sure that they've correctly solved the puzzle—even though they likely haven't—as those who were not drugged, and since I do not know whether or not I am among the 20% unaffected, I should drastically reduce my confidence that I have properly solved the logic puzzle.

Fantl argues that if I (the experimental subject) should bracket my lower-level evidence—i.e., the evidence provided by the logic puzzle that the answer is d.—so too should laypersons bracket the lower-level evidence provided by apparently flawless, sophisticated counterarguments. There are two points to make about the discussion Fantl draws on to make his case. First, in Christensen's own discussion of bracketing, he appears to hedge in a way that Fantl does not. Christensen writes:

In accounting for the [higher-order evidence] about the drug, I must in some sense, and at least to some extent, *put aside* or *bracket* my original reasons for my answer. In a sense, I am barred from giving a certain part of my evidence its due. (2010a, 195, original emphasis)

Note the qualifiers “in some sense” and “at least to some extent.” Christensen similarly qualifies his recommendation to bracket evidence several other times throughout the paper. In contrast, Fantl—who, recall, is relying on Christensen's discussion of bracketing to support one of the key steps in his argument—regularly comments that if you are a layperson confronted with a sophisticated but misleading counterargument which challenges what you know to be true, you should be “unwilling to reduce your confidence in response to the counterargument, even if you will find it apparently flawless” (2018, xv). In other words, by “bracket,” Fantl seems to mean

something like “completely disregard or dismiss.” Of course, it’s difficult to determine exactly how much Christensen means to hedge. It’s likely that how much one should put aside evidence, or not take it into full consideration, will be contextually determined. Depending on the circumstances, for instance, “to some extent” could mean “to a large extent.” Whatever the case may be, the fact that Christensen does qualify bracketing in this way (whereas Fantl does not) should be noted, especially since Christensen takes care to include these qualifiers several times throughout his own discussion.

The second point, which is the one I really want to emphasize, is that there is an option available to laypeople confronted by a misleading but apparently flawless argument that is importantly not available to someone who has been administered a mind-altering drug as part of an experiment. In *Drugs*, I am presumably not afforded the opportunity either to have others check whether I’ve made the correct choice, or to check to see what answer(s) others—in particular, those who have not been administered the drug, and who would thus be much more reliable evaluators—have come up with. In contrast, in the kinds of cases that interest Fantl, there is nothing preventing laypeople from checking with others—in particular, experts who possess the relevant expertise to evaluate the argument—to determine whether the argument is actually flawless or only apparently flawless.

To drive home the point, let’s consider a second bracketing example which Christensen examines in a different paper. Here, Christensen is discussing the phenomenon of hypoxia, which is an impairment of cognitive ability due to lack of sufficient oxygen. As a condition that can affect those at altitudes of 10,000 feet or higher, pilots

and hikers tend to be warned about the condition. Important for our purposes is that victims of hypoxia do not initially realize their judgment has been compromised.

Pilot

You're alone, flying a small plane to a wilderness airstrip. You're considering whether you have enough fuel to make it safely to an airstrip 50 miles further away than your original destination. Checking the relevant factors, you become extremely confident that you do have sufficient fuel—you figure that you've got a fair bit more than the safety-mandated minimum. So you begin your turn toward the more distant strip. But then you notice that your altimeter reads 10,825 feet. You feel completely clear-headed and normal; however you're fully aware of the insidious effects hypoxia can have. Should you trust your recently formed confident judgment about having sufficient fuel, and continue on your path toward the more distant airstrip? (2010b, 126)

Christensen continues: “It seems clear to me that you would be grossly irrational if you were to rely on your recent reasoning about having sufficient fuel to make the more distant airstrip” (2010b, 126).

Clearly, he thinks that a significant decrease in confidence in one's ability to properly assess information is warranted.³⁹ This seems right. Notice, however, that in this example, as in *Drugs*, we are not told that the pilot has an opportunity to check with others (in particular, those *not* suffering from hypoxia, and who have the knowledge and expertise to reliably determine) whether there is enough gas in the tank to make it to the further airstrip. However, were the pilot able to check

³⁹It is worth pointing out that when Christensen does explicitly mention bracketing in this other paper, he similarly qualifies the term: “If such agents must, at least to some extent, bracket or put aside the challenged reasoning in assessing the possibility that that reasoning is flawed...” (2010b, 126).

with a knowledgeable, unaffected person as to whether there is enough gas in the tank, then it seems quite apparent that she should do this and base her belief on the testimony of that person. Similarly, in *Drugs*, if I were allowed to ask one of the people running the experiment whether or not I have solved the puzzle correctly, then my belief about whether I have correctly solved the puzzle should be based on their testimony.

While the experimental subject and pilot ought to seriously doubt—*at the time of the described situations*—their ability to determine whether they’ve reasoning properly, this does not mean that they should not seek out the feedback of those who can properly reason, if that option is available. What licenses them in bracketing their lower-level evidence is simply that, *at the time of evaluation*, there is a good chance they are cognitively compromised, and there is no other reliable and available evidence (e.g., others’ testimony) on which to base their beliefs about their lower-order evidence. As evidence that Christensen would himself be friendly to the point I’m advancing, consider another of his examples:

Reasonable Prudence

I’m a medical resident who diagnoses patients and prescribes appropriate treatment. After diagnosing a particular patient’s condition and prescribing certain medications, I’m informed by a nurse that I’ve been awake for 36 hours. Knowing what I do about people’s propensities to make cognitive errors when sleep-deprived (or perhaps even knowing my own poor diagnostic track-record under such circumstances), I reduce my confidence in my diagnosis and prescription, *pending a careful recheck of my thinking*. (2010, 186, my emphasis)

Notice that in this case, Christensen explicitly points that that a reduction in con-

confidence is required, *but pending a careful rechecking of my thinking*. This could presumably come in the form of my future self rechecking my thinking once I've had more rest, or in the form of another qualified medical professional looking over my diagnosis.

Fantl, by contrast, seems to think that a layperson is licensed in bracketing lower-level evidence provided by sophisticated and apparently flawless counterarguments, not just temporarily, or until a time when one can check with those actually in a position to evaluate the argument, but *forever*. He writes: “[E]ven if there were all the time and resources in the world, you often shouldn’t engage open-mindedly for epistemic and, sometimes, moral reasons” (2018, 129). There are thus two ways in which what Fantl suggests is quite a bit more extreme than what Christensen suggests (again, despite the fact that Fantl is basing this key move in his argument on Christensen’s discussion). First, Fantl does not hedge in his discussion of bracketing, whereas Christensen does. Second, and much more problematically, Fantl seems to think that once bracketing is licensed in a particular situation at a particular time, it is licensed forever. In contrast, Christensen does not suggest this at all. Given the last clause in *Reasonable Prudence*, it appears he would not suggest this.

Being a layperson who knows a controversial proposition does not suffice for forever bracketing a sophisticated and apparently flawless counterargument if you are currently or will later be in a position to seek out the testimony of those who are competent to judge such counterarguments (i.e., experts). I thus maintain that Fantl’s bracketing argument fails. At best, what Fantl has shown is that a layperson who is either presented with or made aware of a sophisticated counterargument ought

to in some sense, and to some extent, bracket or put aside the lower-level evidence provided by the counterargument, but only until she has the opportunity to see how others who have the expertise to evaluate such evidence have weighed in on the matter. Crucially, once we spell out what is involved in bracketing evidence, we realize that to bracket lower-level evidence in the cases under discussion is not to closed-mindedly dismiss that evidence. Instead, it is much closer to remaining agnostic toward such evidence unless or until one is in a better position (e.g., via reliance on expert testimony) to assess it.

Let's take stock. Earlier, I noted that Fantl offered two potential explanations for why a layperson does not lose knowledge simply in virtue of being unable to expose a flaw in a sophisticated counterargument. The first was that the evidence she has for the truth of the relevant proposition often *swamps* or vastly outweighs the contrary evidence the counterargument provides. The second is that she is licensed in bracketing (in Fantl's sense of forever dismissing) the evidence the argument provides. We've now seen that Fantl's bracketing argument fails. Where does this leave us with respect to whether or not a layperson epistemically ought to be closed-minded toward sophisticated counterargument challenging a proposition she knows?

In some cases, it will still hold true that the evidence she has which grounds her knowledge will so vastly outweigh the contrary evidence the argument provides that her knowledge will survive being unable to discover a flaw with the relevant counterargument. But the fact that it's possible for her knowledge to survive does not mean that she is epistemically licensed in closed-mindedly dismissing the argument. There is a difference between knowledge surviving on the one hand, and being epistemically

permitted to be closed-minded on the other. Without the success of his bracketing argument, Fantl has not established that a layperson is epistemically licensed to be closed-minded to relevant counterarguments, even when her knowledge would survive being unable to expose a flaw.

1.2.4 Fantl's definition of open-minded engagement

My final critique concerns certain aspects of Fantl's definition of open-mindedness. How we characterize open-mindedness is important as this will obviously influence whether and when we think we ought be open-minded. Recall that on his view:

You are open-minded toward an argument iff i) affective factors do not prevent you from 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. (2018, 12)

My first complaint concerns the third condition.⁴⁰ If a layperson finds herself in the position where she herself cannot locate a flaw, there are further steps she could take before a reduction in confidence is warranted. More specifically, it seems quite plausible that one counts as open-minded toward an argument if she is willing to reduce her confidence in the argument's conclusion conditional, not just on iii) spending significant time with the argument, finding the steps compelling, and being unable to expose a flaw, but also conditional on iv) doing further research and v) being unable to point out the flaws in the argument on account of one's research.

⁴⁰For ease of discussion, I will presently set conditions i) and ii) aside and assume in what follows that they are met.

Notably, Fantl seems to purposely leave these conditions—i.e., my conditions iv) and v)—out. He writes:

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 ... [Rather, I'm arguing that 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 [sic] you that you often shouldn't be open-minded when you know you're right. (2018, 130, original emphasis)

Recall that he also notes elsewhere:

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. (2018, 129)

He thus seems to think that if you spend time with the argument, think carefully and seriously about it on your own, and cannot find a flaw, then the process has ended. If you are open-minded, your confidence in the falsity of the conclusion will decrease. If you are not open-minded, you will be unwilling to reduce your confidence in the falsity of the conclusion. To see why Fantl is mistaken about this, consider the following:

Climate Science

Carla knows due to extensive research and reliance on expert testimony that climate change is real, a major threat, and largely caused by human activity. Although she is not an expert in the area, she considers

herself to be a bit of an online climate change activist. She regularly posts articles on Facebook concerning recent updates in climate science, articles which offer concrete suggestions on how individuals can reduce their carbon footprints, and so forth. Landon, a climate change denier, comments on one of these posts. He links Carla and others on the thread to a sophisticated argument containing an analysis of changing weather patterns over the past 150 years. According to the argument, there is nothing unusual or alarming about current weather trends. Carla reads the article, but is unable to make heads or tails of the data.

Fantl would say that Carla is open-minded to the counterargument iff she would be willing to significantly reduce her confidence in the falsity of the counterargument's conclusion if, after spending significant time with the argument, she cannot point to a flaw. However, the process of open-mindedly engaging with the argument does not always stop there. Being open-minded is compatible with reserving judgment until one can get assistance, as it were, from those who do have the relevant expertise. Rather than relying simply on her own ability to evaluate the argument, Carla should engage in further research to see whether experts have weighed in on this type of evidence.

At least three things could result from her research. First, she might find experts weighing in on this particular argument, in which case, she could respond to Landon by pointing him to the relevant sources. In other words, even though Carla herself cannot identify the flaw(s) with this argument, experts can and have identified the flaw(s), and Carla can rely on their testimony to know where the counterargument goes wrong. Of course, being a layperson, Carla will likely not fully understand what's wrong with the argument. Even so, she need not fully understand, since

she can rely on the expert testimony and grasp at least in a loose sense why the argument is misleading. For instance, perhaps the experts say that the data is not representative, or that the methodology employed is faulty. This seems sufficient for knowing, not just *that*, but *why* the argument is flawed.

Second, it's possible that what Carla encounters in her research are not responses to this particular argument/analysis, but rather responses from experts on analyses/arguments of this type. In other words, the argument Landon presents Carla with is a token argument of a type which climate scientists have shown to be misleading. Notably, this second possibility is more complicated than the first, for in order to accurately identify the flaw(s) in the argument, Carla would need to accurately determine that this argument is indeed a token of the type debunked by experts. Depending on the circumstances, this may be easy to do, or it may be difficult. However, assuming she can accurately determine that it is a token of a type debunked by experts, it seems like she has all she needs to respond to the challenge provided by the argument Landon references. As before, this would be sufficient for knowing, not just that, but why the argument is flawed.

The third possibility is that Carla comes up empty-handed, failing to find any experts weighing in on this particular argument or type of argument. This is not to say that none have weighed in; she just has not been able to discover any responses. It's at this point that open-mindedly engaging with the argument may require a reduction in confidence that the conclusion of the relevant counterargument is false. This brings me to my next point.

The second complaint I have about Fantl's definition of open-mindedness is his

claim that one must be willing to be *significantly* persuaded by an argument. What is meant by “significant?” Unfortunately Fantl has little to say on the matter. To repeat what was already cited above:

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 arguments 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. (2018, 13)

On the one hand, I can see why Fantl makes this move. He thinks that one should often not be open-minded at all, so there would be no need to fully articulate the details of what’s involved in being open-minded. However, in under-describing what it involves, he subtly stacks the deck against it. For it makes it seem as though one should be willing to be persuaded to a more significant degree than is often required. What I’d like to propose now is that the degree to which one should be willing to be persuaded by contrary evidence will be determined in part by the totality of one’s other relevant evidence. To appreciate the point, let’s modify our earlier example:

*Climate Science**

Carla* is rather unknowledgeable when it comes to climate change. However, the beliefs shes has about climate change happen to be true. For example, she truly believes that climate change is real, a major threat, and largely caused by human activity. One day, Carla* posts an article on Facebook which offers concrete suggestions on how individuals can reduce their carbon footprints. Landon*, a climate change denier, comments on this post. He links Carla* and others on the thread to a sophisticated argument containing an analysis of changing weather patterns over the

past 150 years. According to the argument, there is nothing unusual or alarming about current weather trends. Carla* reads the article, but is unable to make heads or tails of the data.

It seems quite clear to me that the degree to which Carla should be willing to be persuaded by the argument is much lower than the degree to which Carla* should be willing to be persuaded by the argument. Given that Carla is relatively well versed in the climate change debate, and knows certain facts about the reality of climate change, open-mindedness only requires her to be (conditionally) willing to reduce her confidence to a small degree, because the evidence the argument provides has to be weighed against all of the other contrary evidence she possesses about climate change. In other words, if open-mindedness requires a reduction in confidence, it is only to a small extent, because the other evidence she possesses swamps the evidence the argument provides.

By contrast, Carla* is not well versed in the debate. Though she has true beliefs about climate change, they do not amount to knowledge. Her beliefs are, we might suppose, more or less luckily true. Thus, she ought to reduce her confidence in her view to a greater degree than Carla should if she is unable to determine where the argument goes wrong. In short, the extent to which one ought to be willing to be persuaded by an argument is a function of how epistemically well-positioned one is prior to engaging with the argument.

My third and final complaint with Fantl's definition is that open-mindedly engaging with a counterargument might require a willingness to reduce one's confidence in *something*, but not necessarily, or not *only* in the falsity of the conclusion. Suppose

that even after conducting some research, Carla cannot find any experts weighing in on the argument Landon cited. Open-mindedness might require her to reduce, at least to some extent—though, as I have suggested, to quite a small extent—her confidence that the conclusion of the argument is false. It might also (or instead) require her to reduce her confidence in how *complete* or *well-informed* the bases of her beliefs are. In other words, she should reasonably still believe that the argument is misleading. After all, it flies in the face of all of the other good evidence she has which, taken together, grounds her knowledge that climate change is real and threatening. The reduction in confidence that is required, then, is a reduction in confidence in how well she understands the issue, and/or in how well-supported her beliefs are. This is compatible with retaining knowledge, and it still should count as an open-minded response, for she is revising her views in light of new evidence.

All of this is to say that I think Fantl’s discussion of what it means to open-mindedly engage with an argument is under-developed and misleading. As such, it positions the reader to think unfavorably of open-minded engagement. Once we spell out more fully and accurately what open-mindedness involves, Fantl’s contention that laypeople should be closed-minded to relevant counterarguments they lack the expertise to evaluate is even less plausible.

1.3 Conclusion

I’ve argued that Fantl’s position is problematic in a number of ways. The most important lessons to take away from this discussion are the following. First, his bracketing argument fails, and on account of this, so does his argument against

open-mindedness. Second, the general view he's advancing is one that should not be taken lightly, and indeed could even be quite dangerous. Third, there is more to be said about what open-minded engagement amounts to, including the extent to which one should be willing to be persuaded by evidence which challenges what one takes oneself to know. While there may be additional arguments against open-minded engagement, those advanced by Fantl are unpersuasive.

In Chapters 5 and 6, I return to the questions of what being open-minded consists in, as well as whether and how one ought to be open-minded in the types of situations Fantl discusses. The next matter of business, however, is to develop an appropriate framework for thinking about epistemic normativity more generally, especially when it comes to beliefs about morally or politically significant matters.

Chapter 2

A Moral and Social Account of Epistemic Normativity

Getting a handle on what we ought to do epistemically in morally or politically complex situations is no simple matter, and in order to arrive at a better understanding of the role open-mindedness ought to play in our lives, we will have to make a number of stops along the way. For the time being, we will need to shelve the topic of open-mindedness in particular, and turn our attention to the broader topic of epistemic normativity. The reason is this: presumably, a primary reason we care about open-mindedness is that it can promote (and is sometimes even necessary for), say, securing justified beliefs—which we in turn care about because we think justified beliefs are more likely than unjustified beliefs to be true.¹ If this is right, a

¹Indeed, there is a case to be made that even the non-epistemic (e.g., democratic) value of open-mindedness is in many ways parasitic on its epistemic value. Consider: one reason we think it is good for democratic citizens to be open-minded is that, when it comes to political deliberation—as with almost any other kind of deliberation—they are more likely to have or approximate having justified political beliefs if they are willing to engage with each other, and to sometimes change

full story of why we ought to be open-minded includes a story of why we ought to have justified beliefs.²

Of course, one might wonder why, exactly, such an account is necessary. Given that everyone more or less agrees that it is generally good to be open-minded, why try to account for *why* we ought to be open-minded? Why not instead focus on *when* and *how* we ought to be open-minded? I can think of at least two reasons for why accounting for epistemic norms—including “one ought to be open-minded”—will be illuminating. First, insofar as such claims are truly normative, rather than merely evaluative, we should be able to explain why. To echo Hilary Kornblith: “Insofar as epistemologists endorse epistemic norms, it is incumbent upon them explain the source of this normativity” (1993, 358). After all, we can make sense of epistemic *evaluations* or *recommendations* without committing ourselves to the existence of epistemic *demands*. The challenge is therefore to explain why “[they] make *claims* on us ... Or at least, *when we invoke them, we make claims on one another* (Korsgaard

their minds as a result. This is not to say, of course, that this is the only democratic value open-mindedness enjoys—for even if open-minded citizens are not more likely to have or approximate having justified political beliefs than their closed-minded counterparts, open-mindedness arguably promotes tolerance and peace.

²In my discussion of epistemic normativity, I will talk primarily in terms of securing justified beliefs. Two points about this. First, I have opted to talk in terms of *justified* beliefs rather than, say, warranted beliefs, or even knowledge. I prefer justification over warrant given that the former is better suited to capture some of the internalist considerations I wish to track. I also prefer it to knowledge, for even if we ultimately want to talk about knowledge, justification will likely be part of the picture. All that said, I am open to other varieties of epistemic evaluation other than justification, and my account is intended to be neutral in the sense that justification could be replaced with a more suitable variety of epistemic evaluation if necessary.

Second, when I talk about justified *beliefs*, I do not mean to rule out any other important doxastic states or attitudes that are relevant or worth discussing—such as disbelief, suspension of belief, or action-guiding commitments. “Belief” is simply easier to write than “whatever doxastic state is warranted by the evidence” (or something of the sort). Thus, unless otherwise noted, I will use “belief” as a placeholder for its more cumbersome alternative.

1996, 8–9, last emphasis mine).

Second, understanding the (or a primary) source of epistemic normativity puts us in a much better position to understand the extent to which epistemic norms ought to govern our behavior. Suppose I know that I ought to be open-minded. If I don't know what generates or underpins this norm, I'll be in a worse position to understand when, or how, or the extent to which, I ought to be open-minded. In other words, an answer to the question of *when* and *how* we ought to be open-minded will be incomplete without first answering the question of *why* we ought to be open-minded. In what follows, I develop an account of epistemic normativity according to which epistemic norms are underpinned by moral and social norms. One important implication of such an account is that our epistemic obligations are likely stronger and larger in number than we might ordinarily think.

This chapter begins with an overview of the three options Fantl considers for what might serve to ground an obligation to open-mindedly engage with relevant counterarguments. I also explain why he thinks all three fail to generate the obligation to open-mindedly engage. I do not take issue here with his verdicts regarding these options, given that the biggest problem I find with his discussion is that he fails to consider the full range of promising explanations for why one ought to open-mindedly engage with relevant counterarguments. The bulk of this chapter is devoted to defending the view that we cannot accurately determine what we ought to do epistemically without also considering what we ought to do morally and socially. This line thought has frequently been associated with William K. Clifford, though it has received recent attention from Sanford C. Goldberg as well. What Clifford and

Goldberg point out is that understanding ourselves as epistemic agents, including what responsibilities we bear *qua* epistemic agents, simply cannot be divorced from understanding ourselves as moral and social agents.

2.1 Fantl's proposed sources of the epistemic obligation to be open-minded

The first option Fantl considers is that the obligation to open-mindedly engage with relevant counterarguments is role-generated. Certainly, role-generated obligations exist, and it's not unreasonable to think that occupying certain roles generates an obligation to at least engage seriously with relevant counterarguments. Just as doctors have a role-generated obligation to promote the health of their patients, so too is it plausible, say, that formal commentators or participants in a public debate have a role-generated obligation to engage seriously with the arguments put forth in the relevant papers or debates (2018, 105-06). This raises the following questions. Is it plausible that a layperson has a role-generated obligation to engage open-mindedly with relevant counterarguments? If so, what role might that be?

There are perhaps a few potential answers to this question, but the one on which Fantl centers his discussion is that the layperson occupies the role of an arguer, and arguers have a role-generated obligation to engage seriously with relevant counterarguments. To see whether this counts as a plausible response, we'll need to flesh out some details. Some authors, such as Trudy Govier (1999), Jennifer Faust (2008), and Ralph Johnson (2014), have suggested that the primary purpose of argumenta-

tion is to bring about rational assent. In other words, the whole point of entering into argumentative space is to convince your interlocutors on the basis of good reasons that your position is correct, or else to be persuaded by them on the basis of the good reasons with which they provide you. Importantly, rationally persuading one's interlocutors involves seriously engaging with their objections, criticisms, and counter-evidence. After all, if one's interlocutors are persuaded even when one cannot, say, answer their objections, it quickly becomes unclear whether such persuasion is *rational* persuasion or rather persuasion by other means (2018, 107-10).

Fantl sees at least two problems with this suggestion. First, he thinks that Johnson and others have not done enough to show that rational persuasion is indeed the primary purpose of argumentation, for many people engage in argumentation simply to have their own position recognized and accepted. They aren't really open to having their own minds changed; they are only interested in changing the minds of others. There are thus at least two purposes of argumentation, broadly construed: a dialectical purpose and a rhetorical purpose, and Fantl does not see much reason to consider one purpose to be more fundamental than the other (2018, 112-16). Second, even if rational persuasion is the primary goal of argumentation, this by itself does not generate an obligation for arguers to rationally persuade their interlocutors. By way of comparison, the primary purpose of theft is to steal, but thieves are not obligated to steal. Similarly, even if rational persuasion is the primary purpose of argumentation, this does not necessarily mean that arguers are obligated to bring about rational persuasion. In short, a deeper story needs to be told about why assuming the role of arguer generates an obligation. Simply pointing to the primary

function of the practice does not suffice (2018, 110-11).

Next, Fantl considers the possibility that the obligation to open-mindedly engage is grounded in an epistemic obligation to maximize good epistemic consequences. There are at least two ways of developing this line of thought, both of which are tied to freedom of expression and the epistemic goods that can stem from it. First, we might think, along with Mill (1859/2003), that even if we have knowledge—or, in Mill’s terms, true opinions—there is a real danger in not allowing them to be subjected to scrutiny. Namely, the knowledge possessed will eventually become no better than mere prejudice, and those who have true opinions will come to lack a real understanding of the grounds for their opinions. This is what explains why open-minded engagement is necessary; simply providing a stage for people to voice their views, whether or not those views end up being heard or taken seriously, will not suffice.

Second, we might think, along with Michael Scriven (1984), that unless received opinions and values are scrutinized, and the reasons for endorsing them rehearsed, they will eventually be rebelled against, even if they are true or correct. Unlike Mill, what Scriven warns against is that received opinions and values will be abandoned, rather than that they will become prejudices. Even so, these two positions are similar in that they agree that dissenting opinions must be engaged with seriously. Regardless of what the negative outcome may be, engaging seriously with challenges to a received opinion is necessary for subjecting said opinion to real scrutiny.

A related set of points is that open-mindedly engaging with the arguments and viewpoints of disagreeing parties is likely to prevent or mitigate the effects of echo

chambers,³ which in turn can prevent or mitigate the effects of group polarization.⁴ In this way, open-mindedly engaging with the opinions of others, even when one possesses knowledge of the relevant subjects, can provide epistemic benefits to oneself, for one arguably benefits from fewer echo chambers and less group polarization (2018, 116-19). For one, echo chambers can distort one's understanding of how well supported one's own views are, as well as how well supported opposing views are. This is so even when one has knowledge of the relevant issues. Echo chambers can also give one an inaccurate impression of how widely shared one's views are; unpopular (even if true) views may seem well-received when one only considers them from inside an echo chamber. Both of these effects can contribute to group polarization, which also can have bad epistemic effects. For one, those with false or unsupported views will become more and more sure of the truth of their views, and the falsity of opposing views. On the flip side of things, even those with true or better supported views may still suffer from group polarization. As with echo chambers, group polarization can lead one to believe that one's views are better supported than they are, and that those who disagree have little to no basis for their views.

Open-mindedly engaging with the viewpoints of others can also produce good epistemic effects for others. In general, preventing echo chambers and group polarization also benefits others, not just oneself. For instance, when one engages with dissenting opinions in public, others can benefit from the exchange: they can see what reasons are marshalled in favor of the various positions under consideration, evaluate the strengths and weaknesses of each from a distance, and utilize what

³For an excellent discussion of echo chambers and epistemic bubbles, see Ngyuen (2018).

⁴For some discussion of this claim, including cited studies, see Fantl (2018, 120-21).

they've learned in order to shape their own views on the subject.

Despite the merits of these considerations, Fantl thinks they do not suffice to establish an epistemic obligation to open-mindedly engage, at least not in the cases on which he focuses his discussion. First, as he has argued, when a layperson open-mindedly engages with apparently flawless counterarguments that challenge what she knows, she places herself in a bad epistemic position by being willing to reduce her confidence in a proposition she *knows to be true* (2018, 152). In his mind, there is nothing to be gained epistemically in this type of situation. There is only the potential for epistemic harm. While it may generally be true that dissenting opinions should be engaged with so as to avoid prejudice formation or the abandonment of true received opinion, Fantl is arguing that there are exceptions—this notably being one of them. Given that she lacks the expertise to reliably evaluate them, a layperson is epistemically better off not engaging with sophisticated counterarguments open-mindedly, or perhaps even at all. Though he does not make this point himself, Fantl would presumably argue against Scriven that the danger of true received opinion being abandoned can be quite a bit higher for a layperson when she engages open-mindedly with sophisticated counterarguments than when she does not engage open-mindedly, or at all.

Likewise, he argues that while it may be true in general that we should seek to avoid echo chambers and group polarization, there are exceptions to this as well:

[G]roup 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 ... 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.
(2018, 152)

Perhaps realizing that this is a bold position to defend, Fantl does hedge a bit. For instance, he seems happy enough to grant that being in echo chambers can distort how widely received one's own views are compared to opposing views, and that this at least can be epistemically bad. Nevertheless, he insists that distancing oneself from echo chambers enough to get a more accurate picture of the prevalence of various views does not amount to open-mindedly engaging with opposing views (2018, 153). This leads him to conclude that open-minded engagement is not epistemically required or even desirable in the types of cases in which he's interested.

As for the epistemic benefits that others stand to gain by a layperson's engagement with relevant counterarguments, Fantl argues that there is no reason to suppose that such engagement needs to be open-minded. After all, a layperson need not be willing to reduce her confidence if she cannot expose a flaw in her interlocutor's arguments in order for others to see the strengths and weaknesses of the positions under consideration. It might also be argued that it is actually dangerous for a layperson to engage with misleading but apparently flawless counterarguments in public. For in the case where the layperson cannot expose a flaw in the counterargument, others might walk away with the impression that the opposing side's argument is stronger than it is, and that it is actually flawless, rather than merely apparently flawless.

As a final consideration, Fantl discusses the possibility that an obligation to open-mindedly engage with relevant counterarguments is grounded in the rights of one's interlocutors. There are several ways of spelling this out. First, we might

suppose that the right to be heard and be taken seriously stems from political rights rational agents in democratic societies enjoy. For instance, along with Rawls (1993), we might suppose that in order to hold democratic citizens to laws and policies, their opinions and views must be taken into account when forming, revising, or re-evaluating laws and policies. Thus, at least when it comes to issues that are relevant to public law and policy, one's interlocutors have a right to be seriously engaged with. As Rawls remarks, the moral duty of civility that citizens are bound by "involves a willingness to listen to others and a fairmindedness in deciding when accommodations to their views should be reasonably made" (1990, 123). Although fairmindedness is not the same as open-mindedness, an argument could be made that the duty of civility requires a certain amount of open-mindedness as well. Being willing to make accommodations when appropriate presumably involves taking requests for accommodations seriously, and being willing to change one's mind about laws or policies by incorporating accommodations where appropriate.

Amy Gutmann and Dennis Thompson (1990) develop an argument along similar lines, arguing that democratic political legitimacy demands that citizens respect and mutually acknowledge one another, especially when it comes to disagreements over policy. Otherwise, those who oppose certain policies will not really have a say in what those policies are. Instead, they will be forced to abide by them without having the chance to influence or change them. Compared to Rawls, Gutmann and Thompson are a bit more explicit advocates of open-mindedness. They write:

Mutual acknowledgement 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. (1990, 77)

Of course, it is not clear whether such considerations ground an individual obligation to open-mindedly engage with all arguments put forth by, or views endorsed by, all democratic citizens. On the one hand, one might think that individual laypersons have an obligation to open-mindedly engage with the arguments of fellow citizens, at least when one will be voting to enact policy that may negatively affect the citizens presenting the arguments. On the other hand, it could be argued that such considerations ground at most a group obligation, which at least some individual members have to fulfill, but no specific group members have to fulfill.⁵

Furthermore, even if the individual obligation for laypersons exists, this does not mean that laypersons must engage open-mindedly with every relevant counterargument. They would merely seem to ground an obligation to engage with some counterarguments in some special situations—namely, the relevant counterarguments of fellow citizens who may be affected by relevant laws or policies. Moreover, Fantl argues, these considerations do not ground an obligation to engage *open-mindedly* with such relevant counterarguments. At most, they ground an obligation, at least sometimes, to engage *seriously* with such arguments. However, engaging seriously is not necessarily to engage open-mindedly. For one can engage seriously with the

⁵Exceptions plausibly would exist, of course. For instance, individuals occupying special roles—e.g., politicians—might have an obligation to open-mindedly engage with the arguments and views of their constituents. However, the point remains that no layperson who did not occupy a special role would be bound by an individual obligation to open-mindedly engage with the arguments of fellow citizens.

arguments and views of one's fellow citizens by considering them carefully and in as unbiased a way as possible, but still be unwilling to change one's mind, even when one cannot specify what's wrong with them.

Indeed, as was discussed in the previous chapter, Fantl thinks it would be positively irrational for laypeople to engage open-mindedly with sophisticated counter-arguments they know to be flawed:

I shouldn't be silenced simply because my listeners are swayed by affective factors or because my listeners have violated 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 on me. (2018, 152)

At most, Fantl thinks that democratic considerations generate an obligation to engage seriously (rather than open-mindedly) with the views of fellow citizens.

Alternatively, it might be argued that *qua* rational agents, one's interlocutors enjoy an intrinsic right to be heard and to be taken seriously. There are at least two ways of spelling this out. First, one might argue that respecting the dignity of others demands that we take their views seriously and engage with them with respect. For instance, as Kent Greenawalt has suggested, "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" (1989, 153). Of course, Fantl's first response is to say that even if what Greenawalt suggests is true, this does not imply that each of us has to listen to what others say *open-mindedly*. For one can take the views of others seriously without open-mindedly engaging with them. More to the

point, Fantl thinks that even if what Greenawalt suggests is true, this does not mean that every individual layperson is bound by an obligation to engage with the views of every other person they encounter.

Second, following Miranda Fricker (2007), one might argue that the obligation to engage open-mindedly with relevant counterarguments is grounded in the obligation to promote testimonial justice. On Fricker's view, testimonial injustice occurs when a speaker's testimony is not given the credit it deserves due to identity prejudices harbored by the speaker's audience (2007, 28). More specifically, on account of a speaker's perceived or actual social identities (e.g., being a woman, being a person of color), coupled with the hearer's negative prejudices regarding those same social identities, a speaker's testimony can be perceived by her audience to be less credible than it otherwise would be perceived. For instance, when a black woman's testimony that x is perceived to be less credible than a white man's same testimony that x on account of identity prejudice harbored by her audience, the black woman has suffered testimonial injustice. For she is harmed in her capacity as a knower—as a *giver* of knowledge.

Fantl sees two drawbacks to this approach. First, even if the obligation to open-mindedly engage with relevant counterarguments is grounded in the obligation to promote testimonial justice—or epistemic justice more broadly—a layperson would only be obligated to open-mindedly engage with relevant counterarguments when it is plausible that one's interlocutor is likely to be a victim of epistemic injustice. Second, he thinks the obligation to open-mindedly engage with those likely to be victims of epistemic injustice is most plausibly a group obligation, rather than an

individual obligation (2018, 126).

To sum up: Fantl considers three separate options for what might ground a layperson’s obligation to open-mindedly engage with relevant counterarguments. The obligation might stem from the nature of argumentation itself; it might stem from the epistemic benefits of engaging in argumentation; or it might stem from the rights possessed by the layperson’s interlocutors. Fantl finds all of these options wanting. As I mentioned above, I will not presently take issue with any of Fantl’s verdicts regarding these three options.⁶ This is because, even if Fantl’s verdicts regarding these options happen to be correct, I think that he has failed to consider the full range of promising explanations for why one ought to open-mindedly engage with relevant counterarguments.

2.2 Stage-setting remarks

Before laying out the social and moral account of epistemic normativity, it would be useful to specify exactly what I aim to establish. There are three primary goals. First I am to account for the normative nature of claims such as “one ought to have justified beliefs,” or “one ought to be open-minded.” Recall, to say that such claims are normative is to say that they “[they] make *claims* on us ... Or at least, *when we invoke them, we make claims on one another* (Korsgaard 1996, 8–9, last emphasis mine).

For ease of discussion, I will use the term “obligation” to serve as a placeholder

⁶This is not to say that I agree with him. For instance, I think that his claim that no individuals have an epistemic obligation either to prevent (or at least not promote) testimonial justice requires more defense.

for “epistemic oughts” or “legitimate epistemic claims we make on one another.” Importantly, I don’t want to get too hung up on whether or not these sorts of claims always amount to obligations, or whether they are even best understood as obligations. To understand why, it will be useful to consider Walter Sinnott-Armstrong’s thoughts on moral obligation.⁷ According to Sinnott-Armstrong, moral obligations are the kinds of things (whether they be moral demands, moral expectations, or some similar thing) which, when not met, subject one to at least one of the following kinds of sanction: legal sanctions, social sanctions, or private sanctions (i.e., self-directed sanctions made within the privacy of one’s own mind). As he points out, depending on where we draw the boundary lines, being appropriately subjected to a certain kind of sanction may correspond to a failure to meet one’s moral obligation, or it may not.

For instance, according to Mill (1861/1998), being legitimately subjected to any kind of the three mentioned sanctions means one has violated a moral obligation, whereas according to Gert (1998), it is only if one is liable to public sanction (i.e., legal or social) that one has violated a moral obligation. Sinnott-Armstrong goes on to point out that whether Gert or Mill is correct about how to use the term “obligation” is at bottom merely a verbal dispute. What matters is not whether being liable to private sanctions is indicative that one has failed a moral obligation, but rather that this kind of sanction is appropriate for certain kinds of failure—and in general, that certain kinds of sanctions are appropriate for certain kinds of failure.

What I borrow from Sinnott-Armstrong is simply the idea that whether or not

⁷See Sinnott-Armstrong (2005, 193-97).

we use the term “obligation” to refer to the kinds of epistemic claims we make on one another is not what’s most important. Instead, what matters for my purposes is that we do, in fact, make claims on one another to act in certain ways epistemically, that these claims are legitimate, and that certain kinds of sanctions are appropriate when we fail to live up to them. To repeat, then: “epistemic obligation” is simply a placeholder for “legitimate epistemic claims we make on one another,” and it could be replaced by a more suitable or desirable term if necessary.

Second, I aim to establish that epistemic obligations are binding on all epistemic agents, rather than merely some epistemic agents. As the discussion in §1 illustrated, someone might grant that epistemic obligations (including the epistemic obligation to open-mindedly engage) do in fact exist, but that they are binding only on those who occupy certain roles. For example, they might grant that members of a scientific community ought to behave in certain ways, epistemically speaking, given that members of this community rely on each other, work together to promote certain epistemic goods (e.g., knowledge or understanding of certain subjects), and more. These considerations plausibly generate role-specific obligations to not promote or spread false or misleading information, draw hasty conclusions, form beliefs on the basis of insufficient evidence, and so forth. Granting this, however, does not commit one to the existence of obligations that are universally binding regardless of what role one occupies, or what goals one has. For (such a person might say) if someone in the scientific community decided she did not want to be bound by these obligations, she could simply give them up by exiting the profession.

Finally, I aim to establish that epistemic obligations range over, so to speak,

all beliefs. In other words, epistemic obligations concern not just beliefs about so-called important matters, but also beliefs about trivial matters. The distinction between agent-universal obligations and belief-universal obligations is needed given that some who may be willing to grant that epistemic obligations are agent-universal may nevertheless be hesitant to grant that they are belief-universal.⁸

Summing up, the present goals are:

1. To explain what gives epistemic obligations their normative force;
2. To establish that epistemic obligations are *agent*-universal; and
3. To establish that epistemic obligations are *belief*-universal.

2.3 Epistemic obligations are agent-universal

In what follows, I offer two separate arguments. The first is that certain epistemic obligations are underpinned by moral obligations, and the second is that certain epistemic obligations are underpinned by social obligations. To be clear, I do not mean to endorse a clean distinction between the moral and the social, for I think it can be very difficult in practice (and sometimes even in theory) to tease these dimensions apart. There are two reasons for offering separate arguments. First, although the moral and social may often overlap (and are perhaps sometimes inextricably intertwined), they do not always. Second, moral normative pressure is often (though perhaps not always) stronger than merely social normative pressure. For this reason, I think a stronger case for the moral-social account of epistemic normativity can be

⁸See, for example, Goldberg (Forthcoming.) Goldberg's argument will be discussed in more detail shortly.

made by starting with the moral argument, and then offering the social argument by extension.

As I noted before, my view is heavily inspired by both William K. Clifford's (1877/1886) *The Ethics of Belief*, as well as Sanford C. Goldberg's recent work in the area. So, much of what I say closely mirrors elements of their respective accounts.⁹ In what follows, I will flag what I am borrowing from them, as well as where I am either expanding on, or departing from, either of their views.

2.3.1 The moral underpinnings of epistemic normativity

First, let's consider why, at least sometimes, we ought to make sure our beliefs are justified, since doing so is necessary for fulfilling relevant moral obligations.

The Moral Argument

1. We have moral obligations.
2. Meeting our moral obligations involves (1) not causing moral harm, and (2) within reason, minimizing the risk of causing moral harm.
3. Unjustified beliefs are among the notable factors that can cause, or significantly increase the risk of, moral harm.
4. Thus, we have morally grounded epistemic obligations to (as far as we're able) make it more likely that beliefs relevant to the fulfillment of our moral obligations are justified.¹⁰

⁹Broadly, the Moral Argument is inspired by Clifford, and the Social Argument is inspired by Goldberg.

¹⁰Two points. First, here, and elsewhere, when I say that (a) one has an obligation to make it more likely that one's beliefs justified, I mean that (b) one ought to perform belief-influencing actions—such as attending to evidence, being careful and thorough in one's investigation, and so forth—in order to make it more likely that one's beliefs are justified. I put the point in terms of (a) simply because it's less cumbersome than its more accurate counterpart, (b). Second, I add

I take Premise 1 of the Moral Argument more or less for granted. As a reminder, what I mean to take for granted is simply that, *qua* moral agents, there are things we morally ought to do, and that moral agents make legitimate claims on one another to behave (or refrain from behaving) in certain ways. Putting the point in terms of obligations is convenient, but not necessary. Assuming this is safe to take for granted, the work lies in motivating Premises 2 and 3. I will consider the plausibility of each in turn.

Let's begin by considering a specific moral obligation—for instance, the moral obligation drivers have to not hit pedestrians or other cars on the road. Part of what's involved in meeting this obligation is not, in fact, hitting others while one is driving. Additionally, however, it involves not putting others at significant *risk* of being hit. This is why drivers should not text behind the wheel, for even if they do not actually cause any harm, they nevertheless put pedestrians, other drivers, and themselves at significant risk of harm. Similarly, they should not drive if they have good reason to think that the lug nuts on their tires are loose, for the risk of harming themselves or others increases when the risk of one of their tires flying off on the road increases.

Similar points can be made in reference to other moral obligations. Consider, for instance, the moral obligations parents have to their children. Among other things, parents ought to make sure (as far as they are able) that their children have a good education, that they are healthy, that they are being loved and emotionally

the qualification “as far as we're able” in order to account for the fact that a person might do what is reasonably within her power to ensure that her beliefs about a given topic *T* are justified; and yet (perhaps through bad luck), she may nevertheless fail to have justified beliefs about *T*. I submit that such a person has not failed to meet her epistemic obligation.

supported, and more. When it comes to health, parents ought to make sure that their children are getting enough nutrients and sleep, that they're taking their vitamins, that they're not overeating or getting into tide pods, and so forth. Moreover, parents ought to make sure that they are not putting their children at *risk* of bad health. For instance, they should make sure, not only that they are getting enough nutrients, but also that they are getting the vaccinations necessary to prevent serious illness in the future.

Likewise, parents ought to make sure both that their children are getting a good education, and that they are not at risk of being intellectually hindered. In more detail, parents ought to make sure, as far as they are able, that their children are in an educational environment in which they can thrive, that their children are engaging in practices that don't hamper intellectual growth, and that their children are cultivating good intellectual character traits. Importantly, to meet their obligations to their children, not only must parents make sure that their children have access to intellectual goods, and that they are currently engaging in practices that promote intellectual growth, but they must also make sure that they're not allowing their children to be at risk of being intellectually hindered. For example, if parents know that too much gaming hampers creativity and growth over time, parents ought to limit the number of hours per week their children playing video games. Likewise, if they have good reason to think that their children are hanging out with friends who will encourage the development of bad intellectual habits, they should intervene before these bad habits do in fact develop. And so forth.

Notice that in all of these cases, in order to fulfill one's moral obligations, it is not

enough that one does not actually cause moral harm—e.g., one does not in fact hit pedestrians or other drivers, nor does one in fact hamper the physical or intellectual growth or health of one’s children. Additionally, meeting one’s obligations involves ensuring to the best of one’s ability that one has not significantly *risked* causing moral harm. This point generalizes. All else being equal, in order to fulfill one’s moral obligations, one must neither cause moral harm to others, nor act in such a way that one poses a significant risk of causing moral harm to others.¹¹

Of course, risk comes in degrees. And while it is true that meeting one’s moral obligations involves not putting others at significant risk of moral harm, it is not thereby true that meeting one’s obligations involves eliminating the risk of moral harm to others. For instance, drivers should not text behind the wheel because this poses a significant risk of harm to pedestrians and other drivers. However, drivers need not *eliminate* the risk of hitting others, for this would require not driving at all. Likewise, parents should not allow their children to eat as much candy as they’d like, as this can significantly put their health at risk. However, parents need not eliminate the risk of bad health—for instance, by placing their children on an extremely regimented diet and never allowing them to eat candy at all. The point is simply that while meeting our moral obligations involves not causing or significantly risking moral harm, it does not involve entirely eliminating the risk of moral harm. Although it is difficult to determine where, exactly, the threshold of significance lies, this does not affect the truth of Premise 2, which is all I aim to establish here.

Let us move now to Premise 3:

¹¹For ease of discussion, I focus on what we morally owe to others. In doing so, I do not mean to suggest that we lack self-regarding moral obligations or, say, moral obligations to the environment.

Unjustified beliefs are among the notable factors that can cause, or significantly increase the risk of, moral harm.

The rationale for this premise falls quite naturally out of the rationale for the previous premise. For unjustified beliefs are simply among the things that can easily result in actual or significantly risked moral harm. To see why, let us return to the example of driving. A few ways that a driver can fail to live up to her moral obligation to pedestrians and other drivers are: texting while driving, driving while intoxicated, or driving when she knows her lug nuts are likely loose. These are all things she should not do, and others might reasonably expect her to avoid these situations so that she doesn't put them at undue risk of harm.

Alternatively, our driver may fail to live up to her moral obligation by driving with unwarranted confidence. Suppose, for instance, that she has recently heard from reliable sources that certain parts used in her model of car are being recalled. A number of cars with this part have experienced sudden electrical failure, resulting in accidents varying in degrees of severity. Previous to this news, our driver had no reason to believe her car was unsafe. Now, of course, she does. Nevertheless, she decides not to take her car into a dealer to have it checked out. Without any good reason, she believes that her car will not be among the unlucky lot and continues to drive as she did before. If our driver fails to meet her moral obligation to others when texting, then she also fails to meet her moral obligation when she drives with unwarranted confidence. This is so even if she does not end up actually causing harm to others, for she has significantly risked causing them harm.

For further support, it may be useful to bring in some of Clifford's observations

and arguments from *The Ethics of Belief*. At the beginning of his famous essay, he discusses two cases in which the unjustified beliefs of a person or group of people result in actual or potential harm to others. In the first case (which, of course, serves as an inspiration for the driving examples) a shipowner's unjustified belief that his ship is seaworthy results in the death of immigrants who set sail on his ship, which eventually sinks. In the second, a group of religious authorities are wrongfully accused of manipulating the public and indoctrinating children. Even though the truth comes out that the accusations were false and no further harm comes to them, Clifford argues that the accused were nevertheless wrong. In both cases, then, unjustified false beliefs resulted in moral wrongs—in the first case, the immigrants died at sea, and in the second case, the religious authorities were slandered by the false accusations.

Importantly, Clifford insists that even if these beliefs had turned out to be true, the guilty parties still would have failed to meet their moral obligations. Suppose, for instance, the ship owner's belief about the seaworthiness of his ship turned out to be true. Although he would not have been justified in believing that his ship was seaworthy, it could have turned out that his ship was in fact capable of making that voyage and many others afterward. Even so, it still would have been wrong for the shipowner to so believe. Although his belief turned out to be true, it was merely *luckily* true. He should not have so recklessly formed or held his belief, especially when it mattered so much (both practically and morally) that his belief was true.

As before, this point generalizes. In order to meet our moral obligations, we have to make sure that beliefs relevant to those obligations are justified. For having

unjustified beliefs is just one of many ways that can result in actual or risked moral harm. Thus, whenever one's beliefs are relevant to one's moral obligations, one ought to make sure that those beliefs are justified. To put the point in the terms introduced in §2.1: we all make moral claims on each other, whether implicitly or explicitly, to behave in certain ways. Meeting these moral demands often involves meeting certain epistemic demands. Thus, epistemic oughts—at least epistemic oughts which concern beliefs relevant to our moral obligations—are binding, rather than merely suggestive, for they are rooted in moral oughts, which are themselves binding. We can't legitimately opt out of complying with these epistemic demands any more than we can legitimately opt out of complying with moral demands.¹²

For similar reasons, we can see that epistemic oughts are also agent-universal, since the epistemic agents we are concerned with (human beings) are also moral agents. In more detail, all of us are bound by a morally grounded epistemic obligation to make sure that beliefs relevant to the fulfillment of our moral obligations are justified. This obligation is not a role-specific obligation, for being a moral-cum-epistemic agent is what grounds this obligation, and being a moral-cum-epistemic agent is not role-specific; we cannot “opt out” in any real sense of opting out. Thus, morally grounded epistemic obligations exist, and they are agent-universal. This point holds even though what specific moral obligations each person is bound by may differ, depending on their means, socioeconomic status, power (or lack thereof), and other relevant factors.

¹²Of course, we can refuse to comply. But in doing so, others can sanction us for failing to meet the legitimate claims they make on us.

2.3.2 The social underpinnings of epistemic normativity

We've made some progress, though we still have a ways to go. After all, many beliefs are not relevant to the fulfillment of our moral obligations in the ways just described.

This is why we also need:

The Social Argument

1. We have social obligations.
2. Meeting our social obligations involves (1) not causing social harm, and (2) within reason, minimizing the risk of causing social harm.
3. Unjustified beliefs are among the notable factors that can cause, or significantly increase the risk of, social harm.
4. Thus, we have socially grounded epistemic obligations to (as far as we're able) make it more likely that beliefs relevant to the fulfillment of our social obligations are justified.

The first thing to note is the Social Argument's similarity to the Moral Argument. Given this similarity, the Social Argument will be in one sense easy to motivate, since many of the points made above can simply be repeated. In another sense, however, this argument will be somewhat difficult to motivate since, unlike moral obligations, the existence and extent of social obligations—or, at least, something like the social analog of moral obligations—cannot easily be taken for granted. The difficult task, then, is to demonstrate that there are social obligations—at least in the sense that social agents make legitimate claims on one another, whether implicitly or explicitly, to behave (or refrain from behaving) in certain ways, and that the agents on whom the claims are made ought to live up to them. Luckily, Goldberg (Forthcoming) has

undertaken the task of establishing this point. Rather than starting from scratch, then, I can utilize the work he has already done, and expand on his view to suit my purposes.¹³

To motivate the view, Goldberg invokes Karen Jones' (2017) work on trust, as it is especially germane to the issue of what social agents can legitimately expect of one another. As Jones observes, each and every one of us is limited cognitively, temporally, emotionally, and likely in a host of other ways. Given these limitations, we require help from others to carry out the vast majority of our projects, especially if they are to be carried out with any real degree of success. This cooperation and co-dependency is essential for life as we know it, for it allows us to “*extend our agency*” in ways we otherwise could not (2017, 101, my emphasis). She writes:

As finite social creatures other agents are a particular salient source of risk to us, but they also provide a remedy for our finitude, for together we can do what we cannot do alone, whether because the activity itself is necessarily a shared one (waltzing), or because it requires divisions of time, labour, and skill (*most of the activities in our daily lives*). (99-100, my emphasis)

¹³Indeed, my Social Argument shares a great deal in common with Goldberg's “Basic Argument,” stated as follows:

Premise 1 For all subjects S and T, if T is (morally or socially) entitled to expect S to ϕ , then S ought to ϕ .

Premise 2 There are many cases in which given epistemic subjects S and T, T is (morally or socially) entitled to expect S to conform to evaluative epistemic standards in belief-formation and belief-management.

Conclusion There are many cases in which epistemic subjects ought to conform to evaluative epistemic standards in belief-formation and belief-management. (Forthcoming, 4)

Trust when met with trustworthiness allows us to enhance the effectiveness of our agency. As finite social agents we have a pressing interest in being able to do this. However, the interest that we have in being able to extend our agency in this way is not a distinctly moral interest. It is about making agency *effective* by drawing on the agency of others rather than about making it *good*. Extended agential power is an ends-independent value. Whatever our ends, we want to be able to recruit the agency of others: even counter-ethical projects require recruits. (101-102; original emphasis)¹⁴

In short: we necessarily rely on other agents in order carry out projects, to succeed, to learn and grow, and to generally get on with our lives.¹⁵ Because of this, Goldberg argues, social agents are entitled to normative expectations regarding other's cooperativity, at least when these expectations are rooted in a legitimate rule-governed social practice. Let's unpack this.

To begin, I should clarify some of Goldberg's terminology. By "normative expectation" he means "an attitude whereby we hold others responsible...[and where such] expectations can be appropriate even when one has ample evidence that [others will not meet them]" (Forthcoming, 5). By "entitlement" he means "permission" (Forthcoming, 5). Thus, when he says, for example, that one is entitled to expect *X* from other social agents, he means that certain standing norms are in place, which correspond to a legitimate rule-governed social practice, and which thereby license or

¹⁴As Goldberg points out, this last point is important, for if considerations having to do with trust underpin some of our epistemic obligations, and since even immoral projects require trust, then the considerations underpinning some of our epistemic obligations may not always be moral, but may sometimes simply social.

¹⁵Whereas Jones puts the point in terms of the necessity of trusting others, Goldberg opts to put the point in terms of the necessity of being able to rely on others for their cooperation. At least for present purposes, I take these to be two ways of expressing the same point.

permit one to expect others to X . In other words, “the social entitlement to expect derives from the practice itself, as any legitimate practice entitles its participants to hold one another to the rules themselves” (Forthcoming, 6). By a “legitimate practice” he means “an ongoing and recognized practice, [whose] standards are widely acknowledged, and [where] there have been no serious questions as to the propriety of either the practice or its standards” (2017, 2867). A final point to highlight: given that the entitlement to expect certain behaviors from other participants derives from the practice itself, he also points out “that you cannot question the propriety of the expectations without also calling into question the legitimacy of the practice” (2017, 2867).

Some examples will help to tie everything together. We’ll begin by borrowing one of Goldberg’s own.¹⁶ Exceptions aside (e.g., during weeks when a holiday is observed), I am entitled to expect that my garbage and recycling will be picked up from my curb on Tuesday mornings. What entitles me to this expectation is the fact that I participate in a rule-governed, community-wide practice of garbage pick-up, where the standards of the practice are widely acknowledged, and where there have been no serious questions as to the propriety of the practice. So long as I have done what is legitimately expected of me—for example, I bring my bins out to the curb before the trucks are scheduled to arrive, I do not place non-recyclable items in the recycling bin, and so forth—I am entitled to expect that my trash and recycling will be removed Tuesday morning or early afternoon. For having one’s bins emptied on a regular schedule is standard practice, not just in my community, but

¹⁶See Goldberg (2017, 2869).

in communities all over. Since it is a legitimate standard practice, participating in it entitles me to certain expectations, including expectations of others participating in the same practice. Importantly, the expectation that my bins be emptied on time is normative. Even if garbage pick-up has been unreasonably delayed in the past, I still am entitled to expect that it happen on time. For the expectation is grounded in the practice-generated norms, rather than past evidence concerning the timeliness of garbage collection.

Furthermore, if anyone questioned the propriety of my expectation, they would thereby be calling into question the legitimacy of the practice itself. Suppose I express frustration to a friend when my bins do not get emptied until Thursday morning. My friend suggests that I should temper my expectations; after all, sometimes people simply fall behind schedule. A reasonable reply to make, especially if this is not the first time my bins have not been emptied on time, would be to point out that having my bins emptied on a regular schedule is the service for which I'm paying. In other words, a primary reason I participate in the practice is precisely because it entitles me to certain expectations regarding others' behavior (namely, that those employed by the waste company will pick up my trash and recycling in a timely manner). If I am not entitled to those expectations, that counts as a strike against participating at all.

Generalizing the point, we could say that cooperation is legitimately expected when agents are engaged in legitimate transactional exchanges with one another. Importantly, such exchanges are entirely commonplace, and it is not difficult to bring to mind a host of transactional exchanges we engage in with others on a daily

basis wherein cooperation is appropriately expected: purchasing gas on the way to work, coffee during one's morning break, lunch in the afternoon, groceries on the way home from work, and so on. In each of these and similar cases, involved parties are entitled to expect that the other parties will cooperate, as such transactional exchanges are commonplace, have more or less accepted rules, and there have been no serious questions as to their propriety.¹⁷ For example, both buyer and seller are entitled to expect that the other will not try to cheat them out of money; that the other will not waste their time (e.g., by pausing the transaction to complete a drawn out, unnecessary, and unrelated task); that the other will answer any reasonable questions there may be related to the transaction; and so forth.

Transactional exchanges in which cooperation is expected are ubiquitous. But this is just one example of a kind of social practice in which cooperation is legitimately expected. There are also a wide variety of situations in which we regularly find ourselves occupying a certain role, and occupying that role quite obviously requires cooperation. Such situations include small efforts (such as planning a meal with a partner), big efforts (such as being a contributing member of a scientific community), and everything in between. As before, in calling into question the propriety of the

¹⁷To be fair, critics of consumerism and capitalism have questioned the propriety of such practices in various ways. While such criticism may be correct, they do not seem to fully undermine the legitimacy of transactional practices. For the most part, participants in such day-to-day transactional exchanges have not seriously questioned the propriety of such practices, but rather go on with (consumerist) life as they know it. This of course is not to say that there should not be improvements in the consumerist sphere—CEOs of giant corporations should not be making massive salaries while lower-level employees are barely making enough to get by, workers should not be exploited in other ways, and so forth. But to say that there should be improvements on consumerist practices in general (or even in particular instances) does not necessarily or obviously undermine the legitimacy of transactional exchanges in which participants are actively willing to participate, and there have not been any serious complaints about the legitimacy of those particular exchanges.

expectations participants in such a practice are entitled to, one thereby calls into question the legitimacy of the practice itself. For instance, being a contributing member of a scientific community entitles one to expect (among other things) that other participants in the scientific community be honest and thorough, both in the production of and reporting of their research. Calling into question the propriety of such an expectation thereby calls into question the legitimacy of the practice itself. (Why would scientists collaborate with each other in pursuit of epistemic goods like truth and understanding if they were not entitled to expect that participating parties are actually working to promote these epistemic goods?)

Alternatively, planning a meal with my partner entitles me to expect, among other things, that my partner will not sabotage the plan—for instance, by throwing the cooked meal in the trash just before we're about to eat. As before, calling into question the propriety of such an expectation thereby calls into question the legitimacy of the practice itself. Put differently: I would not participate in the practice of planning meals with my partner if I were not entitled to expect that my partner not sabotage our plans. Nor, would my partner plan meals with me if he were not entitled to expect the same.

As with transactional social practices, social practices in which one occupies a certain role is just another example of a type of ubiquitous social practice in which cooperation is legitimacy expected. And there are others. In some cases, it may be that local norms are in place (e.g., cultural or familial) that require agents be cooperative. For instance, given widely acknowledged norms concerning friendship (and perhaps norms concerning being considerate to others more generally), I am

entitled to expect that my friend is ready when I stop by her place to pick her up.¹⁸ If I have to wait 20 minutes in her driveway, sanctioning her in some way upon her arrival—for instance, by calling her out on her tardiness and requesting that she be punctual when I pick her up in the future—would certainly be appropriate, at least assuming she didn't have a good reason to be late.

This expectation is normative, for I am entitled to expect her to be punctual even when past experience tells me that she will not be. In other words, the expectation is grounded in standing norms rather than evidence I have regarding her behavior. What gives me the right to this expectation is that it derives from legitimate, rule-governed social practices (being in a friendship, or interacting and coordinating with others more generally) in which my friend and I are both willing participants. Furthermore, she cannot call into question the legitimacy of my expectation that she be punctual without calling into question (at least to some extent) the legitimacy of the practice and corresponding norms themselves.^{19,20}

Engaging in conversation with others also generally licenses expectations to others' cooperation. (Here Gricean maxims obviously come to mind.) If a conversational partner frustrates my efforts—for example, by bullshitting, or by completely ignoring my questions—that conversational partner has violated a legitimate normative expectation I had regarding their cooperativity. As before, in calling into question

¹⁸It may be that being a friend also counts as occupying a role. It is perfectly fine with me if these categories overlap. All I need to illustrate is that in all of these examples, cooperation for one reason or another is legitimately expected of others.

¹⁹Consider: On the extreme side of things, if my friend takes advantage of our friendship in this and other ways, this may count as a strike against our friendship. On the less extreme side of things, my friend being late, especially if it happens on a regular basis, may count against my agreeing to give her rides in the future.

²⁰See Gilbert (2006, 2013) for relevant discussion.

the legitimacy of the expectation of cooperativity, one thereby calls into question the legitimacy of the practice or standards of conversation itself.

To sum up what Jones and Goldberg are getting at: engaging in social practices in which cooperation is legitimately expected is an integral part of our daily lives. We collectively buy into—and, in many cases, are simply born into—a vast number of practices that demand cooperativity and reciprocity. And we continue to participate in them because in doing so we are thereby able to enhance the effectiveness of our agency. Furthermore, whenever one participates in a legitimate practice that licenses certain expectations of cooperativity, one is expected to be cooperative, just as one can expect that others be cooperative. In effect, then, it is a kind of contractual reciprocity that gives rise to our legitimate entitlements that others be cooperative in their dealings with us, and to others' entitlements that we be cooperative in our dealings with them. Our mutual dependency on one another—that is, the fact that we cannot help *but* rely upon one another in order to get on with our lives—substantiates the idea that there are social obligations.²¹

Now, for Premise 2:

Meeting our social obligations involves neither (1) actually frustrating nor (2) significantly risking frustrating the efforts of those with whom we are cooperatively participating in a legitimate social practice.

Goldberg submits that part of what is involved in being genuinely cooperative—rather than, say, superficially cooperative—is being competent to participate in the

²¹As a reminder, the term “social obligation,” is simply meant to reflect the normative claims social agents make on one another, whether implicitly or explicitly, to be cooperative when they are engaged in legitimate social practices such as those discussed above.

relevant activities. Returning to Jones' quote from above: "Trust *when met with trustworthiness* allows us to enhance the effectiveness of our agency. As finite social agents we have a pressing interest in being able to do this" (2017, 101-102, my emphasis). Insofar as it is in all of our interests to enhance the effectiveness of our agency by relying on others—and it is—then we have a legitimate expectation that others be trustworthy, just as they have an expectation that we be trustworthy. Putting the point more strongly, enhancing the effectiveness of our agency *requires* that those in whom we put our trust are trustworthy. Furthermore, as Goldberg points out, part of what's involved in being a trustworthy participant is being a *competent* participant. For this reason, it is legitimate to expect not just that others be cooperative, but also that others be competent to cooperate; otherwise, relying on others would not be an effective way of extending our agency.

Once this is pointed out, establishing Premise 2 of the Social Argument should be quite easy. For being competent to participate means that one neither (1) actually frustrates nor (2) significantly risks frustrating the efforts of those with whom one is cooperating. The point about risk is important, for even when one does not, as a matter of fact, fail to be cooperative, whenever one puts others at significant risk of frustrating their efforts in a situation in which cooperation is legitimately expected, one has failed to cooperate competently. Hence, one has violated the trust of fellow participants. As was noted above, then, the kinds of considerations that were offered in support of Premise 2 of the Moral Argument can simply be reapplied here. Still, we'll illustrate with an example.

Suppose a friend and I are planning a trip to Europe together, and in order to

make the burden equal, we divvy out certain responsibilities. She might, for instance, be in charge of arranging lodging, whereas I am in charge of budgeting for travel, meals, and other expenses. In agreeing to plan this trip together, we are thereby entitled to expect the other to carry out her various tasks with relative competence. If one of us fails to carry out our assigned tasks competently, then that person has failed to live up to a legitimate expectation of cooperation. Suppose I devote only five minutes to setting up a budget, and that competently setting up a budget would require a minimum of several hours. Given that I have failed to competently carry out my task, I have failed to live up to my friend's legitimate expectation of cooperation.

Importantly, I have failed to live up to this expectation even if the trip ends up running smoothly. For in incompetently setting up the budget, I significantly risked the success of our trip. Just as a motorist fails to meet her moral obligation to ensure the safety of others when she recklessly texts on the road, even when she does not actually cause harm to others, so too do I fail to meet my social obligation to my friend when I fail to devote the requisite time to budgeting, even if the trip ends up running smoothly despite my incompetence. While this is just one example, it is not difficult to see how the point generalizes to any case in which cooperation—hence relevant competency—can be expected. To meet one's social obligations one must neither actually frustrate nor significantly risk frustrating the efforts of those with whom one is cooperatively participating in a legitimate social practice. Social agents are entitled to expect this much of one another.²²

²²As before, it is not necessary that social agents eliminate all risk of social harm (i.e., the frustration of cooperative efforts), as this would be likely amount to not engaging in cooperative efforts at all. Even though it can be difficult in practice to determine where the threshold of significance lies, this does not affect the truth of Premise 2, which is what I aim to establish here.

Finally, Premise 3 of the Social Argument:

Unjustified beliefs are among the notable factors that can cause, or significantly increase the risk of, social harm.

As with Premise 2, this premise can be easily established, for it is just a social extension of the Cliffordian line of thought. More specifically, failing to be doxastically competent is just one way to fail to meet legitimate expectations of cooperation. We can illustrate by modifying the budgeting example. Suppose I am not competent to budget, not because I have not devoted adequate time to budgeting, but rather because I have uninformed or misinformed beliefs about how to budget. Just as I violate my friend's legitimate expectation that I competently carry out my assigned tasks when I do not devote enough time to budgeting, so too do I violate her expectation by having unjustified beliefs about how to budget properly. Furthermore, even if I do not actually sabotage our trip on account of my doxastic incompetence, I have still failed to live up to my friend's legitimate expectations. Suppose that by a stroke of pure luck I managed to set a good budget for the trip, despite having unjustified beliefs about how to budget well. Even though our collective plans had not been frustrated, I still violated the trust of my friend for I had unduly risked the success of our cooperative efforts.

Thus, Premises 1-3 of the Social Argument. To begin, we have brought to light that social agents regularly make claims on one another, whether implicitly or explicitly, to be cooperative when they participate in legitimate social practices. Second, meeting these so-described social obligations involves neither (1) actually frustrating nor (2) significantly risking frustrating the efforts of those with whom we are coop-

eratively participating in a legitimate social practice. Third, in order to meet these social obligations, social agents ought to make sure that beliefs relevant to these social obligations are justified.

The Shared Epistemic Resources Argument

At this point, one might worry that we haven't done enough to motivate the claim that epistemic obligations rooted in social obligations are agent-universal. For all that has been said, the worry goes, these obligations may simply be role-specific obligations that one can easily opt out of complying with by simply exiting the role. For instance, if I do not want to be beholden to my friend's legitimate expectation that I be doxastically competent to budget our trip, I can simply call off the trip or not agree to it in the first place.

I will now present a supplementary argument which I hope will accomplish two things. First, it will lend further support to Jones' claim that cooperative activities make up most of our daily lives, and that social agents cannot easily or realistically opt out of these roles on account of our lives being fundamentally social. Second, it will establish that our socially grounded epistemic obligations are not limited to having justified beliefs about topics relevant to specific cooperative activities. Additionally, social agents have socially grounded epistemic obligations to not damage or threaten the vehicles which make cooperative efforts possible in the first place. In appreciating the breadth of our socially grounded epistemic obligations, it will even clearer why they are not obligations we can realistically opt out of complying with.

Here is the argument:

The Shared Epistemic Resources Argument

1. Epistemic agents necessarily rely on others in order to both (1) develop into mature epistemic agents and (2) have successes (epistemic and otherwise).
2. Developing into mature epistemic agents and continuing to have successes would not be possible without having drawn on (and continuing to draw on) a shared wealth of epistemic resources which has been preserved and added on to over the ages.
3. Given that these epistemic resources are shared goods, epistemic agents who benefit from, and contribute to, these resources are thereby entitled to expect that others not threaten them (for example, by having unjustified beliefs, by spreading misinformation, or by being otherwise doxastically incompetent or insincere).

This argument is inspired by Clifford's following remarks in *The Ethics of Belief*:

[N]o one man's belief is in any case a private matter which concerns himself alone. Our lives are guided by that general conception of the course of things which has been created by society for social purposes. *Our words, our phrases, our forms and processes and modes of thought, are common property*, fashioned and perfected from age to age; an heirloom which every succeeding generation inherits as a precious deposit and a sacred trust to be handled on to the next one, not unchanged but enlarged and purified, with some clear marks of its proper handiwork. Into this, for good or ill, is woven every belief of every man who has speech of his fellows. An awful privilege, *and an awful responsibility*, that we should help to create the world in which posterity will live...

Belief, that sacred faculty which prompts the decisions of our will, and knits into harmonious working all the compacted energies of our being, is not for ourselves, but for humanity. (1877/1886, 342-43 my emphasis)

I find the idea that beliefs—and more generally, our ways of thinking and communicating—are common property (or perhaps something like share epistemic resources) which others can legitimately expect to be protected and perhaps improved upon to be quite insightful. We'll call this *the shared resources claim*. I'll now offer a defense of it, though there are a couple of departures from Clifford to highlight. First, I do not want to attach too much weight to the idea that shared epistemic resources count as property. All I need to establish is that there is something like shared epistemic resources—including epistemic goods (e.g., justified belief and knowledge), ways of sharing epistemic goods (e.g., testimony), and so forth—that (1) we all benefit from and/or rely on, and (2) which we all have a claim on, at least insofar as we can expect that others not abuse or threaten them. Second, Clifford seems to think that the shared resources claim is rooted in moral considerations—that is, he thinks that we morally owe it to each other to preserve our shared epistemic resources. In contrast, and in line with Goldberg and Jones, I don't think that the explanatory story is necessarily moral all the way down. While moral considerations certainly do play a large role, some of the heavy lifting can arguably be done with merely social considerations.

In my view, there are two sub-claims that need to be defended in order to lend credence to the shared resources claim:

Backward-looking claim: Any epistemic successes we have enjoyed (such as knowing or understanding) to date—indeed any successes we have enjoyed more generally—are not ours alone, but are shared by those who made such successes possible.

Forward-looking claim: In continuing to be a part of the epistemic

community which made our past epistemic successes possible, we owe it to others in that community to not threaten the shared epistemic resources.

As I see it, establishing the backward-looking claim is in effect to establish Premises 1 and 2, and establishing the forward-looking claim is in effect to establish Premise 3. I will thus aim to establish these two claims now.

One of the things that Clifford is pointing out is that when we succeed in believing the truth, or in having justified beliefs, this success is often (perhaps always) largely due to the efforts of other epistemic agents in our generation, and of those in generations previous. Our beliefs about science, medicine, morality, religion—indeed, our beliefs about *everything*—are shaped and informed by those who raised us, by those we encounter in our day-to-day lives, and, importantly, by those who lived before us and tested out ideas, theories, and methodologies to determine which are reasonable, which are true, which are worth pursuing, and which are not. Of course, this is not to say that we cannot ever count epistemic achievements—such as knowing or coming to understand—as personal successes, for epistemic achievements are often due to our own faculties, both perceptual and intellectual. It is to say, however, that epistemic achievements would not be possible if not for the help, guidance, and positive influence of others.

If this sounds too strong a claim, reflecting on some considerations might help illustrate the point. First, consider the basic point that our ability to use and understand language—which is arguably a prerequisite for non-perceptual epistemic achievements, such as propositional knowledge or understanding—is possible thanks to the efforts of, and exposure to, those who helped raise us. In other words, our very

capacity for sophisticated epistemic achievements would not exist if not for the help and guidance of others. As for the particular epistemic achievements themselves—for instance, propositional knowledge or understanding—much of our success is due in no small part to others as well. Knowledge and understanding about a great deal of subjects—history, science, politics, mathematics, morality, and much more—is passed on to us by our parents and educators, who expose us to these subjects and who teach us how to navigate the various kinds of knowledge. Credit is likewise due to the historians, scientists, politicians, ethicists, and more, who made contributions in their respective fields; to those who peer-reviewed the contributions to make sure they in fact promote epistemic goods; and to all those who made the success of each of these jobs possible—for instance, those who raised and educated the scientists and ethicists, and who made it possible for these scientists and ethicists in turn to make valuable contributions of their own. Similarly, we owe many of our epistemic successes our contemporaries: to those we interact with in day-to-day life, who share knowledge with us via testimony and who force us to question and articulate our beliefs and assumptions; to news reporters, who share knowledge with us about our communities and the world that we otherwise would not have access to; and much, much more.

In addition to propositional knowledge and understanding, others are also to be credited for much of our knowledge-how: knowledge of how to balance bills; knowledge of how to discern good evidence from bad; knowledge of how to be in a healthy relationship; knowledge of how to ride a bike; knowledge of how to be successful in one's profession; knowledge of how to employ different methodologies in

one's profession, research, or studies; knowledge of how to change a tire; knowledge of how to be a good citizen; knowledge of how to be virtuous, morally and epistemically; and more. This list could go on and on. Yet even this short list illustrates that the kinds of knowledge-how which others have passed on to us, and continue to educate us on, are varied: there is practical knowledge-how, moral knowledge-how, epistemic knowledge-how, and political knowledge-how. As with propositional knowledge and understanding, our knowledge of how to do things is due in no small part to the teaching and positive influence of others.

Even when we are able to get it right on our own, so to speak—for example, when we innovate, make discoveries, or make novel observations—that we are able to get it right on our own is not entirely our own doing. Any discoveries or innovations we produce are made possible in no small part by the wealth of shared epistemic resources which have been built up and refined over the course of human history, and which have been passed down to us by those who have raised and educated us, and by the epistemic community more broadly. If not for this wealth of epistemic goods from which we benefit, those discoveries and innovations would not have been possible. Even some of the most radical discoveries—such as those that result in scientific paradigm shifts—would not have come about if those who made the discoveries had not benefited from the epistemic resources shared by those who came before them and by their contemporaries.

As a final consideration, notice that when things go wrong, epistemically speaking—such as when one is dogmatic, has a good deal of false or unjustified beliefs, has a difficult time telling good evidence from bad, or is otherwise epistemically vicious—

this often due to a *lack* of good epistemic influence, and/or exposure to bad epistemic influence. For instance, those who are gullible are often that way, at least in part, because they have not been properly educated on how to be discerning. Alternatively, consider those who are raised in secluded communities where access to epistemic goods is limited, or worse, epistemic harms are prevalent. Such people often lack epistemic goods and are intellectually hindered or vicious *because* of their intellectual upbringing and epistemic environment. This is not to say epistemic vices are only the result of the negative influences of others. People raised in healthy epistemic environments are perfectly capable of becoming gullible or otherwise epistemically vicious to varying degrees on their own. On the flip side of things, those who are raised in non-ideal epistemic environments, or those who have a fair amount of exposure to bad epistemic influences can develop intellectual virtues despite their circumstances. The point is simply that when we get things right, epistemically speaking, our successes are often not possible without the positive epistemic influence and contributions of others. Even those who are able to develop intellectual virtues in spite of their circumstances typically (arguably always) owe their epistemic successes to some degree or other to others.

Any given individual epistemic agent simply could not have developed into a mature epistemic agent if she had not been given access to (likely a wealth) of epistemic resources. The epistemic resources and epistemic care from which she has benefited were provided by a network of agents in her immediate epistemic community, and each of those agents in turn was provided with epistemic resources and care by those in their immediate epistemic communities, and so on. In short, any

epistemic successes we enjoy to date are not ours alone, but are shared by those who made such successes possible. Furthermore, any successes more generally which *rely on* epistemic successes—which is a great many, if not all of our successes—are shared by those who made epistemic our successes possible. After all, it is quite difficult to succeed in anything without having the requisite propositional knowledge, knowledge how, requisite understanding, and the like.

These points, I hope, should not sound too radical. In fact, I hope that they sound rather obvious. Even so, they are worth bringing to the forefront of the conversation, for they force us to acknowledge just how fundamentally interpersonal our worlds are. In general, having successes and getting by in this world involves having true beliefs, knowledge, and/or understanding about the activities or projects in which one is involved. Thus, not only are the epistemic resources shared, but any successes which the resources have made possible are also shared by the vast network of epistemic agents who have either provided us with epistemic goods or made the production of those goods possible.

This should suffice to establish the backward-looking claim, which brings us to the defense of the forward-looking claim:

In continuing to be a part of the epistemic community which made our past epistemic successes possible, we owe it to others in that epistemic community to not threaten the shared epistemic resources.

Arguably, the backward-looking claim is less controversial than the forward-looking claim. Whereas the former simply amounts to the idea that epistemic resources are shared, the latter amounts to the idea that, in virtue of these resources

being shared, we *owe* it to others to not threaten or abuse them. In other words, the backward-looking claim is descriptive, whereas the forward-looking claim is normative. Nevertheless, I think that reflecting on some further considerations will vindicate the normative claim as well.

Earlier, I discussed Jones' point that most of the activities in our daily lives—from driving, to buying groceries or gas, to having conversations, to picking up one's kids from daycare, to doing one's job or advancing in one's career—require cooperation. Whether we like it or not, we are all intimately caught up in a host of social practices in which cooperation is, to some degree or other, legitimately expected. And succeeding in our efforts often, if not nearly always, requires having justified beliefs about topics relevant to those efforts. Above, it was demonstrated that this need to rely on, and cooperate with, each other as a means to extend our individual and collective agency generates an entitlement to expect that others have justified beliefs, at least when those beliefs are relevant to the success of our cooperative efforts.

The point I want to make now is somewhat different, but nevertheless related in a way that will soon become clear. The point is that is also the case that epistemic agents are entitled to expect that other epistemic agents generally not damage or threaten shared epistemic resources. As a reminder, the epistemic resources in question include not just the knowledge and understanding that have been passed down through history, but also (to borrow Clifford's terminology) our processes and forms of thought, which might include, *inter alia*, our reliance on testimony (both peer testimony and expert testimony), our methods of producing and testing ideas,

our ability to ask critical questions and challenge assumptions (i.e., engage in the space of reasons), and more.

The reason epistemic agents are entitled to expect that others not threaten shared epistemic resources is quite simple: in virtue of such epistemic resources being shared or public, it is legitimate to expect that individuals do not threaten them. Just as members of a physical community are entitled to expect that individuals will not damage, abuse, or threaten public goods or resources—for instance, community members generally are entitled to expect others will not damage or vandalize public parks, roadways, or monuments,²³ that those who hold public office will not use their position for private gain, and so forth—so too are members of epistemic communities entitled to expect that shared epistemic resources will not be damaged or threatened. That this is true should be especially apparent in this age of misinformation within which we unfortunately find ourselves.

We can begin by considering a set of particularly conspicuous examples. At the time of writing, the world is struggling both to understand and mitigate the effects of the the novel coronavirus, COVID-19, which first appeared in Wuhan, China in December of 2019. In the very early stages of the spread of the virus, a great deal of information and misinformation was shared via social media. Some of this misinformation was perhaps to be expected. As a novel virus, there was much to learn about it, and certain predictions or judgments about its nature were bound to change as experts came to better understand it. For instance, proclamations that the

²³I use the term “generally” to allow for certain exceptions—for instance, in cases where damage to, or removal of, public monuments is done in the name of advancing equality and promoting civil rights.

the virus is essentially like the flu were shown to be false when it became apparent that the virus is 10x deadlier than the seasonal flu, and that it spreads at a higher rate.²⁴ The sharing of such misinformation, at least when very little was known about the virus, was perhaps excusable. Other misinformation, however, was almost certainly shared in bad faith, which was decidedly inexcusable.

For instance, at a campaign rally on February 28th, 2020 in North Charleston, SC, President Trump claimed that the Democratic Party was “politicizing” the coronavirus, and openly suggested that it was “the new democratic hoax.”²⁵ Elsewhere, he significantly downplayed the threat of the virus, as well as the extent to which the United States was prepared to face such a threat. However, just a couple of weeks later on March 17th, 2020, Trump claimed: “I felt [the coronavirus] was a pandemic long before it was called a pandemic.”²⁶ Not surprisingly, many are quite angered by Trump’s full-on reversal in narrative, some even going so far as to say that Trump is gas-lighting the nation.²⁷ There are at least two effects of this narrative shift to find troubling—one direct, and the other indirect or downstream. The direct effect is that Trump is unabashedly polluting the epistemic environment by spreading misinformation²⁸ and making baseless claims²⁹ in a time where governments the world

²⁴See Huang (2020).

²⁵Note: he did not directly assert that it was a hoax, but instead suggested that “one of [his] people” claimed this. See C-Span (2020).

²⁶See Rogers (2020).

²⁷See Lemon (2020).

²⁸In addition to being entirely dishonest about what his stance has been on the coronavirus all along, he falsely claimed that tests to determine whether one has the virus would be available to anyone who wanted one. They were not. See Weiland (2020).

²⁹To take one example, Trump has suggested that we may be a “matter of months away” from developing a vaccine—a suggestion medical and infectious disease experts have been quick to correct. See Vazquez (2020).

over are enacting states of emergency in response to the virus, and where having reliable and accurate information about the nature and threat of the virus can be a matter of life or death. The indirect or downstream effect is that Trump polluting the epistemic environment now, especially in a time of crisis when having accurate information really matters, can have rippling effects on public trust in authority well after this particular crisis is dealt with. This is especially so since Trump seemingly sees no need to take responsibility for the false or inaccurate claims he has made.

Of course, this is far from the only time that Trump has polluted the epistemic environment. Given the sheer amount of pollution he has produced in the form of misinformation or baseless claims,³⁰ and given the platform he is afforded by holding the office of the President (which can only serve to amplify the effects of the epistemic pollution he produces), Trump can easily be considered to have a rather large “epistemic carbon footprint.” However, he is far from the only guilty party. The last several years have proven to be a time in which the spread of misinformation online has become so commonplace that it is now often referred to the “age of fake news.” Not only is there plenty of misinformation or mal-information³¹ to go around, but also, social media users often share such information without much thought.³²

As was the case with Trump in particular, many are deeply troubled in general by the fact that we do live in the age of fake news, and for the same reasons. For one,

³⁰See Kessler et al. (2019).

³¹I borrow the term “mal-information” from Kujawski (2019). The term is defined as “Information that is based on reality, [but which is] used to inflict harm on a person, organization or country (e.g. someone using a picture of a dead child refugee (with no context) in an effort to ignite hatred of a particular ethnic group they are against.”

³²For an excellent discussion of anti-truth trends in political discourse, including the role social media platforms play in propping these trends up, see Lynch (2019).

the false, misleading, or baseless claims that are shared on social media platforms and by word of mouth can make tangible differences to people's lives. Suppose I go to my primary care physician to be tested for COVID-19 after hearing that tests are readily available to anyone who wants one. Given that testing kits are not in fact readily available, my trip will almost certainly be a waste of time. Moreover, if I do not already have the virus, I increase my risk of exposing myself to it since it is not unlikely that someone who is infected has recently visited my doctor's office.

Alternatively, suppose my friend and I had planned to vote in the Illinois Primary on March 17, 2020. On the 17th, my friend contacts me to say that because of the outbreak, the Illinois Primary elections will be postponed until further notice. Not having time to check whether this is true, I trust what my friend has told me and carry on with my day. As it turns out, the elections were not in fact postponed, and my friend and I both miss our opportunity to vote. While I certainly shoulder some of the blame for missing the vote (I should have verified what my friend told me when I had the chance, especially knowing how much misinformation is being spread in relation to the pandemic), my friend is to blame as well. She too should have double-checked whether the election was postponed before relaying the message to me, for the success of our cooperative effort depended on having accurate information. This is, of course, one of Clifford's points—a point which Goldberg (and Jones) emphasize, as we have seen. Namely, we are entitled to expect that others have justified beliefs, at least when those beliefs are relevant to their moral obligations or our joint cooperative efforts.

There is, however, a second point that Clifford is making—a point which relates to

his remarks about common epistemic property. Namely, I am entitled to expect that others not produce or spread misleading, false, or baseless claims, even when these claims do not directly affect me. For the more such claims are spread, and the more the epistemic environment is polluted in general, the more epistemic practices—such as being able to rely on experts, or being able to rely on the testimony of others (e.g., one’s peers)—are threatened. In other words, even if particular efforts of mine are not frustrated by particular epistemic pollutants $A_1 - A_N$, I still am entitled to expect that others not contribute to epistemic pollution—for example, by producing or uncritically repeating $A_1 - A_N$. For the more the epistemic environment is polluted in general, the more the very vehicles which make successful cooperative possible in the first place (i.e., trust in experts, reliance on others, and so forth) are jeopardized.

In short, there are at least two reasons the current state of affairs is so deeply troubling. The first is that being in an epistemic environment in which misleading information or misinformation is regularly shared, and critical evaluation is seemingly all too often an afterthought, can quite easily lead to the frustration of one’s efforts—practical, political, moral, or otherwise. The second is that even if certain epistemic pollutants do not directly threaten one’s particular efforts or projects, polluted epistemic environments nevertheless threaten the shared epistemic resources upon which the success of *any* of those efforts and projects rely.³³

In this way, the argument that epistemic resources are shared goods, and that epistemic agents are entitled to expect of each other that the resources not be threat-

³³Of course, I have focused on specific examples of the sharing of information in the age of fake news—and in particular, the sharing of (mis)information in the wake of the COVID-19 outbreak. This is simply because these examples are especially relevant. However the main point I’m making can be easily generalized.

ened, goes hand in hand with the argument that epistemic agents are entitled to expect of each other that they be competent to participate in cooperative social practices. For, as agents who so heavily rely on one another both to develop and extend our agency, we need not only are entitled to expect each other to be competent (including doxastically competent) in our cooperative efforts, but we also are entitled to expect that the mechanisms which make cooperative efforts successful in the first place not be jeopardized.

Above, we considered the worry that socially grounded epistemic obligations may not be agent-universal, given that there's a sense in which they are all role-specific. In other words, the worry is that we don't *have* to comply with them, for we could just exit the particular social roles we occupy, thereby releasing ourselves from these obligations. This worry has been shown to be misguided on two fronts. For one, even if it is true that all of our socially grounded epistemic obligations are in a sense role-specific, we can't well opt out of all of these roles, short of going off the grid. To repeat: we rely on each other fundamentally in order to get on in our lives. Giving up all of these social roles would amount to radically and fundamentally changing the ways in which we get on in the world. Thus, anytime we are engaged in a legitimate social practice—and we regularly and almost necessarily do engage in such practices—other participants are entitled to expect that we be competent, including doxastically competent. Second, aside from any particular social roles we occupy, others are generally entitled to expect that we do not act in ways that will abuse or threaten shared epistemic resources. For all that has been said, however, it is unclear whether these epistemic obligations are belief-universal, for the arguments

given so far pertain only to beliefs which are relevant to the fulfillment of our moral or social obligations.

2.4 The belief-universality of epistemic obligations

There are at least two potential routes for establishing the belief-universality of epistemic obligations. The first can be construed as broadly consequential, the second as virtue-theoretic. The first option is to argue that any unjustified belief—even if it is not obviously relevant to any moral or social obligation—is such that it poses a risk of moral or social harm. Since Goldberg explicitly explores this line of thought, I’ll focus on what he has to say.^{34,35} Borrowing the notion of an externality from economics (i.e., an indirect side effect of some economic activity that affects third parties), Goldberg introduces the notion of a negative epistemic externality:

[A]n epistemic externality obtains when the following conditions are met: (i) a subject S is in doxastic state-type Δ ; (ii) that S is in Δ has (practical and epistemic) consequences on subjects other than S; (iii) among those affected are some subjects who aren’t presently in Δ . (Forthcoming, 13)³⁶

Goldberg’s suggested (though, again, not endorsed) route for establishing the belief-universality of epistemic obligations is to establish that “[t]here are negative externalities involved *whenever* a subject S fails to meet epistemic standards in having

³⁴See also Grimm (2009).

³⁵To be clear, Goldberg is neutral on the question of whether epistemic obligations are belief-universal. He merely proposes an argument for why we *might* think that epistemic obligations range over all beliefs.

³⁶Just as ordinary externalities can be either positive or negative, so too can epistemic externalities. Goldberg is only interested in the negative variety.

doxastic attitude Δ " (Forthcoming, 16, original emphasis). In other words, all doxastic states are such that, if they are poorly formed, they put others at undue risk of social or moral harm. For every doxastic state is such that it *could* be relevant to the fulfillment of one's moral or social obligations. Since no epistemic agent can know in advance which beliefs will be relevant to her moral or social obligations, she ought to make sure that all of her beliefs are justified.³⁷ Others are entitled to expect as much.

I do not think this route for establishing the belief-universality of epistemic obligations is very plausible. The problem is that, while others are entitled to expect that, within reason, we minimize risks when it comes to meeting our morally grounded and socially grounded epistemic obligations, others are not necessarily entitled to expect that we eliminate risks altogether. Earlier, it was pointed out that pedestrians are entitled to expect that drivers minimize the risk of harming them—for example, by not driving drunk. However, pedestrians are not entitled to expect that drivers eliminate the risk of harm altogether by not driving at all. The same reasoning applies here. In other words, I don't think that the fact that any particular poorly-formed belief *might* end up causing moral or practical harm, regardless of how (un)likely it is that the poorly-formed belief will cause harm, suffices to establish the belief-universality of epistemic obligations.

For what it's worth, I think Goldberg is right to point out that it's often difficult to tell in advance which beliefs will be relevant to the fulfillment of our moral and social

³⁷Actually, Goldberg prefers to remain neutral on whether the best way of capturing the point is in terms of having justified beliefs. Instead, he says that others are entitled to expect that our beliefs conform to the correct evaluative standards of epistemology, whatever those standards happen to be.

obligations. However, I don't think that this consideration rescues the plausibility of the consequentialist argument. For even if there may be many cases in which it will be difficult to tell whether a belief will be relevant to the fulfillment of our moral and social obligations, it also isn't difficult to come up with cases in which a poorly-formed belief poses such an extremely small risk to others' practical or moral interests such that the risk hardly seems worth considering. This in turn puts pressure on the thought that mere risk (no matter how small) of practical or moral harm suffices to establish the belief-universality of our epistemic obligations. To take the most obvious kind of case, it is not difficult to come up with examples in which poorly-formed beliefs about trivialities will likely not put others at (any real) risk of practical or moral harm, for these particular beliefs will likely not affect others, either directly or indirectly. Consider:

Hiker

I am hiking deep in the woods and notice two trees off the beaten path that are of roughly the same size and shape. Observing them, I notice that one of them simply looks older to me. I form the belief that this tree is older than the other, despite the fact that I have no experience dating trees. I am, in fact, wrong. The tree I believe to be younger is actually quite older.

It strikes me as quite unlikely that this belief will ever become relevant to any of my moral or social obligations. Moreover, the mere fact that this particular belief *might* end up being relevant to a non-epistemic obligation does not by itself mean that others are entitled to expect that this belief be justified. Others are entitled to expect that the risk of moral and social harm be minimized. They are not entitled

to expect that such risks be eliminated entirely.

Luckily, I think there is another route to take that allows us to circumvent this problem and see quite readily why epistemic obligations are not just agent-universal, but also belief-universal. Rather than focusing on the risks associated with particular poorly-formed *beliefs*, we should focus on the risks of forming or maintaining bad belief-forming *habits*. In other words, the explanation for why epistemic obligations are belief-universal is that, whenever one forms a belief poorly, she indirectly puts others at significant risk of moral or social harm. This is so even when a particular poorly-formed belief will likely never become relevant to her moral or social obligations. The reason is that those same belief-forming habits are at work, so to speak, when one forms beliefs about topics that are obviously relevant to one's social or moral obligations. The very act of forming beliefs on insufficient grounds—even when the belief concerns so-called trivial matters—shapes us into the kind of person who will form beliefs on insufficient grounds in general. And others are entitled to expect that we not be that kind of person.

Clifford would appear to agree that habit-formation is explanatorily central. Here again, he is worth quoting at length:

Nor is it that truly a belief at all which has not some influence upon the actions of him who holds it. He who truly believes that which prompts him to an action has looked upon the action to lust after it, he has committed it already in his heart. If a belief is not realized immediately in open deeds, it is stored up for the guidance of the future. It goes to make a part of that aggregate of beliefs which is the link between sensation and action at every moment of all our lives, and which is so organized and compacted together that no part of it can be isolated from

the rest, but every new addition modifies the structure of the whole. No real belief, however trifling and fragmentary it may seem, is ever truly insignificant; it prepares us to receive more of its like, confirms those which resembled it before, and weakens others; and so gradually it lays a stealthy train in our inmost thoughts, which may someday explode into overt action, and leave its stamp upon our *character* for ever. (1877/1886, 342, my emphasis)

[I]t is not the risk only which has to be considered; for a bad action is always bad at the time when it is done, no matter what happens afterwards. Every time we let ourselves believe for unworthy reasons, we weaken our powers of self-control, of doubting, of judicially and fairly weighing evidence. We all suffer severely enough from the maintenance and support of false beliefs and the fatally wrong actions which they lead to, and the evil born when one such belief is entertained is great and wide. But a greater and wider evil arises when the credulous character is maintained and supported, when a habit of believing for unworthy reasons is fostered and made permanent. (1877/1886, my emphasis, 344-45)

His emphasis on the importance of developing and maintaining good intellectual character is crucial. Avoiding the negative effects that bad belief-formation has on our intellectual character is equally as important as avoiding the (actual or merely risked) negative effects that come from having particular unjustified beliefs.

Returning to *Hiker*, consider how a hiking buddy might respond if she comes to know that I have this belief about the age of the trees. When she asks why I was so confident about their relative ages, I report that the one just *looked* older to me. We might naturally expect my friend to criticize my belief, assuming she knows that I have no good reason to be confident about such judgments. Suppose

I push back and say: “What does it matter? It isn’t going to make any difference to anyone! I couldn’t even take you to those trees if I wanted to; I don’t know how to get back there.” Despite my excuses, it would be entirely reasonable for her to sanction me. That this belief likely won’t make a difference to anyone is beside the point. For if I formed *that* belief rashly, it’s likely I will form *other* beliefs rashly. And many of those other beliefs actually will make a difference to others. Sure, it might have mattered more that I formed this belief on the basis of good evidence if it would likely have made a direct social or moral difference to others. For instance, if someone were to need to take my word about the age of those trees, it would have mattered more than I knew what I was talking about.³⁸ But even if this particular belief makes no difference to others, it’s not the belief itself that counts in this case. It’s the fact that *whenever* I form beliefs rashly, it makes it more likely that I’ll form other beliefs rashly—beliefs which in fact will matter to others.

Generalizing the point, once it matters that *some* beliefs (indeed, a great deal of our beliefs) are justified—namely, beliefs relevant to our moral and social obligations—it matters that *all* beliefs are justified. The reason for this is that we cannot promote good belief-forming habits when the beliefs in question clearly matter, but be lax when it comes to forming beliefs about relatively trivial matters. After all, as proponents of doxastic involuntarism have long been pointing out, we don’t have control over belief formation in that way. That is, we cannot decide that we will form our beliefs well in cases A - N, but not in cases A₁ - N₁. We simply don’t have that kind of doxastic control. At best, we can indirectly control or influence what kinds of beliefs

³⁸Notice, too, that the sanction I would reasonably be subjected to would be harsher in the case when my belief would likely make a direct moral or social difference to others’.

we have by trying to develop and maintain good belief-forming habits—for example by spending time with evidence relevant to our beliefs, trying to resist making hasty conclusions, trying to counteract or lessen the influence of our implicit biases, and so forth. In short, the only way we can ensure that our beliefs are justified when it matters, so to speak, is by cultivating good epistemic habits. But, if that’s so, then the content of the belief is in a sense irrelevant to the importance of making sure the belief is justified.³⁹

More precisely: Epistemic agents ought to have justified beliefs when those beliefs are relevant to their social and moral obligations. In order to make sure that these beliefs are justified, epistemic agents need to cultivate and maintain good belief-forming habits. Since the same belief-forming habits are employed when (a) one forms a belief that will likely be relevant to one’s moral or social obligations and (b) when one forms a belief that will likely be irrelevant to one’s moral or social obligations, epistemic agents ought to try to ensure that all of their beliefs are justified, regardless of the content of those beliefs happen to be. In a strange way, then, it appears that doxastic involuntarism—which is typically understood to be a problem for views according to which we have epistemic obligations⁴⁰—actually serves to establish their belief-universal nature.

In the next chapter, I argue that there are two important implications of treating epistemic normativity as being intimately tied up with moral and social normativity.

³⁹This is not entirely accurate. As I will argued in the next chapter, the more significant an issue is non-epistemically, the more important it is that beliefs relevant to that issue are justified. But saying this is consistent with saying that even beliefs about insignificant issues ought to be justified.

⁴⁰For some discussion, see Alston (1988), Feldman (1988), and Peels (2017).

Chapter 3

Further Implications of the Moral-Social Account

I will now argue for two further claims, both of which are natural extensions of the view developed in the previous chapter. First, non-epistemic considerations make it such that we sometimes ought to *form* justified beliefs about matters of social, political, or moral importance. Second, the importance of meeting our epistemic obligations is roughly proportional to the strength of the corresponding non-epistemic obligations. I will also say a few words about what meeting these epistemic obligations looks like in practice. My hope is that with these further arguments on the table, it will be even clearer why articulating the social and moral underpinnings of epistemic normativity is important. For in better understanding the intimate relationship between epistemic and non-epistemic norms, we will be in a much better position to understand the nature and extent of our epistemic obligations.

3.1 Obligations to have or form beliefs

It is not uncommon for epistemologists to argue that we have epistemic obligations. For instance, Stapleford (2013) argues that we have an imperfect epistemic obligation to deliberate evidence relevant to our beliefs.¹ Peels (2017) argues that “one has an epistemic intellectual obligation to perform belief-influencing action *A* if one believes that it would be epistemically bad not to perform *A* in that it leads to or maintains certain false beliefs” (106). Johnson (2018) and Lackey (2018a) argue that we have epistemic obligations to voice dissent in certain contexts. Gunn (2018) argues that we have an epistemic obligation to respect and promote the epistemic agency of others. And the list goes on. For the most part, however, epistemologists agree that we do not have any epistemic obligations to *have* or *form* beliefs about any given topic. In other words, some epistemic obligations may constrain *how* we ought to form beliefs—we ought to not be hasty in forming beliefs; we ought to consider a wide variety of evidence relevant to topics under consideration; and so forth. We may also have epistemic obligations to act in ways that promote healthier epistemic environments—e.g., by voicing dissent, or by acting in ways that allow others (and ourselves) to develop and express full epistemic agency. However, most would deny that there are agent-universal epistemic obligations to form justified beliefs about certain topics.²

¹Actually, Stapleford is not so much endorsing that we have such an obligation. He is merely interested in this paper to show that it’s feasible we do have positive epistemic duties or obligations. His dialectical opponent is Nelson (2010), who argues that epistemic obligations can only be negative.

²Emphasis should be placed here on “agent-universal,” for it is not all that uncommon to hold that there are *role-specific* epistemic obligations to form beliefs about certain matters. Peels (2017), for instance, claims that students and professionals have epistemic obligations to seek out certain

One notable exception to this is Gideon Rosen, who claims:

We are under an array of standing obligations to inform ourselves about matters relevant to the moral permissibility of our conduct: to look around, to reflect, to seek advice, and so on. The content of this obligation varies massively from case to case. But no matter the details, whenever a person acts badly from some sort of ignorance, the question will always arise whether has discharged these ‘epistemic’ obligations. (2002, 63)

Another notable exception—which should not come as much of a surprise by this point—is Goldberg (2017), who argues that there are cases in which a person *should have known* that p , despite the fact that she was not in a position to know it given her current evidence. What makes it the case that she should have known p , despite not being in a position to know it, is that others are entitled to expect that she know p , and where what licenses this right to expect is either participation in a rule-governed social practice, or else the basic moral and epistemic expectations we have of one another. Goldberg calls these “the ‘epistemically interesting’ cases of the *should have known* phenomenon” (2017, 2864). They can be contrasted to should-have-known cases in which a person actually has evidence that p , but simply fails to recognize or pay heed to p . To appreciate the difference, consider:

Suppose that you have some rare condition whose symptoms are bizarre but extremely uncomfortable. Here there would be no presumption that

evidence. For instance, detectives have a role-specific epistemic obligation to seek out evidence relevant to their cases; students have a role-specific epistemic obligation to understand class material; women who decide to get pregnant have a role-specific obligation to acquire information about pregnancy; and so forth. Interestingly, Peels does suggest briefly that there are agent-universal “moral intellectual obligation[s]” to *not* form beliefs about particular topics (2017, 101). For instance, all else being equal, one ought not spy on others.

your family doctor would know about the condition or its treatment. (She may never have encountered that condition before, or even read about it in medical school.) But if she fails to recognize that, given the bizarre and extreme symptoms she's observing, she ought to consult with the relevant medical experts, then you might well say that she "should have known" to do so. (Whether this case is an epistemically interesting one depends on what she knew about the conditions under which she ought to consult with others.) (Goldberg 2017, 2865)

The arguments I present in this chapter are similar to Goldberg's in many respects, though I take our goals to be distinct. Whereas he aims to establish that there are epistemically interesting cases of the *should have known* phenomenon, my goal is to establish that meeting our non-epistemic obligations sometimes requires forming (justified) beliefs about topics which are relevant to our non-epistemic obligations, but about which we currently lack beliefs. Some cases I discuss may qualify as epistemically interesting by Goldberg's lights, and others not. While determining whether the cases I discuss are epistemically interesting is a worthwhile pursuit, it is nevertheless orthogonal to my current goal.

What I aim to establish is similar in spirit to Rosen's remarks as well. However there are some differences. First, whereas Rosen is strictly interested in how we ought be informed on matters of moral importance, I also discuss the importance of being informed on matters of social and political importance. Second, and perhaps more notably, whereas Rosen seems to take for granted that we ought to inform ourselves about matters relevant to our moral obligations and simply leaves it at that, I offer an explanation for why this is so.

Here is why I think such an explanation is needed. On the one hand, it may

seem obvious that you ought to be informed on matters that are relevant, say, to your moral conduct. Likewise, it may seem obvious that others are entitled to expect that you are. On the other hand, it is important to flesh out the details of why this is so, given that what this line of thought recommends is actually, I think, quite demanding in comparison to views according to which our epistemic obligations are strictly negative, or views according to which meeting our positive epistemic obligations does not involve forming beliefs about new topics. Moreover, fleshing out the details of the interplay between our epistemic and non-epistemic obligations can provide us with better sense of what staying informed on matters relevant to our non-epistemic obligations actually entails. This is especially important as it provides a route for dealing with the fact that the view is demanding in comparison to other accounts of epistemic normativity.

In the remainder of this section, I argue that at least sometimes, others are (to borrow Goldberg's phrasing) entitled to expect that we form justified beliefs about topics that are relevant to our moral, social, and political obligations. In line with comments I've made previously, I do not wish to suggest that there is a clear dividing line between these categories, as they often overlap, and the lines between them often can be blurred. There is perhaps a case to be made that one or more of these categories can be subsumed by the others. Nevertheless, since I do not want to make any assumptions about this, and since I want to leave open the possibility that there are cases which clearly fall under the umbrella of one category and not the others (e.g., social, but not moral), I will offer separate, though analogous, treatments of each.

3.1.1 Topics relevant to one's moral obligations

Earlier, we discussed how meeting one's moral obligations often requires having justified beliefs about topics relevant to these obligations, given that lacking such justified beliefs often means that one will cause (or will significantly risk causing) moral harm. For this reason, one must do one's best to ensure that one's beliefs about these relevant topics are justified. What I'd like to draw our attention to now is that it can be the case that the topics one needs to have justified beliefs about in order to meet one's moral obligations are not topics one has beliefs about at all. This might be because one simply hasn't given thought to the matter, or because, while one has given brief thought to the matter, one neglects to give it the attention it deserves. Let's look at some examples to illustrate the point.

There plausibly exists a moral obligation to not prop up or turn a blind eye to social inequalities. Of course, figuring out what roles one plays in systems of power and oppression can be quite difficult, given that one can at once be a beneficiary and victim of injustice, and given that simply carrying on as normal in a society in which there is social inequality often just is to help prop up those inequalities. So, determining whether one meets this obligation, or even how best to meet it, can be difficult. Furthermore, there are some (indeed many) moral agents who are not in any real position to address or change those inequalities—perhaps because they themselves are victims of inequality, or because they simply lack the power to make any real positive difference. That said, when it is in their power to do something about those inequalities—particularly if they benefit from them—moral agents ought not help prop them up, or turn a blind eye to them. Consider:

Ignorant Ian

Ian is an upper-level manager at a company employing thousands of people. Despite its official stance of treating all of its employees with equality, there is unjust wage disparity. Moreover, the fact that there is wage disparity is fairly noticeable. White people on average make more money than people of color, and men on average make more than women and those who are gender non-conforming. Given his position, Ian bears partial responsibility for making sure this kind of problem does not exist. Nevertheless, he has not given the issue any serious thought. He does not systematically look into how wages and salaries are determined, whether there are patterns as to who gets promotions or raises (and who doesn't), or even whether there is any notable disparity in general.

If anyone has an obligation to not help prop up social inequalities, Ian does.³ And yet, despite his professional responsibilities, and despite being in a position to effect positive change, he simply does not give the issue any thought. He has failed to do what he morally ought. One obvious explanation for this failure is that he lacks beliefs about topics directly relevant to the obligation. Accordingly, I submit that Ian ought to form beliefs about topics relevant to his moral obligation.

Notice, crucially, that the fact that he has not given the issue any thought certainly does not absolve him. Given that wage disparity is now a commonly discussed problem, and given his professional rank and responsibilities, wage disparity in his company is something Ian should have beliefs about. Since he lacks beliefs about

³The reader might wonder whether this is the best example to use, given that Ian's obligation seems to be part moral, part professional. However, I don't think the fact that this obligation is part professional in any way diminishes the fact that the obligation still has a moral dimension. After all, presumably the reason it is a professional obligation at all is because it derives from a moral obligation.

this matter, he has an epistemic obligation, deriving from his moral obligation, to seek out information on the matter. Not only ought he seek out relevant information, but he also ought to form *justified beliefs* about the matter. Put differently: since he does not have justified beliefs about this matter—let alone beliefs about it at all—he has an obligation to form beliefs about it. And, these beliefs ought to be justified, since having unfounded, biased, or poorly informed beliefs will certainly not suffice for meeting his moral obligation. For having unjustified or poorly informed beliefs about the topic will either cause moral harm to those who are victims of the unjust wage disparity or else put them at significant risk of moral harm.

Of course, sometimes circumstances are not so black-and-white. Let's muddy the example a bit by supposing instead that Ian has given some thought to the matter. He knows that wage disparity is an issue for many companies, and he is aware that part of the responsibility falls on him to make sure that his company does not contribute to this general problem. On account of this, he does take some time to look into the matter and notices that there is, in fact, wage disparity. He comes to realize that white people on average make more money than people of color, and men on average make more than women and those who are gender non-conforming. However, engaging in a merely cursory investigation of why the disparity exists—for example, by briefly looking at one or two specific instances of wage disparity between two people holding the same position—he simply chalks up the disparity to differences in qualifications, education, experience, and work ethic. In other words, although he notices the wage disparity, he determines (or perhaps even decides) that it must not be unjust or in any other way problematic. With his mind more or less

made up, Ian resolves to look no further into the issue.

The problem here is not just that he doesn't look at more than a couple of examples. He also fails to even consider the fact that there might be factors currently unknown to him which contribute to the disparity. He fails to actually investigate the issue, and instead simply turns his attention elsewhere. As before, Ian has not met his obligation to not contribute to or help prop up social inequalities. Although he did give the issue some consideration, he did not *really* look into it. Chalking up the wage disparity in his company to factors such as differences in qualifications and work ethic is inexcusable, especially when his investigation was so superficial, and especially given that those who fight against social inequality regularly point out that these considerations are often incorrectly cited as a means to justify wage disparity. In short, Ian went through the motions of looking into the matter and then relied on the very excuses he should have known to be wary of in order to dismiss the problem.⁴

At this point, I want to pause to consider in more detail what prevented Ian from doing what he morally ought to have done. One contributing factor was his belief that there is no problematic or unjust wage disparity in his company. This belief was based on a hasty and superficial investigation, and as such, was unjustified. However, the problem goes deeper. For, in order to have a justified belief about

⁴To reiterate one of the points made above: whether this case would count as epistemically interesting by Goldberg's lights would depend on whether or not Ian in some sense recognized that his investigation was not as thorough as it should have been. If he didn't recognize, it would count epistemically interesting; if he did recognize, then it wouldn't. Likewise for other cases I discuss. To repeat: my aim is not to determine whether or not these cases count as epistemically interesting, but rather to establish that the subjects in these examples lack beliefs about topics they ought to have justified beliefs about.

whether there is unjust wage disparity in his company (for ease of discussion, let's call this topic *T*), Ian would have needed to investigate further, and then form justified beliefs about topics that are themselves relevant to *T*. For example, in order to have justified beliefs about *T*, Ian would need to have justified beliefs about how race, gender, and gender expression impact one's chances of moving up social and professional ladders—hence contribute to the problem of wage disparity. Additionally or alternatively, he would need to have justified beliefs about how being socialized in virtue of having particular social identities can impact whether one is a victim of, or beneficiary of, wage disparity. Additionally or alternatively, he would need to need to have justified beliefs about how social networking, implicit biases, and stereotypes more generally have an impact on who is perceived to be more or less qualified, or to have a better or worse work ethic—and thus, how these factors in turn have an impact on wage disparity.

In short, in order to have justified beliefs about the topic which is what we might call *directly relevant* to the fulfillment of his moral obligation—where this directly relevant topic is whether or not there is unjust wage disparity in his company—Ian also has to have justified beliefs about topics which are what we might call *indirectly relevant* to the fulfillment of that obligation. These indirectly relevant topics are ones such as those mentioned—i.e., how socialization, stereotypes, biases, networking, and so forth contribute both to wage disparity and perceptions of it. In other words, having justified beliefs about the directly relevant topic(s) requires having justified beliefs about the indirectly relevant topics. The latter are a precondition of the former.

What this tells us, I submit, is that fulfilling his moral obligation would not only require having justified beliefs about the topic directly relevant to his moral obligation. Because he does not have beliefs about indirectly relevant topics, and because having justified beliefs about those topics is (at least in this particular case) a precondition for having justified beliefs about the directly relevant topic, it turns out that meeting his obligation also involves forming justified beliefs about those indirectly relevant topics. Meeting others' legitimate expectations requires it.

Let's consider another example to help bring out the point:

Reckless Rita

Rita believes that global warming is a hoax. She also engages in a variety of behaviors that are damaging to the environment: she does not recycle, she regularly litters, and is in general a wasteful person. Most egregiously, she owns a manufacturing company which produces a great deal of carbon emissions. Rita's friends have urged her to be more responsible and have pointed her to a number of sources which explain why her behavior and her company is damaging to the environment, and which offer suggestions for reducing one's carbon footprint. Rita thanks her friends for their concern, but informs them that she hasn't taken climate change seriously in the past, and she isn't about to start now. Ignoring their requests, she continues to carry on in her typical unsustainable and self-absorbed fashion, remaining generally unaware of how her actions contribute to global warming.

It seems clear to me that Rita ought to take a look at the articles her friends have suggested. In carrying about her life in ignorance, she fails to meet her moral obligation to not pose excessive risk to the environment. Since others are entitled to

expect that she meet this moral obligation, Rita ought to form justified beliefs about how her actions impact the environment.

Like Ian, then, Rita lacks justified beliefs about topics which are relevant to one of her moral obligations. However, it's not just that she lacks *justified* beliefs; she also lacks beliefs at all about some of these relevant topics. Let's unpack this. We've stipulated that Rita has an unjustified belief about a topic which is directly relevant to her moral obligation. This directly relevant topic is whether her behavior contributes to climate change (again, we'll call this topic *T*). Rita's belief about *T* is that she is not contributing to climate change. Given that Rita's belief about *T* is unjustified, in order to meet her moral obligation, Rita would need to form beliefs about topics which are themselves relevant to *T*. These topics might include: what climate experts have to say about carbon emissions; the amount of carbon emissions her manufacturing company produces, and the impact that has on the environment; how excessive waste and lack of recycling negatively impact the environment; and more. Importantly, these are topics about which she currently lacks beliefs. In choosing to bury her head in the sand and carry on in life without informing herself on matters that make a moral difference, Rita fails to meet one of her moral obligations. And insofar as meeting her moral obligation requires having justified beliefs about these topics, Rita ought to have justified beliefs about them. Others are entitled to expect that she do so.

To generalize the points being made so far: To fulfill their moral obligations, moral agents sometimes need to have justified beliefs about topics which are either directly or indirectly relevant to their moral obligations. In such cases, when moral

agents lack justified beliefs about the relevant topics, they ought to perform belief-influencing actions so that they may form justified beliefs about such topics, as this is a prerequisite for meeting their moral obligations.

3.1.2 Topics relevant to one's social obligations

A similar argument can now be advanced which concerns topics relevant to one's social obligations. For ease of discussion, let's modify our earlier example in which my friend and I plan a trip together. Before, we discussed how each of us was entitled to expect the other to competently carry out her assigned tasks. For instance, given that I agreed to be in charge of budgeting for travel, meals, and other non-lodging expenses, my friend has a right to expect that I am reasonably competent to budget—which means that, among other things, I have justified beliefs about how to budget. Now, suppose that I have uninformed or misinformed beliefs about how to budget for this trip. Suppose, for example, that I simply suppose without looking into the matter that the exchange rate of dollars to various European currencies is 1:1. As a result, I underestimate how much money needs to be set aside for travel, souvenirs, and so forth. In virtue of failing to live up to her legitimate expectation that I can carry out my assigned tasks competently—i.e., in virtue of having unjustified beliefs about relevant topics—I have violated my friend's legitimate expectations.

As was discussed with respect to moral obligations, then, sometimes what causes one to fail to meet one's social obligations is that one lacks justified beliefs about topics which are relevant to those obligations. When this is the case, meeting one's social obligations necessitates forming justified beliefs about these matters. Above,

we made the distinction between matters that are directly relevant versus indirectly relevant to one's moral obligations, and I proposed that lacking justified beliefs about matters of either variety can cause one to fail to meet one's obligations. The same point applies here.

In this particular case, it's safe to assume that I have beliefs about topics directly relevant to my social obligation, including: what budgeting involves, how much time it takes to budget, and so forth. It's just that the beliefs I have about these directly relevant topics are unjustified in virtue of being un- or misinformed. In order to fulfill my obligation, then, I would need to have justified beliefs about these topics. And in order to have justified beliefs about *these* topics, it very well may be the case that I need to form (justified) beliefs about topics which are themselves relevant to these directly relevant topics—i.e., indirectly relevant topics.

Suppose I have not even considered the fact that there might be other expenses of which I am currently unaware that need to be factored in. For instance, I have not considered the possibility that there will be foreign transaction fees; that in order to charge our electronic devices, adapters will have to be purchased; that if we want to use our mobile devices, we will have to purchase temporary international plans or else rely entirely on WiFi hot spots; and so forth. And yet, in order to have justified beliefs about how much money to set aside in the budget, I need to have justified beliefs about these topics—topics about which I currently lack beliefs. Thus, I ought to form justified beliefs about these topics in order to meet my social obligation—i.e., in order to meet my friend's legitimate expectations.

After all, when my friend places her trust in me to competently carry out my

tasks in planning the trip, she is entitled to expect that if I am not competent to fulfill my tasks, I either become competent or else take on other tasks which I can competently complete. As evidence, consider how my friend might reasonably react when we run into financial trouble on our trip as a result of my ignorance and budgeting incompetence. If I attempt to exonerate myself by claiming that I didn't know any better, it would be appropriate for her to point out that this is no excuse. If I didn't know how to budget properly, I should have learned or else let her know that I wasn't up for the task.

Moreover, even if by a sheer stroke of luck I manage to set aside enough money in our budget and do not, in fact, undermine the success of our cooperative efforts, my friend would still be within her rights to sanction me for my lack of competence. Suppose it comes up in conversation that I really didn't know what I was doing when I set the budget for our trip, and there were a lot of expenses we encountered on our trip that I did not anticipate. Rather than basing the budget on research I can conduct, or on past experience, I made an uneducated but lucky guess. I think it would be entirely appropriate for my friend to sanction me—for instance, by insisting that I not be in charge of finances if we ever take trips in the future unless I put in the effort to become competent to budget.

As before, the lesson can be easily generalized: To meet their social obligations, social agents sometimes need to have justified beliefs about topics which are either directly or indirectly relevant to their social obligations. In such cases, and when social agents lack justified beliefs about these relevant topics, they ought to perform belief-influencing actions so that they may form justified beliefs about such topics,

as this is a prerequisite for meeting their social obligations.

3.1.3 Topics relevant to one's political obligations

Finally, an argument focusing on political obligations can be similarly advanced. To begin, it would be useful to say a few words about why political considerations are being brought up now, and not before. On the one hand, it seems fine to me to suppose that political obligations can be reduced to social obligations, moral obligations, or some combination of the two. It is for this reason that I did not provide an independent argument for how epistemic obligations can be grounded in political obligations. On the other hand, given that I am interested in whether and how we ought to be open-minded with respect to issues that are morally or politically controversial, it is important that I eventually bring political obligations to the forefront of the conversation. Consider:

Uninformed Uma

Uma plans to vote in an upcoming local election. One of the items included on the ballot is a motion to address the growing homeless population in her community—a topic Uma knows relatively little about. She knows that there are homeless people in her community—she sees them regularly on her commute to work—but she does not know which local structures or measures (if any) are in place to address homelessness. Though she has the time, Uma does not look into the proposal before voting. Moreover, the little information that will be included on the ballot will certainly not suffice to bring her up to speed on what the proposal actually entails. On election day, Uma votes in favor of the proposal, thinking that this will benefit the homeless population. In fact,

the proposal is largely detrimental to that population.

It seems clear to me that democratic political agents have a political obligation to vote responsibly. To be democratic *participants*, political agents need to have justified beliefs about the matters they're voting on—or, at the very least, they need to have informed action-guiding commitments which govern how they vote. If they do not, they would simply be going through the motions; more or less randomly selecting boxes on a voting ticket hardly counts as participating. Put differently, democratic political agents are entitled to expect that fellow political agents be informed when they cast votes. Because Uma was an uninformed voter—and especially given that her vote negatively impacted the community she wished to help—Uma has failed to meet her political obligation. Meeting her political obligation would have required her to form beliefs about topics she currently lacks beliefs about, including: what homelessness in her community looks like specifically, what measures (if any) are already in place to address it, or how the motion (if approved) would change anything. She ought to have formed justified beliefs about these topics before she voted, and others are entitled to expect that she had done so.

As before, the point generalizes. Because I am a political agent, it generally will matter quite a bit that I have justified beliefs about, say, the harms of discriminatory practices, the travel ban, prison reform, the legitimacy of socialist practices in a liberal democracy, the conditions under which speech should be sanctioned, whether we owe reparations for past injustices, and more. For being a democratic citizen requires participation, and political agents ought to have beliefs about important political matters.

Of course, what it means to be a responsible political agent varies by context. The era, country, and community one lives in, one's socioeconomic position, the various demands one has to face on a day-to-day basis, one's social identity, and much more, all affect the extent to which one can exercise political agency, the political obligations one is bound by, and the strength of those obligations. These are considerations that have to be kept in mind, for it would be a mistake to assume that we can make blanket statements about political obligations "we" are all bound by. That said, it is, I think, certainly possible to make fairly general claims about how our political obligations are related to epistemic obligations. All I aim to establish here is that meeting one's political obligations—whatever those might be—sometimes requires having justified beliefs about topics that are either directly or indirectly relevant to those obligations. And, when one lacks beliefs about these topics, one ought to form belief-influencing actions so that they may form justified beliefs about such topics, as this is a prerequisite for meeting one's political obligations.

3.2 Proportional strength of obligations

The second implication of taking the Cliffordian line of thought seriously is that the importance of meeting our epistemic obligations is proportional to the importance of meeting the non-epistemic obligations which underlie them. In the moral case, for example, the greater the moral harm that is at risk, the more it matters that one's beliefs about relevant topics are justified. Compare the following two cases:

Promise Breaking

Pablo promises his father that he will go to breakfast with him on Saturday. Pablo is a very forgetful person, and knows this about himself. There is a significant chance that he already has plans Saturday, but he does not check his schedule. He believes it is open. As it turns out, he has already promised to go to breakfast with his best friend that day. When Saturday comes, he will have to break his promise either to his father or to his friend.

Promise Breaking*

Pablo promises his father that he will bring him to his chemotherapy appointment on Saturday morning. Pablo is a very forgetful person, and knows this about himself. There is a significant chance that he already has plans Saturday, but he does not check his schedule. He believes it is open. As it turns out, he has already promised his best friend that he will watch her children as she attends a potentially career-changing meeting with a client that morning. When Saturday comes, he will have to break his promise either to his father or to his friend.

I think it is fairly safe to say that a greater wrong is committed in Pablo*. Whereas in the first case, Pablo's promise breaking merely frustrates the breakfast plans of either his father or friend, in the second case, Pablo will seriously fail either his father or friend, both of whom are depending on him to help them make important appointments. As evidence that it matters more that his beliefs are justified in the second case, all we have to do is consider how much stronger the appropriate sanctions would be in that case. The takeaway here is simply: Insofar as it matters more, morally, that we keep certain promises as compared to others, it also matters more that our beliefs relevant to keeping those promises are justified.

Consider also the following pair of cases:⁵

Allergy

My niece has a mild nut allergy. When I babysit her, I have to take care in order to prevent her from having a mild allergic reaction. In particular, I have to avoid using dishes and silverware that may have come into contact with nuts. Additionally, I have to make sure that my beliefs about what foods contain, or may have been exposed to, nuts are justified. For instance, I have to read labels on food packaging so that I don't feed her anything that might contain nuts while having unwarranted confidence that she will not be harmed.

Allergy*

My niece has a severe nut allergy. If she is exposed to nuts this can send her into anaphylactic shock. When I babysit her, I have to be extremely careful. In particular, I have to avoid using dishes and silverware that may have come into contact with nuts. Additionally, I have to make sure that my beliefs about what foods contain, or may have been exposed to, nuts are justified. For instance, I have to read labels on food packaging so that I don't feed her anything that might contain nuts while having unwarranted confidence that she will not be harmed.

If we are inclined to say that I should be much more careful in general in the case where my niece has a severe allergy—given that she is at risk of being much more seriously harmed—then I think we should say that I ought to be more careful epistemically, since being epistemically irresponsible is one of the ways in which I can allow serious harm to come to her. Again, the lesson to draw from these examples is that when there is risk of greater moral harm, the more it matters that

⁵This pair of examples is adapted from one appearing in Smith (1983, 551).

one's relevant beliefs are justified. Although I've chosen cases which involve moral obligations to illustrate the point, the same lesson applies to cases involving other types of non-epistemic obligations. In general, then, the stronger one's non-epistemic obligation, the stronger one's corresponding epistemic obligation.

As I mentioned earlier, one advantage of highlighting this implication of the view is that it provides us with a better sense of what staying informed on matters relevant to our non-epistemic obligations actually entails. This is especially important as it provides a route for dealing with the fact that the view is demanding in comparison to other accounts of epistemic normativity. In the next section, I'll say a few words about what meeting our epistemic obligations involves.

3.3 Meeting our epistemic obligations

My aim thus far has been to establish that what one epistemically ought to do is intimately related to what one ought to do morally, socially, and politically. In doing so, I have focused the discussion on the epistemic obligation to make sure that one's beliefs are justified. Of course, this is to speak in rather general terms, for making sure one's beliefs are justified cannot be reduced to a single activity. Rather, it involves a variety of actions a person ought *to perform* in order to (try to) ensure that her beliefs are justified, as well as specific actions she ought *not perform*. Actions one ought not perform include: closing inquiry too soon; actively avoiding evidence or sources of evidence that one is unjustifiably biased against; and so forth. Actions one ought to perform include: open-mindedly engaging with evidence relevant to

one's inquiry; seeking out a variety of information (from a variety of sources) about a subject before forming beliefs or opinions about it; considering the credibility of one's sources of information; and so forth. Like others,⁶ then, I think that our epistemic obligations concern belief-influencing actions. This way of understanding epistemic obligations assumes that (given doxastic involuntarism) we cannot *control* our beliefs. At best, we can perform actions that will *influence* how we form and maintain beliefs.

Now, when it comes to negative epistemic obligations, these are easy enough to know how to meet. One simply refrains from performing epistemically bad belief-influencing actions.⁷ However, when it comes to positive epistemic obligations, things are not quite as obvious. Given that meeting one's positive epistemic obligations involves actually performing, rather than refraining from performing, a belief-influencing action, how does one know when enough is enough? For instance, assuming one has an epistemic obligation to seek out evidence relevant to some topic *T*, how does one know when one has sought out enough evidence? To help the discussion along, I will sketch some parallels between positive epistemic obligations and positive moral obligations. In doing so, I do not mean to suggest that the analogy between moral and epistemic positive obligations is perfect. Even so, the comparison

⁶See, e.g., Peels (2017).

⁷Of course, the line between positive and negative epistemic obligations (or actions, for that matter) is not so clear cut. After all, one sometimes has to put in active effort to make sure that one does *not* engage in bad belief-forming practices. E.g., if a person is prone to be a wishful thinker, remedying this would presumably involve active effort on her part to train herself to form more accurate beliefs. It should be noted that this point carries over to negative moral obligations. For instance, the obligation to not steal is typically considered a negative moral obligation. One simply needs to refrain from stealing in order to fulfill the obligation. However, for kleptomaniacs, refraining from stealing can be quite difficult. Thus, not stealing actually requires active effort on their part.

will be useful for present purposes.⁸

Typically, positive moral obligations are treated as *imperfect* obligations, meaning that, under most circumstances, there is no specific action a person has to perform in order to meet her obligation. Instead, there are a number of actions she could perform which, individually or jointly (depending on the case), would suffice for meeting the obligation. In other words, a person typically has some leeway or discretion at her disposal when it comes to meeting her imperfect moral obligations. For instance, the obligation to donate to charity is imperfect. We do not have to constantly act toward meeting it, as we lack the temporal and financial resources to do so. Instead, there is some leeway as to when, and by how much, we donate. To fulfill the obligation, we simply have to regularly donate to those less fortunate to the extent that we are able. Relevant, of course, will be how well off we are ourselves; we are only bound by a obligation to donate if we have some to spare. Furthermore, the more we have to spare, plausibly the more we ought to give away.

Because we enjoy a certain amount of leeway when it comes to meeting the obligation, we rarely have a obligation to donate in a particular instance. That is, when we refrain from donating in a particular instance we typically are not thereby failing in our obligation to donate to charity. For example, I do not necessarily fail to meet my obligation to donate if I decide not to donate \$1 to a local charity at the grocery store check-out. Instead, my obligation to donate to charity at best—at least under normal circumstances—gives me a reason to donate to some charity or other(s), on a regular basis, as my means allow. Of course, there may be cases in

⁸Others have drawn similar comparisons. See, e.g., Stapleford (2011) and Lackey (2018a).

which I fail to meet my moral obligation if I fail to donate in a particular instance. Imagine I cross paths with a homeless person who is clearly starving, and there is nobody else around willing to provide her with a meal. If I selfishly decide to ignore her, despite having good reason to suppose she will die soon without food, then I fail to meet my moral obligation. This type of situation, however, is typically the exception when it comes to meeting one's imperfect obligations.

Similar considerations apply to imperfect epistemic obligations. For ease of discussion, we'll return to our familiar budgeting example. In order to meet the social obligation I have to my friend, I ought to do some research in order to make sure I prepare a good and accurate budget. This socially grounded epistemic obligation, like the moral obligation to donate to charity, is an imperfect obligation, for I have some leeway when it comes to how I fulfill it. In other words, in order to meet the obligation, I do not necessarily have to look at any particular sources of information. Instead, meeting this obligation simply requires seeking out some reliable sources of information or other on how to budget well, as well as some reliable sources of information or other concerning extra costs to be prepared for when traveling internationally. Moreover, while I ought to spend enough time researching to be sufficiently knowledgeable to prepare a good budget, I am also allowed some discretion when it comes to when and how I research.

We can illustrate further by returning to another familiar example. Recall Ian, who has a moral obligation to not promote or help prop up unjust wage disparity. To meet his obligation, Ian ought to spend sufficient time looking into how wages are determined at his company, taking care that his sources of information are reliable.

Since he is not already knowledgeable on the subject, he should also investigate the ways in which unjust wage disparity might look superficially unproblematic. Given that this morally grounded epistemic obligation is imperfect, however, he has some discretion as to how he goes his investigation.

Finally, it is worthy pointing out the following similarity between moral and epistemic imperfect obligations. When it came to the moral obligation to donate to charity, although a person generally gets to choose (within some constraints) where, or to whom, she donates her money, there are going to be cases in which failing to donate to a specific person or cause suffices for violating her obligation. So too with imperfect epistemic obligations. Suppose Ian is presented with a detailed document prepared by a number of employees at his company who do not identify as white males (that is, then, the very people who on average make less at the company than white males). This document details the wage disparity in the company, and it explains why the disparity is unjust. If Ian decides not to use this report as one of his sources of information in his investigation, he would arguably thereby fail to meet his epistemic obligation, for it is too relevant and available for his consumption for him to reasonably ignore it.

To sum up: comparing the kinds of epistemic obligations with which I am concerned to imperfect moral obligations provides us with a good sense of what meeting them involves. Here are a couple of points to take away from the discussion. First, given that positive epistemic obligations are best understood as imperfect, we generally are allowed some discretion when it comes to meeting them. This is especially important to note, given that we only have so much time and energy to devote to

meeting our positive epistemic obligations. Second, in line with what was said in the previous section, the strength of one's epistemic obligations to perform belief-influencing actions A_1 - A_n is directly related to the strength of the non-epistemic obligations one has. Understanding this helps us determine which imperfect epistemic obligations we ought to prioritize meeting. For example, it will matter more that I have justified beliefs about topics relevant to a strong moral obligation than that I have justified beliefs about topics relevant to some relatively weak social obligation. Thus, I ought to devote more time and energy to investigating topics relevant to the moral obligation, as it is more important all things considered.

Chapter 4

Objections and Replies

This chapter is devoted to addressing a handful of objections that could be leveled against my view. In addition to the obvious benefit, addressing these objections provides me with the opportunity to clarify certain aspects of my view.

4.1 I make claims about the obligations we have.

Who is the “we?”

In a recent paper, Jennifer Lackey (2018b) considers views according to which we have an epistemic obligation to object to what we believe to be false or unwarranted—especially when what is objected to seems morally or epistemically harmful. Drawing on Charles Mills’ work, Lackey builds the case that a constraint on any such theory is that “one’s obligation can be directly influenced by one’s social status” (2018b, 82). As applied to the obligation to object in particular, she concludes that “people have

different obligations [and] are under various kinds of normative pressure to express their dissent, depending on who they are and what social position they occupy” (2018b, 92). Thus, the strength of the obligation, whether one is excused from meeting the obligation, and whether one has the obligation at all, depends on one’s social status. For, among other things, one’s social status has all kinds of relevance as to whether one’s objection will be taken seriously, what one risks when one objects, and whether one has sufficient time or energy to object.

Lackey illustrates with the following example. Suppose a white male tenured professor and a black female tenure-track professor both overhear a colleague making a blatantly sexist remark. It’s fairly obvious that someone should call out the remark. However, to say that both professors have an equal obligation to object completely disregards the fact that the female professor has much more to lose by calling out her colleague than the male professor. Not only might the male professor be more effective if he speaks up, but also, if she speaks up, she puts herself in a vulnerable position. She might be perceived as too sensitive, or as too willing to complain. Alternatively, or additionally, it might have a negative impact on her social standing in the department, for she might be treated coldly in the future by the person she called out or by those close to him.¹

Thus, accounts of an obligation to object should not begin by making general claims about the universal duty to object, and then allow exceptions for those in “less than ideal” circumstances. For being in less than ideal circumstances is often the norm, not the exception. As Lackey observes, the closer we look at actual

¹For more discussion, see Lackey (2018b, 91-94).

circumstances, the more “we find ourselves recognizing that objecting is often a luxury, one that not everyone can afford to make” (2018b, 82). More generally, what any individual ought to do is determined by their circumstances, by what is afforded to them (or not), by what they would have to risk or sacrifice (or not). Important facts about power, marginalization, oppression, and agency should thus be treated as central, rather than peripheral, considerations in any theory of obligation.

This stance on accounts of obligations is compelling. Perhaps unfortunately for me, however, it does appear to inspire an objection to my account. Namely, in arguing that epistemic obligations are agent-universal, am I not denying that people have different obligations and are under various kinds of normative pressure, depending on who they are and what social position they occupy? To borrow a phrase from Mills: in arguing that there are agent-universal epistemic obligations, am I not relying on ideals to the marginalization or exclusion of the actual?²

This worry is important to address for a number of reasons, one of which is that doing so allows me to elaborate on a point I’ve touched on at various points of the dissertation, but which I should state as clearly as possible. Namely, I certainly do not want to make the mistake of starting out with general claims about what we (in an idealized sense of “we”) all owe each other, and then abstract away from the ideal in order to accommodate the fact that oppression and power differentials directly affect what individuals owe (or don’t) to each other. I agree with Mills, Lackey, and others who insist that any theory that does this will wind up with distorted conceptions of obligations, including who is bound by those obligations and under

²See Mills (2017, 75).

what circumstances. At the same time, I think it would also be a mistake to assume that the fact that epistemic agents occupy unique positions in systems of power and oppression means that we cannot offer a unified account of how our epistemic obligations are intimately related to our non-epistemic obligations.

One of the central claims for which I've been arguing is that what epistemic obligations one has is dependent on one's social and moral obligations—*whatever those might be on an individual-by-individual basis*. In other words, establishing a general connection between epistemic obligations and non-epistemic obligations is entirely consistent with claiming that which obligations one has will depend on one's positionality. Although the details of specific examples will make all of the difference as to whether one is bound by a specific moral, political, or social obligation, the fact remains that every epistemic agent—whether marginalized or privileged³—will have non-epistemic obligations of various flavors and strengths. And, whenever having justified beliefs about topics relevant to these obligations is necessary for meeting said obligations, that agent has a corresponding epistemic obligation to have justified beliefs about said topics.

Here is a different way of articulating the point. In Chapter 2, we considered the possibility that epistemic obligations underpinned by moral and social obligations may not be agent-universal. We saw that this was misguided on two fronts. For one, epistemic obligations underpinned by moral obligations are agent-universal since all

³To be clear, I do not mean to suggest that marginalization and privilege are mutually exclusive. Nor do I mean to suggest that they are one-dimensional. Often times, a person will be marginalized or privileged to varying degrees, and across different dimensions. For instance, one might be privileged insofar as she is an educated professional with financial stability, yet marginalized insofar as she is a black queer woman. Furthermore, she might be more or less marginalized or privileged in certain contexts or social groups than others.

epistemic agents (at least the non-idealized epistemic agents about whom I'm theorizing) are also moral agents. And moral agents can't well decide to simply stop being moral agents. Thus, even though what moral obligations various individuals are bound by will vary from case to case, all moral agents have moral obligations, and in virtue of this, corresponding epistemic obligations. This is the sense in which morally grounded epistemic obligations are agent-universal. Second, epistemic obligations underpinned by social obligations are also agent-universal since, although social agents can often decide to exit particular roles in order to relieve themselves of obligations specific to those particular roles, they cannot exist all social roles without going entirely off the grid. Thus, even though what social obligations various individuals are bound by will vary from case to case, all social agents have social obligations, and in virtue of this, corresponding epistemic obligations. This is the sense in which socially grounded epistemic obligations are agent-universal.

Let's modify Lackey's example in order to illustrate the point. Absent any defeating conditions or considerations, it's clear that the male professor ought to object to the sexist remark made by the colleague. It could be argued whether the female professor ought to object; perhaps she is excused from her obligation, or perhaps she does not have it all. In any case, what she ought to do and what her male colleague ought to do are not the same. Let's suppose she is not obligated to object. Even so, this does not mean that there will not be other instances in which she ought to object to something that has been said. For instance, suppose she is out to dinner with two colleagues who hold no noticeable position of power relative to her. When one of them makes a homophobic remark, she ought to object to what was said.

Of course, things are not necessarily so simple, and the modified example is not perfectly analogous to Lackey's original example. For someone who enjoys having their contributions respected to a high degree *by default*—such as a white tenured male professor—stands less to lose *by default* than someone who only in certain contexts might enjoy having their contributions respected to a high degree—such as a black woman early in her career. Even if her contribution in this context would be respected to a similar degree to the male professor's contribution (that is, were he to object to the sexist remark in the original example), it still holds that a person who lacks certain privileges will regularly face battles—such as being stereotyped, or not taken seriously—that those who are largely privileged do not have to face. And whether or not one has to regularly face and navigate such oppressive forces can make a difference as to what one ought to do in any given context. After all, time, energy, and willpower are limited commodities. And when one has to expend them on a regular basis just to get by, this limits how much one has to give when “obligation calls.”

The bottom line I aim to establish is simply: whether one's specific non-epistemic obligations are shared by others, most (if not all) agents have some non-epistemic obligations of some form or other. Whatever one's moral, political, and social obligations amount to, whenever having justified beliefs about relevant topics is necessary for the fulfillment of these obligations, one ought to have justified beliefs about these topics. In aiming to establish the ways in which epistemic obligations and non-epistemic obligations are intimately related, I thus do not commit myself to claiming that specific non-epistemic obligations (and specific corresponding epistemic obliga-

tions) are had by all, or even most, agents. As I stated in the introduction of this dissertation, the account I'm developing is meant to be general in the sense that a primary goal is to elucidate the ways in which what one ought to do epistemically (whatever that amounts to in each person's case) is intimately bound up with what one ought to do morally, politically, and socially (whatever that amounts to in each person's case). Doing this is entirely consistent with prioritizing the fact that which specific obligations people are bound by may vary wildly based on the disadvantages and/or advantages they suffer from and/or enjoy.

4.2 The account is too demanding

I've argued that epistemic agents, in virtue of also being social, moral, and political agents, are entitled to expect quite a lot of one another epistemically. They are entitled to expect that other agents' beliefs are justified, especially when these beliefs concern matters of moral, political, or social importance. They are entitled to expect that others *have* (justified) beliefs about certain topics. And they are entitled to expect that other epistemic agents not threaten the shared epistemic resources which make cooperative efforts possible in the first place. In comparison to other accounts of epistemic normativity, my view is rather demanding. In fact, one might worry that it's too demanding.

By way of response to this worry, I will now examine an account of epistemic obligation recently proposed by Rik Peels. This examination provides the reader with the opportunity to compare my account to a less demanding account of epistemic

normativity. It will also provide me with the opportunity to explain why accounts of epistemic normativity should be relatively demanding.

4.2.1 Peels' account of epistemic obligation

There are two features of Peels' account that should be noted upfront. First, throughout the discussion, he takes care to emphasize that a primary constraint on his theory is that it not be over-demanding. In his view, any view of epistemic obligation (and presumably obligation more generally) “on which we violate the vast majority of our obligations and are blameworthy or excused for that is implausible: to be excused should be an exceptional situation, not the normal situation” (2017, 103).

Second, the central case of epistemic obligation on which he focuses his analysis are what he sometimes refers to as “epistemic non-contingent intellectual obligations,” and at other times (and more plainly) as “purely epistemic obligations.” What this boils down to is that Peels more or less treats as peripheral the following kinds of cases: role-specific intellectual obligations (for instance a police officer's professional intellectual obligation to carefully and impartially evaluate evidence—which are often contingent, given their role-specific nature); and moral intellectual obligations (for example, “a non-contingent moral intellectual obligation to not spy on [one's] niece” (2017, 101)—which, though they are often non-contingent, are nevertheless part moral in nature). Instead, and to repeat, he wants to focus his attention on epistemic non-contingent intellectual obligations—i.e., purely epistemic obligations. With these preliminary remarks in place, let's take a look at the specifics of his account.

As a launching point in the development of his account, Peels considers William James' famous observation that being a good epistemic agent involves knowing the truth and avoiding falsehoods. He sees at least three ways to interpret this claim:

1. One could believe more true propositions and fewer false propositions [relative to what one would believe were one not trying to be a responsible epistemic agent].
2. One could believe more true propositions and an equal number of false propositions [relative to what one would believe were one not trying to be a responsible epistemic agent].
3. One could believe an equal number of true propositions and fewer false propositions [relative to what one would believe were one not trying to be a responsible epistemic agent]. (2017, 102)

He rejects the first and second interpretations off the bat since he thinks claiming that there are any agent-universal obligations to add beliefs to our belief sets would be far too demanding.⁴ After all, there are infinitely many true beliefs we could come to have. Many of these truths are mundane, it would be difficult to determine which to add, and we could never discharge our obligations since there would always be more true beliefs we could add.

Interestingly, he also thinks that the third interpretation is problematic, at least in its current form. As he points out, most of us hold a great deal of false beliefs. Part of this presumably stems from the fact that we are products of our upbringings and surroundings, and that sometimes what we are taught to believe is more rooted in convention than truth. Couple this with the facts that we can mis-perceive

⁴Recall, though, that he thinks there are plausibly role-specific (e.g., professional) epistemic obligations to form justified beliefs about certain topics.

things; that we are biased in numerous ways; that over the course of human history we have been mistaken about a great number of things; that we are not perfectly rational agents, and that we regularly make mistakes when reasoning; and so forth. This makes for many people having many false beliefs. Now, Peels does grant that “presumably, for many of these false beliefs, there is *some* action we could perform, such as careful investigation or consulting experts, which would rid us of those [false or unjustified] beliefs” (2017, 103). Even so, he thinks if this third interpretation of James’ claim is true, we would either be blameworthy or excused from a large number of our epistemic obligations. Here is why.

In the first place, we do not, as a matter of fact, perform such belief-influencing actions on a regular basis. So, if we were bound by an epistemic obligation to perform such belief-influencing actions, it would seem as if we would be blameworthy or excused much (if not most) of the time. Moreover, even if we wanted to perform such belief-influencing actions, we’d likely not know where to begin, since from our perspectives all of our beliefs are true. In more detail: we might know that a great number of our beliefs are either false or unjustified. However, we typically cannot tell which beliefs those are from our own perspectives. If we thought of any particular belief *B* that it was false or unjustified, presumably that would at least go some way toward changing our doxastic attitude regarding *B*. Thus, even though we could perform belief-influencing actions to address our false or unjustified beliefs, since we don’t know which beliefs are false or unjustified, “[w]e would have to check virtually *all* of our beliefs” (2017, 103, original emphasis). The bottom line is that, if the third interpretation were correct, we would either be excused or blameworthy for failing

to meet a majority of our epistemic obligations. Given that he thinks that this in itself is a strike against a theory of obligation, he maintains that this interpretation, at least as it stands, is incorrect.

One problem with it, as he sees it, is that it makes epistemic obligations too objective, and as such, generally too difficult to meet. For this reason, he proposes an account of epistemic obligation which can meet the theoretical constraints he has in mind. On his view:

[O]ne has an epistemic intellectual obligation to perform belief-influencing action *A* if one believes that it would be epistemically bad not to perform *A* in that it leads to or maintains certain false beliefs. (2017, 106)

Notice that this way of characterizing epistemic obligations—which can be understood as a modified version of the third interpretation of James’ quote—has a very subjective lean to it. Importantly, Peels does want to maintain that there is room in his account for some objective epistemic obligations. However, he seems to think that the vast majority of our epistemic obligations are best understood in subjective terms. Moreover, it would seem that objective epistemic obligations only come into play when an epistemic agent becomes aware of specific problems with her beliefs. Peels illustrates by discussing cases of “doxastic discrepancy”—e.g., when an epistemic agent comes to believe that one or more of her beliefs were unreliably formed, or that some of her beliefs are in tension with one another. In such cases, Peels maintains that “one’s *objective* epistemic intellectual obligation is to do that—or one of the actions—which will remove the doxastic discrepancy, whereas one’s *subjective* epistemic intellectual obligation is to do that which from the subject’s perspective

renders it sufficiently likely that the doxastic discrepancy is removed” (2017, 107).⁵

Of course, a somewhat obvious objection to this proposal is that when it comes to non-epistemic obligations, both objective and subjective considerations matter. For instance, it’s quite common to hold that moral agents ought to avoid not just subjective moral badness, but also objective moral badness. So why think epistemic obligations would be any different? Peels’ response to this concern is worth quoting at length:

When it comes to non-epistemic obligations, we have a pretty good idea of what counts as, say, a morally bad action and what does not. And that is because as soon as we know what properties an action has, we often only need to reflect on them in order to see whether the action is morally good or bad—in fact, such moral insight is often direct and does not require much reflection. However, when it comes to epistemic obligations, things are quite different. For many of our beliefs, such as virtually all of our perceptual beliefs, we cannot find out merely upon reflection whether they are true or false, and it is often not clear whether a belief is irrational or not and whether it is reliably formed or not. This is simply a matter of how we are constituted. Given [this] ... we have good reason to think that epistemic obligations and non-epistemic obligations are in this regard different ... Moreover ... my argument leaves room for at least some *objective* epistemic obligations. (2017, 105-06, original emphasis)

⁵Another thing to point out: though he does not go into too much detail about this, Peels thinks that subjective epistemic obligations should be put solely in terms of what a subject believes, rather than what she *should* believe. Thus, an epistemic agent is not culpable for failing to believe things she should believe. See (2017, 104-105).

4.2.2 Response to Peels

Compared to Peels' account, my account of epistemic obligation is quite demanding. By his lights, this should make my account implausible. I think he is mistaken. In what follows, I'll explain how my account more accurately captures the nature and number of our epistemic obligations. I'll begin by elaborating on how my account appears over-demanding in comparison to his. I'll then explain why his account is not nearly demanding enough.

To see why my account is comparatively quite demanding, let's return to Peels' three interpretations of James' claim that it is epistemically good to believe truths and avoid falsehoods. Peels denies that we are epistemically obligated to add true beliefs to our belief sets. Moreover, he thinks that—except in cases when we have a reason to think that a particular belief or other is false or unjustified—we are not epistemically obligated to do anything to address false or unjustified beliefs. In contrast, I have argued both that (1) we ought to take measures to make sure the beliefs we do have are justified, especially when these beliefs are relevant to our non-epistemic obligations, and (2) we ought to add justified beliefs to our belief sets when these beliefs would be about topics relevant to our non-epistemic obligations. Moreover, since it's plausible that many, if not most, people fall short of (1) and (2), this would mean that on my view, a great many people are either blameworthy or excused for not meeting their epistemic obligations.

The first issue to address is Peels' suggestion that knowing what one's non-epistemic obligations are, as well as how to fulfill them, is quite easy. Concerning moral obligations, he remarks: "moral insight is often direct and does not require

much reflection,” and that “as soon as we know what properties an action has, we often only need to reflect on them in order to see whether the action is morally good or bad” (2017, 105). This seems to me to be quite plainly false. Certainly, there are times when determining what is morally right or wrong is quite simple, and there is general consensus on what one morally ought to do. That said, figuring out what one morally ought to do can be quite difficult. This is something that professional ethicists—indeed, most anyone who has spent time carefully considering ethics and metaethics—can agree upon. It can be quite difficult to determine what our moral obligations are. And even if we can determine what they are, it can be quite difficult to know how to meet them, when we are excused from meeting them, how much effort we have to exert to meet them, and more. As those who have taught courses in ethics know, a common objective of such courses (which can often be a cause frustration for students) is not to leave the course having all of the answers about what’s morally right or wrong, but rather to leave the course having a better appreciation for the complexity of moral issues, coupled with an ability to think through them more carefully and critically. This is not to endorse a form of moral relativism. It is simply to appreciate the fact that morality is difficult, and that moral insight is often not direct but rather does require a good deal of reflection.

Similar points can be made for other non-epistemic obligations. For example, although it is sometimes easy to determine what one politically ought to do, it is also often difficult to know what our political obligations are, how to best meet them, when we are excused from meeting them, and more. So, the reason we think both objective and subjective considerations matter when it comes to our non-epistemic

obligations cannot be that it's simply easy to tell what our non-epistemic obligations are, as well as how to meet them. This puts pressure on his contention that epistemic obligations should be thought of primarily in subjective terms, whereas non-epistemic obligations should not. Contrary to what he suggests, I submit that accounts of epistemic obligations and non-epistemic obligations should both feature subjective and objective considerations prominently. The fact that he thinks of epistemic obligations as being primarily subjective not only makes his account not nearly demanding enough, but—more to the point—it means his account fails to accurately capture the nature and extent of our epistemic obligations. I will now give additional reasons to support this point.

My second complaint is that Peels' primary constraint on accounts of obligation is far from obviously true. That many people may be blamed or excused for not performing some action *A* does not mean that there is no obligation to perform *A*. We can illustrate using examples we've already discussed. For instance, there plausibly exists for many people a moral obligation to not contribute to pollution and global warming—namely, those who have some degree of power to not contribute. And yet, many people do not put in as much effort as they reasonably could in order to avoid contributing to pollution and global warming. Importantly, the fact that there is more they could do, but don't, does not mean that the obligation does not exist. It simply means that many people are either blameworthy or excused for not fully meeting it.

The same points apply to the obligation to not contribute to, or be complicit in, discriminatory practices. Assuming this obligation exists, many people regularly

fall short of meeting it. A person might sometimes be excused from meeting the obligation. As we've discussed, speaking out against sexist or racist remarks can expose one to certain risks on account of one's social identity. Other times, someone will be blameworthy for failing to meet it. Those who benefit from racist or sexist power structures often accept the benefits they receive while doing very little to reform these power structures. Sexist and racist remarks are often overlooked or ignored rather than addressed. To take another example, democratic citizens have a political obligation to be informed voters, yet many people fail to do the research necessary for being adequately and accurately informed on the issues on which they vote. Sometimes they may be excused for not being informed; other times they may be blameworthy. But the fact that many people fail to be adequately informed does not indicate that the political obligation to be an informed voter does not exist.

The bottom line is: the existence of some moral or political obligation O is perfectly compatible with the fact that many of those bound by the obligation are either blameworthy for, or excused from, not fully or regularly meeting O . If this holds in the moral and political cases, why think the epistemic would be any different? Moreover, the fact that many of us may either be blameworthy or excused for not meeting our epistemic obligations is a result to be expected once we understand that our epistemic obligations are intimately tied to our non-epistemic obligations. After all, assuming it is plausible that many of us fail to fully or regularly live up to a wide variety of our non-epistemic obligations, this would in turn suggest that many of us fail on the epistemic front as well. For presumably at least one reason we fail to meet these non-epistemic obligations is that we don't have enough justified beliefs

about topics relevant to these non-epistemic obligations. We are too uninformed or misinformed. (The voting example is a clear case.)

My next criticism concerns Peels' point that even if we know that there may be a great deal of beliefs we hold which are either false or to some degree unjustified, from our perspectives, we often can't tell which beliefs those might be. And, given that we often can't tell which beliefs those might be, we won't know where to begin in order to improve our epistemic situations. I agree with Peels that we can't well just go checking all of our beliefs in order to determine which of them are false or unjustified. This would be a virtually impossible task. However, the fact that we often can't identify which of our beliefs specifically are false or unjustified does not mean that we have no epistemic obligation to do anything about our epistemic circumstances. I have three suggestions to offer.

First, even if we cannot determine on the basis of reflection alone which of our beliefs may be poorly formed, we regularly encounter evidence which conflicts with one or more of our beliefs, and which gives us cause to examine the bases of our beliefs.⁶ Peels seems to acknowledge this much when he explains that an epistemic agent ought to "remove doxastic discrepancies," which can arise when a subject acquires evidence that one or more of her beliefs were unreliably formed, or that some of her beliefs are in tension with one another. However, he seems to think that these kinds of situations are the exception, rather than the norm. In contrast, I think we regularly encounter evidence which conflicts with one or more of our beliefs. When this happens, we typically ought to revise those beliefs in light of the

⁶Or, if we have already examined the bases of our beliefs, to re-examine them.

new evidence.

Second, we can engage in activities and behaviors aimed at making us more epistemically virtuous. This by itself will not result in our suddenly figuring out which of our beliefs are problematic and which are not. What this can do, however, is help us in the following ways. Moving forward, when we form new beliefs, we will be less likely to add more false or unjustified beliefs to our belief sets than we would were we less epistemically virtuous. Moreover, when it comes to beliefs we already hold, we will hopefully be more likely to spot a problem when those beliefs come under direct consideration. For instance, suppose I am having a conversation with a friend about my stance on whether or not we have moral obligations to the planet—a topic I have not thought about for quite some time. It seems safe to suppose that I would be more likely to question and re-assess the bases of my beliefs—and, that if my beliefs are not well-founded, I would be more likely to recognize this and come to revise them—if I am more epistemically virtuous than if I am less epistemically virtuous. Therefore, even if we have a difficult time determining which of our beliefs are false or unjustified, there are still steps we can take toward reducing the number of false or unjustified beliefs in our belief sets.

A final point is that even if it might be difficult in general to determine which of our beliefs are false or unjustified, there is good reason to suppose that we should be especially attentive to the quality of our beliefs about moral and political matters. For one, moral and political issues are more often than not quite complex. This in itself is a reason to think that we ought to approach these issues with epistemic

caution.⁷ If a person finds herself thinking that the beliefs she has about a complex moral or political issue are on the whole justified, and she is not an expert in the relevant matter, then she should probably actively consider her reasons for thinking she has a good grasp on such matters. Questions she could ask herself include: How much time have I spent considering the issue? On what bases have I formed my beliefs? Is there more information relevant to the issue I could discover or consider that might give me a better understanding of the issue? And so forth. This type of examination is especially called for if a person finds herself thinking that the beliefs she has about a *range* of complex moral and political issues are all justified.

To appreciate the importance of this point, it would be useful to remind ourselves of the discussion in Chapter 1 concerning implicit biases and related factors which can cause our reasoning to go awry. To repeat some of what was stated there: We know that these biases and other problematic belief-influencing factors exist, that they are pervasive, and that many people consider themselves to be exceptions to the rule. Moreover, while we know that our thinking is flawed in general, and that in general we think the grounds for our beliefs are stronger and less influenced by biases and other affective factors than they are, we should expect that matters will be worse when it comes to moral or political topics. For one, such topics are often emotionally charged. Moreover, such topics are often likely to bring to mind—whether consciously or unconsciously—the very stereotypes these biases embody, and these stereotypes can influence judgments in ways that are far less than apparent.

For instance, discussions about the travel ban may bring to mind negative stereo-

⁷Beliefs about complex moral and political issues can thus be contrasted with perceptual beliefs, which we can generally assume are justified and true (absent a reason to think otherwise).

types about immigrants, and these stereotypes may negatively influence a person's judgments in ways she is not even aware of. Important to keep in mind is that this can happen even when the person issuing the judgments consciously disavows these stereotypes. Similarly, discussions about policies proposed by members of a certain political party may bring to mind stereotypes about members of that political party, and these stereotypes may negatively influence a person's judgments in ways she is not even aware of. And so on. So even though it can be difficult to tell from our own individual perspectives which of our beliefs have been poorly formed, once we take into account that our thinking is biased in a number of ways, as well as that there are good reasons to think that biases are especially likely to influence beliefs about specific kinds of issues, we can see that there is reason to focus our critical energy on beliefs that concern those issues.

Of course, whether or not beliefs about moral and political issues are less likely than other types of beliefs to be unjustified or influenced by bias is an empirical issue. So I don't want to exclusively rely on this point. However, I don't need to. As I've argued, while it is important to make sure that one's beliefs are justified in general, it's more important to make sure that one's beliefs are justified when they are relevant to one's non-epistemic obligations. And the stronger or more significant one's non-epistemic obligations are, the more important it is that beliefs relevant to those obligations are justified. Far from not knowing where to start our "belief-checking," as it were, I submit that—in addition to engaging in epistemically good belief-influencing practices in order to become more epistemically virtuous in general—we ought to spend time trying to ensure that the beliefs we have which

are relevant to our non-epistemic obligations are justified. Thus, even if it does not turn out that beliefs relevant to our non-epistemic obligations are especially likely to be biased or influenced by affective factors, we should nevertheless be focusing our critical energies on making sure these beliefs are justified, since so much hangs (morally, politically, and socially) on whether or not they are.

My final complaint with Peels' account is his theoretical focus on so-called purely epistemic obligations. As should be unsurprising by this point, I consider his treatment of purely epistemic obligations as central—and by default other intellectual obligations (e.g., morally grounded epistemic obligations) as peripheral—to be a serious mistake. This mistake is especially important to address considering that Peels is not alone in this. (Indeed, this seems to be the norm, rather than the exception.) Richard Feldman, for instance, also invokes the notion of the purely epistemic when developing his own account of epistemic obligations. On his view, epistemic obligations are “obligations to believe to which the practical benefits of beliefs are not relevant. They are obligations that arise from a purely impartial and disinterested perspective” (1988, 236).

To be clear, I agree with Feldman that the non-epistemic benefits of believing particular propositions are in an important sense epistemically irrelevant. For instance, that believing some *particular* proposition p would result in my being a better moral agent makes no difference to whether I should believe p from the epistemic point of view. From the epistemic point of view, what matters is, e.g., that I believe propositions that are supported by the evidence I do (or should) have. This point is important to emphasize, given that we don't want to make the mistake of assuming

that epistemic normativity is *reducible* to non-epistemic normativity. In other words, I agree with Feldman that what one morally ought to believe, and what one epistemically ought to believe, can come apart.⁸ Nevertheless, as the arguments in the previous two chapters have illustrated, non-epistemic considerations are extremely relevant to our epistemic obligations in at least two important ways. For one, one's non-epistemic obligations determine in large part what one's epistemic obligations are. Moreover, one's non-epistemic obligations determine in large part the strength of one's epistemic obligations.

What this boils down to is that I think the tendency to put the “purely epistemic” on a pedestal is wrong-headed, and that the sooner we realize this the better. Focusing on epistemic considerations without giving due weight to non-epistemic considerations can only result in a distorted conception of epistemic normativity as it pertains to epistemic agents living in the actual world (as compared to, say, idealized epistemic agents). Epistemologists like Peels who aim to account for epistemic normativity in a meaningful way for epistemic agents living in the real world (which Peels undeniably is),⁹ and who treat purely epistemic obligations as the paradigm of intellectual obligation, simply build their theories on faulty foundations.

⁸For discussion, see (1988, 236-37).

⁹Consider some of his remarks: “The purpose of this book is to develop an *account of responsible belief*. By an ‘account of responsible belief’ I mean an informative analysis in terms of necessary and sufficient conditions of what it is to believe responsibly” (2017, 3). Elsewhere, “[W]hat we believe often makes a crucial difference to how we act. Thus, we should take our doxastic household seriously by forming responsible rather than blameworthy beliefs. However, if we want to believe responsibly, we need at least some grasp of what it *is* to believe responsibly. An account of responsible belief, therefore, provides us with the material from which we can derive guidelines for our cognitive behavior” (2017, 7).

4.3 A problem of regress

A further objection is that my account invites a problem of regress. To see why, let's return to the distinction between topics that are "directly relevant" versus "indirectly relevant" to one's non-epistemic obligations. When it comes to meeting one's non-epistemic obligations, there are beliefs one could have about topics that are relevant to said obligations because they concern the obligations themselves. These would count as directly relevant. For instance, beliefs one has about the harms of discriminatory practices in the workplace—e.g., beliefs about wage disparity—are directly relevant to one's moral obligation to not contribute to oppression and discrimination. In addition to this, there are beliefs one could have that are relevant to one's non-epistemic obligations, not because they concern the obligations themselves, but because they are relevant to topics that do concern the obligations themselves—e.g., implicit biases, networking, and socialization.

I argued that having justified beliefs about topics that are directly relevant to one's non-epistemic obligations requires (at least sometimes) having justified beliefs about topics that are indirectly relevant. But once we accept this, wouldn't it stand to reason that having justified beliefs about topics that are indirectly relevant to one's non-epistemic obligations requires having justified beliefs about topics that are in turn relevant to *them*? For example, in order to have justified beliefs about implicit biases, networking, and the like, won't one need to have justified beliefs about topics such as evolutionary psychology, social hierarchies built on class, gender, and race, and more? And in order to have justified beliefs about *these* topics, wouldn't one need to have justified beliefs about topics relevant to them? And so on.

Worrisome as this problem may seem at first, it has a relatively straightforward solution, for a good portion of the epistemic work can be offloaded to experts. In other words, we don't need to become experts ourselves on, say, implicit biases, or the insidious ways in which people are harmed or benefited on the basis of their social identities, in order to have justified beliefs about such topics. We can simply find out, and rely upon, what relevant experts have to say on such matters. Of course, determining which experts to trust can certainly be difficult in and of itself. But doing the work to figure out which experts to trust is far easier than trying to become an expert oneself on a host of complicated topics.¹⁰ Moreover, identifying which experts to trust is not a new problem. Nor is it a problem just for my account. Rather, it is a problem for any view that recommends that laypersons rely on expert testimony.

Even so, a worry remains. Namely, even if we can avoid the problem of regress, it can be difficult to determine when we ought to be seeking out more evidence, or examining the quality of beliefs we already hold. To reiterate one of Peels' points, even if we know that in general we have a good number of false or unjustified beliefs, from our perspectives, it won't always be obvious when more epistemic work is called for.

There are a number of responses to make by way of addressing this epistemological problem of determining when one ought to do more epistemic work. The first is simply that we can always in some sense be aware of our fallibility, and constantly endeavor to improve our belief-forming habits in general. Thus, even if it is generally

¹⁰For relevant work on expertise, see Coady (2012) and Goldman (2001), which explores the problem of figuring out which experts to trust, including in cases where experts disagree.

difficult to determine when one ought to, say, seek out more evidence on a topic, becoming more epistemically virtuous is one way to become better at determining when seeking out more evidence is necessary. In other words, determining when we need to do more epistemic work is difficult. One way of becoming better at determining when we need to do more epistemic work is to become more epistemically virtuous in general, as the epistemically virtuous agent is a better judge of such things relative to her less virtuous counterparts.

In addition to striving to become more epistemically virtuous in general, there are action-guiding principles we can use to determine which of our specific beliefs we should be especially ready to scrutinize. First, having one's beliefs directly challenged presents one with a relatively easy opportunity to examine the bases of those beliefs. Far from having no idea where to start, one can begin by scrutinizing beliefs that other call into question. Second, as a general rule, the more complicated a matter is, the more one should hesitate to believe that one's beliefs about such topics are sufficiently justified. Thus, when one has the opportunity to examine specific beliefs about complicated topics—especially when those beliefs are directly challenged—one should scrutinize them more carefully. Let's consider a pair of examples to illustrate the point.

First, suppose my friend and I are discussing which dog breeds are recognized by the American Kennel Association (AKC). I claim that Goldendoodles are not among the recognized breeds. My friend expresses some uncertainty and questions the basis for my belief. In response, I report that I've recently checked the AKC's official website and found that the breed is not recognized by the organization. All I have to

do to determine that my belief is justified is simply recall checking the website. Of course, it's possible that Goldendoodles have become officially AKC registered in the time since I checked the website. But if there were any doubt in my or my friend's mind about the truth of my belief, all we would have to do is visit the AKC's official website to verify my claim. The point is that both determining that my belief is justified and verifying the truth of my belief are not complicated affairs.

Contrast this to a case in which a friend and I are discussing the US voting process, and there comes a point in the discussion where she presses me on my beliefs regarding the fairness of the electoral college. Given the complicated nature of the topic, and depending on how strong my beliefs are, I likely should pause to consider what reasons I have for my beliefs, how I came to have them in the first place, and whether my beliefs might be biased in certain ways. For instance, I might ask myself something like the following series of questions. Has the electoral college recently benefited or hindered the political party with which I am affiliated? Would my attitude towards, or beliefs about, the fairness of the electoral college be different had it not recently benefited or hindered my political party? How much do I actually know about the electoral college, including which populations it may have benefited or harmed in the past? On what sources of information did I rely in coming to have beliefs about the electoral college? Are my views based on the testimony of reputable historians, political scientists, and the like, or are they simply based on random bits of information I've encountered in discussions or news clips? And so forth.

Unlike the previous example, determining whether my beliefs about the electoral college are justified could likely be a complicated matter. Furthermore, if there were

any doubt in my friend's mind about the truth of (or strength of justification of) my beliefs, looking into matters to verify the truth of (or strength of justification of) my beliefs would itself be difficult and time consuming. At the very least, it would certainly not be as simple as quickly consulting a website. As a general rule, then, the more complicated the matter, the more likely it is that one ought to examine the bases of one's beliefs about that matter, and perhaps seek out further evidence to both determine whether one's beliefs are justified and hopefully strengthen the justification of one's beliefs.

Of course, examining one's views in this way is not always necessary. For example, we might suppose that I have done my homework—I'm politically informed, and I know quite a bit about the US electoral process both presently and historically, including how the electoral college may or may not have been used to marginalize or benefit certain voters. If I am so informed, I can be more rightly confident that my beliefs about the fairness of the electoral college are justified than if I were not so informed. And yet, *even when I am so informed*, given the complicated nature of the topic, there is likely room for improvement on my views. Thus, even when there is less (perhaps far less) reason to examine my views when I am well-informed on a topic, it still holds that the more complicated a topic is, the more likely that one's beliefs about that topic should be examined, and the more likely it is that further evidence regarding the topic should be sought out.

Third, in line with what I have been arguing throughout, I submit that one should devote what time, energy, and other resources one has to examining the beliefs one has regarding issues that matter non-epistemically. For example, suppose I am not

sure whether I have done enough research on how to budget for a trip well, nor am I sure whether I've done enough research on the ballot options for the upcoming election. Given that what I vote on—and whether I vote responsibly—matters more than how well I create a budget, I should devote my time (or, at least, more of my time) to researching the ballot options. In other words, one should prioritize the scrutinizing of beliefs which concern matters of non-epistemic importance. For, as I have argued, in order to fulfill our non-epistemic obligations, our beliefs about those topics ought to be justified. Furthermore, the more significant an issue or topic is morally, politically, or socially, the more it matters that our beliefs about said topic are justified.

At this point, one might object that even if we restrict our focus to topics which are non-epistemically significant and/or complicated, we still have the problem of knowing when to continue or stop seeking out more evidence relevant to one's beliefs. Thus, the problem of regress does not go away. After all, there is an enormous amount of evidence which could be directly or indirectly relevant to any one of our non-epistemic obligations. Furthermore, as finite creatures, we have limited time, attention, cognitive and other resources to spare on examining our beliefs, and to engage with evidence relevant to our non-epistemic obligations. So there's no way that all of this evidence could be thoroughly examined, or really examined at all. I concede this point, and have two responses to offer. First, the fact that there will always be more epistemic work we in some sense could do is just something we have to accept. Just as our limited capacities constrain what moral and political

obligations we end up having, or the lengths one has to go to meet them,¹¹ so too do our limitations constrain what our epistemic obligations end up being, as well as the lengths we have to go to meet them.

Second, although the non-epistemic significance of certain topics makes it such that we should devote more time and energy to making sure that our beliefs about these topics are justified—which, in turn, invites the problem of regress—so too does the non-epistemic significance of these topics make it such that we ought to eventually put a halt to epistemic work. For making sure that beliefs relevant to our non-epistemic obligations are justified will do little to no good if in doing so we become too paralyzed to act. For instance, although I should do the epistemic work to try to ensure that my beliefs about the issues on the upcoming ballot are justified, I also need to stop investigating the relevant issues and candidates at a certain point if I am to vote. In short: just as the non-epistemic significance of certain topics can threaten to push us up the ladder of regress, so too can they require that we stop advancing up it at some point or other.¹²

In conclusion, to say that it is difficult and time consuming to examine and perhaps strengthen the bases of one's beliefs is different from saying that one cannot know when one should be seeking out more evidence or examining one's beliefs. On the contrary, as I am arguing now (and contrary to what Peels contends), we can

¹¹For instance, if I had unlimited time and monetary resources, I would be morally obligated to do more to help those less fortunate than I currently do. Since I do not, less can be legitimately expected of me, morally speaking.

¹²Of course, this is not to say that once a person puts a halt to epistemic work (so that she can act to meet her non-epistemic obligations) she should never pick up the epistemic workload again. As I've emphasized, topics related to one's epistemic obligations are often complex, and our understanding of these issues is frequently evolving. But needing to put a halt to epistemic work in order to act is consistent with continuing that epistemic work in the future.

have a good idea about when we should be examining the bases of our beliefs: (1) we should examine the bases of our beliefs when our beliefs are challenged, or when we encounter evidence which conflicts with our beliefs; (2) the more complicated a topic is, the more likely it is that we should examine the bases of our beliefs relevant to that topic; and (3) we should prioritize examining the bases of beliefs about topics that are relevant to our non-epistemic obligations. Moreover, becoming more epistemically virtuous in general will both make one a better judge of when one ought to examine the bases of one's beliefs, and it will make one a better judge of the strength of the bases of one's beliefs.

4.4 These are not epistemic obligations

Finally, it may be objected that the obligations I've been discussing are not really epistemic obligations underpinned by social and moral obligations. Rather, they are moral and social obligations that have epistemic components. By way of a first response to this objection, I should point out that it does not matter all that much to me whether we call these "epistemic obligations," or "non-epistemic obligations that have epistemic components." For what we *call* them seems to be a linguistic matter. The reason I have been calling them epistemic obligations is because—to borrow a phrase from Peels—"whether or not we meet them makes a difference to what we believe" (2017, 101). To me (and to Peels), this seems sufficient for counting as an epistemic obligation. If the reader disagrees, she can call them something else. All that matters for my purposes is that we agree that whether or not we meet these

obligations makes a difference to what we believe.

The second point is that it is consistent with the view I have been developing that it will often be difficult to differentiate between different kinds of obligations. We can certainly define what it means for something to be a social obligation, as distinct from a moral obligation, as distinct from an epistemic obligation. In practice, however, it may be very difficult, if not impossible, to treat a particular obligation as strictly (or even primarily) moral, as strictly (or primarily) epistemic, and so on. In other words, arguing over whether something counts as a moral obligation, rather than an epistemic obligation, may be pointless. For the kinds of obligations I've been discussing might be best understood as hybrids.¹³ Although I do not want to take a stand on the issue here, I am open to this possibility.

¹³Conceiving of epistemic obligations this way is something that feminist epistemologists have been advocating for more and more recently. See Fricker (2007) and Gunn (2018) for some discussion.

Chapter 5

Open-mindedness

The purpose of the last several chapters has been to get a better handle on epistemic normativity by way of understanding and appreciating its relationship to non-epistemic normativity. As I explained earlier, an answer to the question of whether and when one ought to be open-minded will be incomplete without first answering the question of why one ought to be open-minded. Understanding the (or a primary) source of epistemic normativity thus puts us in a much better position to understand the extent to which epistemic norms (including “one ought to be open-minded”) should govern our behavior.

The view I’ve been advancing suggests a quite natural answer to the question of when one ought to be open-minded. In particular, moral, social, and political agents are entitled to expect that others be virtuously open-minded with respect to evidence, arguments, and the like that are relevant to their non-epistemic obligations. For, as virtue epistemologists have emphasized, knowledge and justified belief are

the products of intellectual virtue.¹ When others are entitled to expect that one has justified beliefs or knowledge about a certain topic, they are thereby generally entitled to expect that the beliefs one has (or will come to have) about that topic are formed through intellectually virtuous acts, or through the exercise of virtuous intellectual skills, including open-mindedness. Moreover, on my view, the more non-epistemically significant a topic is, the stronger others' entitlement to expect that beliefs about that topic are justified. It is therefore a consequence of this view that, in general, the more non-epistemically significant a topic is, the stronger others' entitlement to expect that one (virtuously) open-mindedly engages with evidence, arguments, and the like that are relevant to that topic.

In Chapter 1, we encountered Fantl's account of open-mindedness, which I argued was misguided on a number of fronts. In this chapter, I offer an original account of open-mindedness which I believe more accurately captures the conceptual core of open-mindedness. On my view, to be open-minded is to be critically receptive to relevant intellectual options. Having an accurate conception of open-mindedness in place will put us in a good position to tie the various aspects of our overarching discussion together in Chapter 6.

5.1 Preliminary remarks

Within the field of philosophy alone, quite a few conceptions of open-mindedness exist. In Chapter 1, we encountered Fantl's. Of course, there are others:

John Dewey: [Open-mindedness] includes an active desire to listen to

¹See, e.g., Zagzebski (1996, esp. pp. 270-72) and Sosa (1991, 276-77).

more sides than one; to give heed to facts from whatever source they come; to give full attention to alternative possibilities; to recognize the possibility of error even in the beliefs that are dearest to us. (1933, 30)

William Hare: Open-mindedness is properly thought of as a kind of critical receptiveness in which our willingness to consider new ideas is guided by our best judgment with respect to the available evidence. (2003, 84)²

Heather Battaly: [W]hat makes one open-minded is the disposition to engage seriously with (relevant) intellectual options. (Battaly, 2018b, 266)^{3,4}

Jason Baehr: 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. (2011, 152)

While each of these definitions has merit, I think there is nevertheless room for improvement. Rather than adopting any of these particular views, then, I will defend an original view of open-mindedness, according to which:

²Here, following Bertrand Russell (1928/2004, esp. p. 130), Hare highlights *critical receptiveness* as being explanatorily central. Hare also characterizes open-mindedness in the following way, which is meant to capture the spirit of the definition quoted above: “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, especially when there is some reason why such evidence and argument might be resisted by the individual in question” (2003, 76).

³In the two papers I cite of Battaly’s, she is primarily interested in closed-mindedness, though she does remark on open-mindedness. Moreover, some of the important points she makes in her discussion of closed-mindedness can be carried over to a discussion of open-mindedness. For example, what counts as an intellectual option, as well as what it means to engage seriously with an intellectual option, will remain the same whether we are discussing closed- or open-mindedness.

⁴Elsewhere, Battaly highlights the importance of being *able* to take seriously relevant intellectual options (2018b, 275-278). Like Baehr (2011), she thinks that one might be willing, but unable, to engage seriously with relevant intellectual options, and that this inability can prevent one from being open-minded.

To be open-minded is to be critically receptive to relevant intellectual options.⁵

This view obviously draws inspiration from some of the accounts listed above, most notably Hare's and Battaly's. The remaining sections of this chapter illustrate the similarities and differences between these two accounts and mine. Before getting to that, I want to mention two stage-setting remarks.

First, I focus my discussion on the virtue of open-mindedness. In doing so, I do not mean to suggest that open-mindedness always counts as a virtue.⁶ It's simply that analyzing the virtue is more conducive to the goal of determining when others are entitled to expect that one be open-minded. Second, since it would not otherwise be obvious, I do want to mention at the forefront that my way of thinking about open-mindedness is also influenced by Dewey, in addition to Hare and Battaly. What I find particularly insightful about Dewey's discussions of open-mindedness is that he situated them within broader discussions of education. For, as a philosopher of education, he was interested in how to best raise young minds to be good inquirers—how to train them in the scientific spirit.

That he emphasized the importance of open-mindedness to education—particularly in the context of cultivating inquiring and scientific spirits in young minds—is, I think, quite insightful, for it suggest two important ways about how we ought to be thinking of open-mindedness. First, open-mindedness is integral to inquiry. All else being equal, when one is inquiring, one ought to be open-minded. The second

⁵As I explain below, being critically receptive to intellectual options will sometimes require first seeking out, and then being critically receptive to, relevant intellectual options.

⁶See Battaly 2018a for an excellent discussion of how closed-mindedness might sometimes be a virtue.

(and related) point is that open-mindedness is not the *only* thing that is integral to inquiry. This is important to realize, especially in light of some of the worries that tend to be raised in reference to open-mindedness—e.g., that open-minded inquiry can be epistemically risky, especially if one is not competent in, say, discriminating between good and bad evidence. Underlying worries like this, however, is the failure to recognize that proper inquiry involves not just open-mindedness, but also other intellectual traits or skills—for example, intellectual humility, competency in discriminating between good and bad evidence, and more. Accordingly, objections such as “open-mindedness can be risky; after all, one doesn’t want to be so open-minded that one’s brain falls out!” are in a sense misguided; underlying them is the assumption that the value of open-mindedness can be divorced from the value of other virtuous traits or skills that are also central to inquiry.

An analogy might help illustrate the point. Consider some factors which contribute to making a person a skillful swimmer. Strength is obviously important, but so too is form. Being physically capable only gets a person so far if she does not know proper stroke techniques, how to take breaths so that breathing does not interfere with her strokes, how often to take breaths, how to control her breathing, and so forth. All of these factors are necessary for making a skillful and fast swimmer. Taken in isolation, though, each of these factors only go so far toward contributing to swimming excellence. Furthermore, focusing on one of these factors in isolation (and at the expense of the others) can sometimes lead to bad results.

Suppose a person was never taught how to swim, but is determined to learn. She decides that the best place to begin is by focusing her energy on building muscle.

Before stepping foot in a pool, she dedicates a few weeks to strength training at her gym in order to build her core, upper-body, and lower-body strength. Eventually, she decides she is ready to give swimming a go. At her next trip to the gym, she heads to the 25-meter pool and begins making her way down a lane. While she is able to make it some distance down the length of the pool, she quickly begins to struggle to stay afloat due to her complete lack of knowledge of stroke technique. A lifeguard ends up having to jump in to save her.

The obvious moral of the story is: while strength is important for swimming, one must also know *how* to swim in order to actually be able to swim. But knowing how to swim is not always sufficient for being able to swim, either. Lack of strength can also lead to bad results. Suppose a person has recently recovered from a serious illness which has caused severe muscle atrophy. Even if this person was previously a good swimmer, she might similarly find herself in need of saving at the deep end of a pool if she attempts to swim without giving herself the time to rebuild muscle.

While the analogy is not perfect, it does help to put us in the right frame of mind for thinking about open-mindedness. Just as certain factors which are necessary for making someone a successful and skilled swimmer can be dangerous when taken to be individually sufficient, so too can certain factors which are necessary for making a person a good inquirer be dangerous when taken to be individually sufficient. In other words, once we realize that we ought to think of open-mindedness as it relates to proper inquiry, and once we recognize open-mindedness as one of a number of traits or skills that are central to proper inquiry, objections such as the one raised above begin to lose their force. After all, it should come as no surprise that inquiry might

not lead to epistemically desirable results when one possesses only one of the virtuous attitudes or skills which is central to proper inquiry. I mention this point both to put us in the right frame of mind for thinking about open-mindedness, and to stave off a family of worries that are often raised in discussions about open-mindedness.

5.2 The objects of open-mindedness

One of the first issues to address is: what exactly is one open-minded toward? Put differently: what are the objects of open-mindedness? On one end of the spectrum are authors who restrict the scope of “object” in their definitions quite a bit. For example, Fantl (2018) is primarily interested in open-mindedness as it pertains to evaluating arguments, and he restricts his definition to fit that aim.⁷ More inclusively, Hare (2003) understands the objects of open-mindedness to be arguments and evidence. On the most inclusive side of the spectrum sit authors like Battaly (2018b), who understands the objects of open-mindedness to be “intellectual options.” According to Battaly, the category of intellectual options includes, but is not necessarily limited to: “sources, questions, ideas, evidence, methods, experiences, inquiries, and goals” (2018b, 265).

As for myself, I understand the objects of open-mindedness to be anything that falls in Battaly’s category of “intellectual options.” Since she does not provide a strict definition of the term, I will do so here (though, of course, much of what I say is just to elaborate on her own discussion of the term). Given that I will be characterizing

⁷That said, he does not deny that one can be open-minded to things other than arguments, such as experiences, ideas, and more.

intellectual options in terms of inquiry, I'll kick off the discussion with that topic.

Following Michael Lynch, I treat inquiry as “the range of epistemic practices we engage in when asking and answering questions, whether banal...or sublime” (2009, 225), and where the point of answering and asking questions is to uncover the truth,⁸ to gain understanding⁹ or knowledge,¹⁰ or to form justified beliefs.¹¹ Of course, not all inquiries take the same general shape. In addition to some being more serious than others, they can also take different objects. Some take evidence (or arguments) as their objects. E.g., “Is this good evidence (or a good argument) for p ?” Or, “Does this argument (or evidence) help me determine whether p ?” Others take ideas as their objects: “Is this a good idea (e.g., for some project)?” Others take methodologies or sources as their objects: “What methodologies or sources are good to employ or rely on?” Still others take questions or *other* inquiries as objects: “Is this a good question to ask?” Or, “Is this a good inquiry to pursue?” Or, “What questions (or inquiries) are worth asking (or pursuing)?”

By way of illustration, let's consider the general ways in which an epistemologist might conduct her research. Before pursuing any particular project or question, she might seriously consider the merits of pursuing certain projects or questions over others. For example, she might consider the merits of pursuing the question, “Can particular social identities limit or enhance epistemic agency?” Or, “Is there such a thing as reasonable peer disagreement?” Of course, this does not mean that she needs to actually pursue all these questions, nor does it even require that she find

⁸See Williams (1973) and Lynch (2009).

⁹See Pritchard (2008).

¹⁰See Kelp (2014), Millar (2011), and Williamson (2000).

¹¹See Davidson (2005), Feldman (2000), and Rorty (1995, 2000).

all of them in the end worthy of pursuit. Her goal here is simply to determine which questions or projects are worth pursuing in the first place.

Alternatively, our epistemologist might inquire about how to best investigate certain questions. Suppose she determines that it is worthwhile to investigate whether or not having particular social identities can affect epistemic agency. With that settled, she might then endeavor to determine which *sources* are worth paying attention to, and which *methodologies* are effective for addressing this question. For instance, she might consider the merits of reading work produced by philosophers of color, work produced by female philosophers, work produced by non-cisgendered philosophers, work produced by white male philosophers, and more. She might also consider the merits of reading work produced in the continental tradition, work produced in the analytic tradition, work which utilizes the methodologies of critical race theory and/or feminist theory, work produced in non-philosophical traditions, and more. When engaging in this kind of inquiry, her goal is not to figure out which questions are worthy of pursuit, nor is it to actually investigate or answer the questions she deems worthy of pursuit. Rather, the goal is to figure out how best to go about answering the questions under consideration.¹²

Finally, our epistemologist might inquire about *evidence* (which can come in the form of testimony), *arguments*, and *ideas*. Suppose she has determined not only that it is worthwhile to investigate whether or not having particular social identities can affect the exercise of epistemic agency, but also that reading classic works in critical race theory is a good way to go about investigating this question. For these reasons,

¹²Admittedly, this line will not always be easy to draw. For which methodology one pursues might influence which questions one deems worthy of pursuit, and vice versa.

she reads W.E.B. Du Bois' (1903) *The Souls of Black Folk*, and Frantz Fanon's *Black Skin, White Masks* (1952/2008). Engaging with these works in this way is to inquire about the ideas, arguments, and evidence contained within these works.

Notice that, though the object of each type of inquiry is slightly different, there is something unifying them all. Not only are they all objects of inquiry. What's more, they all either (1) have propositional content (questions, projects, evidence, ideas, arguments), (2) are sources of propositional content (sources), or else (3) are ways of evaluating propositional content (methodologies). In this sense, all of the objects of inquiry I have discussed so far are *intellectual*. This separates them from what we might call non-intellectual objects of inquiry, which are varied and numerous. For example, one might inquire about places, cars, houses, and more. However, since none of these are related to propositional content in the way just described, they do not count as intellectual. Thus, on my view, intellectual options are:

Intellectual Option: An intellectual option is anything that can count as an intellectual object of inquiry, where to count as "intellectual," the object must (1) have propositional content, (2) be a source of propositional content, or else (3) be a way to evaluate propositional content.

5.3 Relevant intellectual options

The next topic to address is that of relevancy. To begin, it would be useful to to explain why a relevancy restriction is wanted in the first place. In other words, why think that someone who is critically receptive to irrelevant intellectual options would not count as being open-minded? Battaly offers a helpful example to illustrate:

Imagine that a police detective in a remote town is currently investigating the homicide of a local woman—one whose husband is well known to have been convicted of domestic battery on more than one occasion. Suppose our detective transcends her default belief that the husband did it in order to consider the possibility that David Bowie’s ghost did it. Our detective, it turns out, is a super-fan of Bowie. (2018b, 265)

While there are a number of ways we could describe the detective, Battaly points out that we should not describe her as being open-minded. More apt descriptors might include “imaginative, fanciful, or naive” (2018b, 265). Hare would appear to agree. Discussing cases in which similarly far-fetched claims or possibilities are put on the table, he remarks: “in the absence of anything in the way of serious evidence ... receptiveness to such claims, in the sense of according them any probability whatever, is indistinguishable from gullibility” (2001, 44, ft. 21). I think Battaly and Hare are correct to insist that open-mindedness does not involve being receptive to any intellectual option, no matter how improbable or irrelevant that intellectual option might be. For we certainly want to distinguish open-mindedness from traits or qualities like gullibility, naiveté, and (to borrow a term from Dewey) *empty-mindedness* (1933, 30).

This leaves us with the task of specifying what, exactly, makes an intellectual option relevant. One promising way to get a handle on the issue is to consider an example of an intellectual option that is intuitively relevant, and then consider what might explain why the option counts as relevant. Consider:

Vaccinations

Vincent is a new father and has some reservations about vaccinating his infant. (He has heard rumors that vaccinations are dangerous, that they

cause autism, and more.) He asks a few friends in his parenting circle about the matter, and they assure him that there is nothing to these rumors. To back up their claims, they recommend some peer-reviewed papers published in reputable medical journals for Vincent to take a look at. Having good reason to think that both his friends and the papers they recommend are reliable, Vincent decides to follow up on his friends' recommendations.

It is quite clear that the papers Vincent's friends recommend are intellectual options that are relevant to his inquiry. The question is: what, exactly, is it about the papers that makes them relevant? There are at least a couple of options on the table. The first is that these papers are relevant in the sense that the results contained within are accurate. Or, at least, whatever is contained within is, as a matter of fact, likely to aid in Vincent's inquiry. That is, even if the papers do not directly provide him with the answers he is searching for, they may lead him to the answers. Thus, one explanation for why an intellectual option counts as relevant is that it is:

Objectively Relevant: An intellectual option is objectively relevant if it is either at least part of the end product of one's inquiry—e.g., a true answer or explanation—or else if it is likely to aid in inquiry.

While this is a good explanation of what makes the recommended papers relevant, it is not the only explanation to offer. Not only will the papers in fact provide Vincent with the answers he seeks (or else aid him in his inquiry), but also, we might suppose that Vincent has *good reason* to think that the sources his friends recommend are reliable. For instance, suppose he knows that his friends are trustworthy and scientifically literate, that peer-reviewed papers published in reputable

medical journals are typically good sources of information, and so forth. A second explanation, then, for why an intellectual option counts as relevant that it is:

Evidentially Relevant: An intellectual option is evidentially relevant if one has good evidence for thinking that (1) it is likely to be at least part of the end product of one's inquiry—e.g., a true answer or explanation—or else that (2) it is likely to aid in inquiry.

Notice that for an intellectual option to count as relevant in the evidential sense, one needs to have *good* reason for thinking that (1) or (2). One can be mistaken about whether one actually has good reasons, and in these cases, intellectual options one might think are relevant are not necessarily relevant. For instance, suppose that Vincent is researching the risks of vaccinations on the internet and, for some unknown reason, has it in his mind that every 4th article he scrolls past will be a reliable source of information relevant to his inquiry. Taking his method of choosing articles to be reliable, Vincent decides to print out the 4th, 8th, and 12th articles that appear in the search results. Whether or not these articles are relevant in the objective sense, they are certainly not relevant in the evidential sense. Since Vincent's method of choosing articles is entirely arbitrary, and since this is the sole basis he has for choosing these articles, he decidedly does not have good reason for thinking these intellectual options will aid him in his inquiry. Thus, these intellectual options are not evidentially relevant, though he takes them to be.

One way to describe the difference between these two accounts of relevancy is to say that the former captures externalist intuitions about relevance, whereas the latter captures internalist intuitions about relevance. In other words, the former concerns what the world, as a matter of fact, is like, whereas the latter concerns what an agent

has good reason to think the world is like. Now, since the recommended papers appear to be both objectively and evidentially relevant, one might be tempted to think that in order for an intellectual option to count as relevant full stop, it should count as relevant in both senses. I think we should resist this temptation.

To motivate this thought, I want to take a brief detour to discuss the topic of epistemic defeat. Doing so will lend support to my contention that there are different, legitimate senses in which an intellectual option can count as relevant. Broadly speaking, a defeater in epistemology is either a belief or a proposition that can negatively affect the positive epistemic status one's belief that p would otherwise enjoy. Defeaters are generally understood to come in two flavors: there are propositional defeaters, which track externalist intuitions about when beliefs amount to knowledge, and there are mental state defeaters, which track internalist intuitions about justification (and, indirectly, knowledge). On the one hand, a propositional defeater is a true proposition p , where this proposition has the following properties: (1) the agent is unaware of p but (2) if the agent were to become aware of it, the agent would no longer be justified in believing some other proposition q , which she justifiably believes. Importantly, propositional defeaters are typically not understood to undermine justification; they are simply meant to prevent knowledge. In other words, while the agent *would* lose justification for her belief that q were she to learn of the defeater, so long as she is unaware of it, her true belief that q may still enjoy justification. The presence of the defeater simply prevents her justified belief from amounting to knowledge. Consider Alvin Goldman's famous Gettier-style case:

Fake Barn County

Henry is driving in the country side with his son. For the boy's edification Henry identifies various objects on the landscape as they come into view. "That's a cow," says Henry, "That's a tractor," "That's a silo," "That's a barn," etc. Henry has no doubt about the identity of these objects ... Moreover, each object is fully in view, Henry has excellent eyesight, and he has enough time to look at them reasonably carefully, since there is little traffic to distract him ... [However,] unknown to Henry, the district he has just entered is full of papier-mâché facsimiles of barns... Having just entered the district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. But if the object on that site were a facsimile, Henry would mistake it for a barn. (1976, 772-73)

While Henry's true belief that he is looking at a barn is widely understood to count as justified—after all, he is stipulated to have great eyesight, he is not rushed in his judgment that he is looking at a barn, and so forth—he is also widely understood to lack knowledge. This is due to the existence of the propositional defeater: <Henry is in fake barn county>.

On the other hand, a mental state defeater is a belief that p , where this belief negatively affects the justification an agent's belief that q would otherwise enjoy if not for p . Of course, a loss or prevention of justification will typically amount to a loss or prevention of knowledge. Importantly, though, the focus is on a loss or prevention of justification, rather than knowledge. To illustrate, we can return to (a slight adaptation of) Paxson and Lehrer's "Tom Grabit" case, which we discussed in Chapter 1.¹³

¹³See Paxson and Lehrer (1969, especially pp. 228-31).

Tom Grabit

One day I see someone who looks like Tom Grabit steal a book from a local library. My belief is true and justified (we might suppose that I am familiar with what Tom Grabit looks like). However, later that day, Tom's mother—who happens to be a compulsive liar—tells me that Tom is out of state, and that his identical twin brother, Tim, is the one who stole the book. As a matter of fact, Tim does not exist, but I do not know this, nor do I know that Tim's mother is a compulsive liar.

Because I do not know that Tom's mother is a compulsive liar, my justification for believing that Tom stole the book is taken to be negatively affected. For my belief that Tom stole the book to regain the positive epistemic status it once enjoyed, I would have to become aware of the fact that Tom's mother is a compulsive liar, that Tom doesn't exist, or both.

I mention epistemic defeaters primarily because I think that seeing what counts as a defeater sheds light on what counts as a relevant intellectual option. After all, if it is widely accepted that q is an epistemic defeater for p , then plausibly q should also count as a relevant intellectual option. For if a proposition or a belief can negatively affect the epistemic status of one's belief that p , then said proposition or belief should count as relevant to one's inquiry whether p .

Consider my true belief that Tom stole the book from the library. I submit that this is a case in which an intellectual option (i.e., "Tim stole the book") is not objectively relevant, but is nevertheless evidentially relevant. In more detail, it cannot be objectively relevant since Tim does not exist. However, it still counts as evidentially relevant since, absent some reason to think that Tom's mother might be lying, I presumably have good reason to take her at her word: mothers typically know

their children as well as, or perhaps better than, anyone else; mothers typically do not have reason to lie about their children; most people are not compulsive liars; and so forth. Importantly, what this example illustrates is that there can be intellectual options that are relevant in the evidential sense, but not relevant in the objective sense. Moreover, the intellectual option does not fail to be relevant full stop even though it is not objectively relevant, for it still manages to negatively impact the epistemic status of my belief.

There are also cases in which an intellectual option is objectively relevant, but not evidentially relevant. Consider again *Fake Barn County*, in which a propositional defeater exists for Henry's belief that he is looking at a barn. In this case, Henry is unaware that he is in fake barn county, so he has no reason to think that the possibility that he is in fake barn county is relevant at all to his belief that he is looking at a real barn. Thus, <Henry is in fake barn county> is not evidentially relevant. However, it is objectively relevant since it is a matter of fact that he is in fake barn county, and this fact is relevant to his belief that he is looking at a barn. This is evidenced by the fact that the defeater undermines Henry's knowledge that he is looking at a barn.

So far, we have identified two ways in which intellectual options can be relevant. Though they may be relevant in both senses, they may also be relevant in one sense, but not the other. While we have made progress, this picture is not yet as nuanced as it needs to be. This is because there is another type of defeater sometimes discussed in the literature—one which doesn't quite fall either into the category of mental state defeaters or propositional defeaters, but rather lies somewhere in between. If

we assume such defeaters are legitimate—and I think that we should—this would suggest that there is another sense of relevancy we have yet to capture. Consider Gilbert Harman’s example:

Political Assassination

A political leader is assassinated. His associates, fearing a coup, decide to pretend that the bullet hit someone else. On nationwide television they announce that an assassination attempt has failed to kill the leader but has killed a secret service man by mistake. However, before the announcement is made, an enterprising reporter on the scene telephones the real story to his newspaper, which has included the story in its final edition. Jill buys a copy of that paper and reads the story of the assassination. What she reads is true and so are her assumptions about how the story came to be in the paper ... Jill has justified true belief ... But she does not know that the political leader has been assassinated. For everyone else has heard about the televised announcement. They may also have seen the story in the paper and, perhaps, do not know what to believe; and it is highly implausible that Jill should know simply because she lacks evidence everyone else has ... Her knowledge is undermined by evidence she does not possess. (1973, 143-44)

If Harman is right that falsehoods that pollute the epistemic environment can undermine one’s knowledge, then I think we ought to count such falsehoods as relevant intellectual options. Notice, however, that <a secret serviceman was killed> is not evidentially relevant for Jill (she has no reason to think that it is true), nor is it obviously objectively relevant. If this is correct, then we need a different explanation for why it counts as relevant. What is going on in this case, it seems, is that most members of Jill’s community widely believe that it is true, and presumably with good reason. After all, they presumably believe things on the basis of

news broadcasts quite regularly, and under normal conditions, this is a reliable way of gathering information. It just so happens that their environment is not normal. Furthermore, their believing this with good reason is arguably what prevents Jill from having knowledge.¹⁴ It seems, then, that there is a third sense of relevancy to take into account:

Socially Relevant: An intellectual option is socially relevant if it is widely believed (and with good reason) by member's of one's epistemic community that the intellectual option (1) is likely to be at least part of the end product of inquiry—e.g., a true answer or explanation—or else that (2) it is likely to *aid* in inquiry.

As was the case with evidential relevance, adding the qualification “with good reason” is important. For just as one can be mistaken about whether one has good reasons for thinking that an intellectual option is likely to be the product of inquiry, or else likely to aid in inquiry, so too is it possible that members of one's epistemic community can be similarly mistaken. Suppose that members of Jill's epistemic community came to believe that the secret serviceman was assassinated, not because they saw a TV report, but because they all happened to dream it and then came to form their beliefs on this basis. Presumably, the widely shared belief that the secret serviceman was assassinated rather than the political leader would not prevent Jill from knowing the political leader was assassinated, since Jill's community member's beliefs were not formed on the basis of good reasons.

Let's take stock. We have identified three distinct ways in which an intellectual option may count as relevant: it may be objectively relevant, evidentially relevant, or

¹⁴Consider: if their beliefs were not formed on the basis of good reasons, then it would be unclear whether Jill in fact lacks knowledge.

socially relevant. Under normal circumstances, an intellectual option that counts as relevant in one sense will often count as relevant in the other senses as well. However, as our discussion has revealed, an intellectual option will not always be relevant in all three senses, but may only be relevant in just one or two senses. That said, if we find all three defeaters cases compelling—i.e., if we think that the epistemic status of the highlighted beliefs in all three cases was negatively affected—then we should treat all three senses of relevancy as individually legitimate.

This discussion has also revealed that intellectual options are not relevant full stop; rather, they are relevant *relative to* specific inquiries. Recall Vincent. The papers his friends recommended are both objectively and evidentially relevant, relative to his inquiry about the connection (or lack thereof) between vaccinations and autism. They are plausibly socially relevant as well. After all, members of Vincent’s epistemic community presumably have good reason to think that these papers are relevant to whether vaccinations cause autism. However, these papers would decidedly not count as relevant relative to some other inquiries. For example, suppose instead that Vincent wants to determine whether or not co-sleeping is safe. The papers recommended by his friends would not, as a matter of fact, aid in this inquiry; nor would Vincent have good reason for thinking they would; nor would members of Vincent’s epistemic community have good reason for thinking that they would. So, these papers would not count as relevant in any of the three senses relative to this inquiry.

Combining everything that has been said thus far, we have the following definition:

Relevant Intellectual Option: An intellectual option counts as relevant relative to inquiry *I* just in case it is either evidentially relevant, objectively relevant, or socially relevant to *I*.

While we have certainly made progress, there are a couple more qualifications to add, for what counts as both evidentially relevant and socially relevant will be relativized even further. Let's start with evidential relevancy. Notice that what counts as good reasons for one person will not necessarily count as good reasons for another. Take any two epistemic agents. Not only may they have different evidence available to them, but also—due to potential differences in upbringing, values, background beliefs, and more—they might end up weighing the same evidence very differently. Consider:

Herbal Remedies

Veronica strives to use “natural” remedies to treat physical ailments whenever possible. She was raised by parents that strongly believed that physicians are too eager to prescribe medications, and that many physical ailments can be treated using herbs, oils, and the like. Sheena, in contrast, is quite suspicious of natural remedies. Both of her parents are physicians, and they instilled in her a distrust of natural medicine. As a matter of pure coincidence, both Veronica and Sheena have cuts that are beginning to look infected, and are trying to determine whether they should go see a doctor. While browsing on the web, they both happen upon an article titled “Herbal Remedies for Bacterial Infections.” Given their differences in background beliefs, values, and their ways of weighing evidence, this article counts as a evidentially relevant for Veronica but not Sheena.

Notice that, given that their inquiry is the same, the article will either be objectively

relevant for both, or not. The point is simply that whether the article counts as evidentially relevant for them depends on features unique to them as individuals.

Likewise, whether an intellectual option counts as socially relevant will depend on features unique to the agent's epistemic community. Of course, how we determine what an agent's epistemic community is an issue in and of itself. For it is plausible that what counts as an agent's epistemic community can vary depending on the inquiry in question. For example, if I am inquiring into how best to interpret Kant's *Groundwork of the Metaphysics of Morals*, my epistemic community is presumably the philosophical community, or perhaps even something narrower. If, however, I am inquiring into the pros and cons of a certain federal regulation, my epistemic community would include at least my fellow citizens, and plausibly more. Thus, whether an intellectual option counts as socially relevant relative to an agent's inquiry will be contextually determined. The point is simply that, whatever one's epistemic community happens to be relative to one's inquiry *I*, what counts as socially relevant will be determined by features of that epistemic community.

5.4 Critical Receptivity

The next order of business is to specify what it means to be critically receptive to relevant intellectual options. I model this portion of my account of open-mindedness after Hare's. As I see it, there are two primary strands running through Hare's view of serious engagement:

The Socratic Strand

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, especially when there is some reason why such evidence and argument might be resisted by the individual in question. (2003, 76)

The Russellian Strand

Open-mindedness is properly thought of as a kind of critical receptiveness in which our willingness to consider new ideas is guided by our best judgment with respect to the available evidence. (2003, 84)

I'll begin by discussing the Socratic strand. I'll then turn to the Russellian strand, explaining both how it relates to the Socratic strand, as well as why I ultimately opt to characterize open-mindedness in terms of critical receptivity.

5.4.1 The Socratic strand

Hare considers Socrates to be an exemplar of open-mindedness, most notably on account of his “insistence that we must go where the wind of the argument takes us” (2009, 7).¹⁵ Hare (2009) identifies three hallmarks of open-minded engagement in the Socratic style. I will address each in turn.

The first hallmark

First...to follow the argument in the Socratic sense is to make reason and evidence the linchpin of serious inquiry and reflection. Socratic inquiry,

¹⁵Hare borrows this from Plato's Republic 394d.

like open-minded inquiry, seeks, and is guided by, the best reasons to be found. (2009, 7)

In other words, there is a motivational component to the virtue of open-mindedness: namely, what's driving the open-minded person in inquiry is pursuit of epistemic goods, such as understanding and knowledge. This fits in quite well with my contention that the objects of open-mindedness are intellectual objects of inquiry. For, assuming that acquiring epistemic goods is the goal of inquiry, and that the only objects of open-mindedness are intellectual objects of inquiry, then when one is open-minded, one seeks to acquire epistemic goods. This reveals that the relationship between epistemic goods and open-mindedness is, at least in a loose sense, bidirectional: (1) When we are open-minded, we seek epistemic goods (motivational component), and (2) when we inquire, we ought to be open-minded (because, all else being equal, being open-minded will get us closer acquiring epistemic goods).

It might be objected that (2) presupposes that we will be able to recognize good and bad intellectual options as such. After all, if we are not able in a given instance to, say, tell good reasons apart from bad, then engaging with these reasons will not reliably help us acquire epistemic goods. It may even take them further out of our reach. It is certainly the case that one's ability to accurately assess intellectual options can be hampered for one reason or another. In some cases, this will be because an agent has not epistemically matured, as it were. In other cases, an agent's epistemic environment may be polluted or hostile, which means that (perhaps through no fault of her own) it may be difficult for her to reliably evaluate relevant intellectual options. Indeed, it may be difficult to determine in such environments

which intellectual options are relevant in the first place.

However, it was in anticipation of worries such as this one that I emphasized early on that the value of open-mindedness cannot be divorced from the value of other intellectual virtues central to inquiry. So, the presumption is simply that, when there is not a deficit of other intellectual virtues, being open-minded typically puts one in a better position to acquire epistemic goods. This is what's at the heart of the first hallmark.

The second hallmark

Second, to follow the argument is to track the fortunes of the inquiry, taking into account alternative and conflicting possibilities, looking at all sides impartially, raising objections and considering replies, paying careful attention to what is to be said for and against the proposition, and trying to ensure that nothing germane to the issue is ignored or excluded. (2009, 7)

There are three things that stand out to me about this second hallmark. The first does not so much speak to what it means to engage open-mindedly with an intellectual option, but rather to which intellectual options we ought to engage with. It is worth commenting on briefly.

First, Hare thinks that one ought to engage with a range of intellectual options relevant to the inquiry. One must not only consider reasons for accepting a proposition (to use Hare's example), but also consider reasons against accepting it. Now, as Hare puts it, one must try "to ensure than nothing germane to the issue is ignored or excluded." In other words, one must aim to engage with *all* intellectual options

relevant to the inquiry. This seems quite demanding, because there will potentially be a host of intellectual options relevant to any given inquiry. However, it is not as demanding as it might first appear—at least not on my preferred reading. While open-mindedness might demand that we in some sense aim to engage with every intellectual option relevant to an inquiry, open-mindedness does not necessarily require that we actually engage with all relevant intellectual options. Practical, moral, and perhaps other considerations weigh against doing so.

On account of this, it's quite plausible that we are regularly excused for not engaging with all intellectual options relevant to our inquiries. This is assuming, of course, that a plausible constraint on a theory of open-mindedness is that it should be possible for someone to be, or at least closely approximate being, virtuously open-minded. In light of all of this, I suggest we modify Hare's claim in the following way: With practical and other constraints in mind, we ought to engage with the intellectual options that are most relevant to our inquiry. Like many other things, relevancy comes in degrees. And so, it may very well be good enough to engage with the intellectual options most relevant to an inquiry, given that we have limited time and mental energy to devote to inquiry.

The second point to make about this second hallmark—and this does speak directly to what it means to engage open-mindedly—is that we ought to devote time and careful attention when we engage with relevant intellectual options. It is not enough to briefly consider them. Rather, we must evaluate them carefully and thoroughly.¹⁶ Although there is no set amount of time one must spend evaluating an

¹⁶In suggesting that open-minded engagement involves careful and thorough evaluation of relevant intellectual options, I do not mean to conflate open-mindedness with either careful evaluation

intellectual option, carefully and thoroughly evaluating it often does take some time. At the very least, one must not be hasty in one's evaluation. One caveat to make is that, as was discussed above in reference to the *number* of intellectual options one engages with, it is plausible that here too some leeway is permissible. When an epistemic agent evaluates the intellectual options most relevant to her inquiry, practical considerations may weigh against her evaluating them all as carefully as she possibly can. That said, she should strive to evaluate them as carefully as she can, given practical and other constraints.

The third and related point is that one must strive to be *impartial* in her evaluation. As Hare writes elsewhere, open-minded inquiry involves a resistance to biases and drawing hasty conclusions:

[G]enerally, the open-minded outlook seeks to avoid resistance to evidence, hasty conclusions, wishful thinking, bias, prejudice, pressure to conform, and every factor that compromises our ability to get to the truth of the matter. (2009, 6)

Surely, there is only so much we can do to prevent our judgments from being clouded by bias, prejudice, and the like. Still, we have to strive to the best of our ability to not succumb to such influences, and to not be too quick in forming and revising views. Suppose I know that I am biased against Republican news sources. When I read an article written by a Republican, I ought to remind myself that the mere fact that it was written by a Republican is not grounds for taking it less

or thorough evaluation. To borrow a line from Baehr, while we can treat these virtues or skills as distinct, we can nevertheless treat them as "facilitating" one another (2011, 155-56). This sits well with understanding open-mindedness as one of numerous factors which contribute to successful inquiry.

seriously. If I end up rejecting something that was said in the article, it ought to be because I have good and independent grounds for doing so.

One final remark: I think that Hare's third point is weaker than it should be. On his view, open-mindedness merely requires that one *seeks* to avoid biases, prejudice, and other such similar factors. However, putting it this way suggests that being open-minded is compatible with failing to avoid these factors. I think this is mistaken. Like Battaly, I think that sometimes a person's biases, prejudices, and the like will be strong enough to prevent her from accurately assessing evidence and other intellectual options, even if she tries hard to do so.¹⁷ To illustrate, suppose that, try as I might, I cannot help but think that the Republican article is rubbish simply because it was written by a Republican. If I hadn't known who the author was, my evaluation of it would have been very different. Unfortunately for me, my inability to overcome my biases prevents me from engaging open-mindedly with that relevant intellectual option. Even though I try to be open-minded, I fail to be.

The third hallmark

Finally, to follow the argument is to be willing to follow it through to its conclusion, to be prepared to accept whatever results from the inquiry.
(2009, 6)

In other words, it is not enough for one to take time to evaluate or consider an intellectual option, as well as strive to mitigate the affects of biases, prejudices, and the like. In addition to this, one must be willing to appropriately respond to the intellectual options one evaluates in one's inquiry. Fantl (2018) also emphasizes

¹⁷See Battaly (2018b, 275-78).

the importance of being willing to follow the argument wherever it may go. As he elaborates, and as is implicit in Hare's own view, if, for example, one carefully evaluates an argument, judges its merits and shortcomings as objectively as possible, and finds it to be compelling and flawless, then she ought to be "appropriately moved" by it.¹⁸

Following a suggestion from Fantl, it might help to put the point negatively.¹⁹ Imagine that Paula is an atheist, that she is biased against religious folk, and that she is presented with an argument by a religious person. Paula is aware of her biases, and she does her best to keep them in check. For example, she repeatedly reminds herself that the fact that a religious person is offering the argument in no way affects the strength of the argument. She carefully evaluates the argument, and is unable to detect a flaw. If Paula resists a change in her beliefs (or if her beliefs simply fail to be affected by the argument), despite the fact that she is unable to find a flaw, then Paula is not open-minded. The point is simply that, if engaging with an intellectual option should result in a change in one's beliefs, but one's beliefs are not affected as a result of engaging with that intellectual option, then one is not open-minded with respect to that intellectual option.

Of course, being open-minded is compatible with one's doxastic states not being *immediately* affected by an intellectual option. So, if the reason Paula's doxastic states are unaffected by the argument is that she wants to do more research before reaching any conclusions, this would be compatible with her being open-minded. For,

¹⁸Oddly enough, Fantl reads Hare as omitting the requirement that one's doxastic states ought to be responsive to intellectual options in this way. I think Fantl simply misreads Hare.

¹⁹Fantl illustrates using a different example. For his discussion, see (2018, 11-12).

in this re-imagined case, she has judged that there is more relevant evidence that bears on the topic, and that she should engage with it before drawing any conclusions. That said, if the research she undertakes later on corroborates the conclusion of the argument, then Paula's beliefs should be affected at that point.

Summing up

To sum up, open-mindedness according to the Socratic strand is characterized by:

1. Seeking epistemic goods;
2. Evaluating the intellectual options most relevant to one's inquiry (one must be thorough and careful in this evaluation);
3. An active (and to no small degree successful) effort to prevent one's judgment from being clouded by bias, prejudice, and the like;
4. And appropriately adjusting one's beliefs in response to engaging with relevant intellectual options.

5.4.2 The Russellian strand

With the Socratic strand laid out, we can turn our attention to the second line of thought influencing Hare's account of open-mindedness:

The Russellian Strand

Open-mindedness is properly thought of as a kind of critical receptiveness in which our willingness to consider new ideas is guided by our best judgment with respect to the available evidence. (2003, 84)

According to Russell, genuine inquirers must strike a fine balance between openness and scrutiny. Put differently, genuine inquirers must be at the same time (1) willing and able to be receptive to relevant intellectual options relevant to one's inquiry, and (2) be resistant to having one's doxastic states be affected by those intellectual options until they have been carefully examined.²⁰

Here is how I take the Russellian strand to be related to the Socratic, as well as why I prefer to characterize open-mindedness in terms of critical receptivity. In the first place, I submit that the "receptivity" aspect of the Russellian strand can capture everything the Socratic strand has to offer, but that it can also do more. In other words, we should understand "being receptive" as being willing to take relevant evidence and argument into account in forming or revising our beliefs and values, especially when there is some reason why such evidence and argument might be resisted by the individual in question. In addition to this, the Russellian strand can capture more than the Socratic strand.

In the first place, there are advantages to emphasizing the importance of receptivity. In particular, it allows us to accommodate the fact that a person can count as open-minded even if she does not engage with every intellectual option relevant to her inquiries. Here is what I mean by this. I am claiming that when one is open-minded, she is (critically) receptive to relevant intellectual options. It is important to remember, however, that practical and other constraints will place limitations on how many intellectual options one actually can engage with. In other words, being receptive is understood as maintaining an openness of mind, keeping in mind that—

²⁰See Hare (2001, 40).

due to cognitive, temporal, emotional, and other limitations and considerations—one can only engage with so many intellectual options at a given time.

If it were not for all of the limitations and constraints mentioned, she would engage with those intellectual options. At the same time, she cannot engage with all intellectual options relevant to all of her inquiries, and so the number of intellectual options she “receives,” as it were—i.e., the number of intellectual options she actually engages with—will be far fewer than the number of relevant intellectual options she is receptive to. We can thus distinguish between receptivity on the one hand, and receiving, so to speak, on the other. I submit that as long as a person is receptive in the sense that she would engage with all relevant intellectual options were these defeating factors not in play, then she can count as open-minded.

This aspect of my account marks one point of departure from Battaly, who argues that it will sometimes be necessary to fail to be open-minded if we hope to get on with our lives.²¹ On her view, being open-minded requires one to seriously engage with relevant intellectual options. Depending on one’s inquiry, there could be a plethora of relevant intellectual options, and it would be virtually impossible to seriously engage with them all. Consider students who need to finish their dissertations. At some point, they will have to refrain from seriously engaging with intellectual options relevant to their projects, for if they engaged with every relevant intellectual option, they would never make any progress. Finishing projects sometimes demands closing ourselves off to relevant intellectual options. Alternatively, if I am trying to learn about some complex topic, such as human contribution to climate change, I cannot

²¹See Battaly (2018a, 36-37).

seriously engage with every relevant intellectual option, for there is a multitude of these. It will often be the case that seriously engaging with every intellectual option relevant to all of our inquiries will be impractical, or worse, impossible. On her view, one is being closed-minded toward the relevant intellectual options with which one does not seriously engage.

On the critical receptivity view, however, failing to engage with certain intellectual options can be compatible with being open-minded. We'll begin by considering what Hare has to say on this issue. Though he does not articulate the point exactly this way, he argues that ignoring certain relevant intellectual options can actually be compatible with being open-minded. His reason for this is that he thinks that "[o]pen-mindedness operates at different levels" (2003, 82). In more detail:

It is possible...to ask in an open-minded way about the wisdom of turning to alternative possibilities now compared with the wisdom of staying the course. It is simplistic to suggest that open-mindedness automatically dictates that one's present inquiries must be abandoned or one's attention diminished. Open-minded reflection might well suggest that there is compelling reason to continue one's present investigation and to defer a more detailed examination of criticisms and objections even if we suspect that these may have some merit. (2003, 82)

He would thus maintain that being open-minded is compatible with refraining from exploring every argumentative avenue and addressing every objection, *so long as* one has open-mindedly considered whether engaging with those options or ignoring them is the wiser course of action.

I am inclined to agree with Hare on this point. We can be *receptive* to all kinds of relevant intellectual options. But, given our limited capacities, in practice we can

only “receive” so many. So long as I engage with a good deal of the intellectual options most relevant to my dissertation, and so long as I engage with some of the higher-order intellectual options (e.g., “Should I consider every objection that could be raised against my view?”), then I can still count as being open-minded, even if I do not engage with every relevant intellectual option at the lower level (e.g., “Do objections $X_1 - X_n$ pose problems for my view?”).

The second advantage to defining open-mindedness in terms of critical receptivity goes back to a criticism I had of Fantl in Chapter 1—namely, that he failed to accommodate the fact that the degree to which one’s beliefs should be influenced by relevant intellectual options is a function of the totality of one’s evidence. By defining open-mindedness in terms of critical receptivity, however, we can easily accommodate this fact. For the level to which one should be receptive to relevant intellectual options is determined in part by the evidence one has that those intellectual options are conducive to the attainment of epistemic goods. Recall Carla and Carla*. Carla had much stronger reasons than Carla* for believing that human activity is a major contributor to global warming. Accordingly, I submitted that when presented with contrary evidence, Carla’s doxastic states should be affected to a lesser degree than Carla’s*. Highlighting the “critical” aspect of critical receptivity allows us to accommodate this. Carla should be more critical and less receptive to the contrary evidence as compared to Carla*.

In the next section, I discuss what it means to *fail* to be critically receptive to relevant intellectual options. Doing so will help shed further light on my account, given that that some interesting features of open-mindedness come out in such a

discussion.

5.5 Failing to be Critically Receptive

In explicating the nature of closed-mindedness, Battaly explores various ways in which one can fail to be open-minded. The discussion in this section will be modeled after hers.

5.5.1 Failing to be receptive at all

As Battaly points out, there are at least a handful of ways in which one can fail to be open-minded. For one, a person can fail to be critically receptive to a relevant intellectual option because she fails to be receptive to it at all. Within this category, there are a number of subcategories.

First, and perhaps mostly obviously, one can fail to be receptive to relevant intellectual options because one *purposefully ignores* them:

Elitist Ellen

Ellen is working on a paper concerning epistemic injustice, but refuses to read any relevant literature that is not written by a philosopher. In her opinion, only philosophers make real contributions. When Ellen's friend recommends a paper written by a sociologist, Ellen does not read it—in fact, she refuses to even skim it—solely because it was written by a non-philosopher. Ellen is not critically receptive to that intellectual option.

Ellen ignored a relevant intellectual option because of her elitist and mistaken assumptions about who makes valuable contributions and who does not. One way of

diagnosing the problem is to say that she is too critical, and not receptive enough. However, thinking that one has nothing to gain from engaging with relevant intellectual options is not a necessary condition for counting as ignoring them. A person might also ignore a relevant intellectual option because she knows the option is in tension with what she believes, and she might not want to run the risk of coming to doubt one or more of her current beliefs. In any case, when one ignores a relevant intellectual option, one knows that they exist but fails to be receptive to them.

One might also fail to *seek out* relevant intellectual options:

Essential Oils Mary

Mary believes that using combinations of essential oils can cure most illnesses that are normally treated with antibiotics. Deep down, she suspects that there may be scientific research out there that would be in tension with her beliefs, but she does not seek it out because she is comfortable with her beliefs and does not want to run the risk of her beliefs changing after engaging with the research.

Mary does not ignore the relevant intellectual options in question (i.e, scientific studies which are in tension with her beliefs), since it is not as though she is made aware of them and then chooses not to engage with them. Rather, she simply fails to seek out whether there are any such intellectual options.

Third, one may fail to be open-minded because she fails to *generate* relevant intellectual options. To illustrate, we can modify Battaly's detective example:

Detective

A small-town detective is working to solve a woman's murder. She considers and weighs the evidence she has before her, and comes up with

the hypothesis that the woman's daughter is the murderer. Now that she has a suspect in mind, the detective does not attempt to come up with further suspects—i.e., she fails to generate alternative hypotheses of who the murderer might be—but rather is content to stick with her first hypothesis, despite the fact that the evidence did not conclusively incriminate the daughter.

This detective fails to be receptive to relevant intellectual options, for she fails to generate alternative hypotheses. Given that the evidence did not conclusively incriminate the daughter, the detective should generate (and then critically evaluate) alternative hypotheses before reaching any decision.

Battaly (2018a) also argues also one can fail to be open-minded simply in virtue of being ignorant of the existence of certain relevant intellectual options. One example she takes to illustrate this point is borrowed from Wayne Riggs:

Oblivia

Oblivia is a stereotypical “ugly American” who travels abroad. She sees people in other countries acting in ways that disclose values and belief systems different from her own, but she is oblivious to these disclosures. It simply fails to occur to her to explain the sometimes puzzling-to-her behavior of these other people in terms of a different belief or value system. She simply experiences them as behaving “strangely.” Since she perceives no alternative cognitive standpoints, her willingness and ability to seriously consider such standpoints is rather moot. She will be unable to actually consider those standpoints because she is unaware of their existence. What Oblivia lacks is a certain “sensitivity” to various kinds of cues that there is a cognitive standpoint on some matter that is distinct from her own. (Riggs 2016, 24)

What's preventing Oblivia from being open-minded is obviously her obliviousness;

it acts as an internal block, so to speak, preventing her from engaging with relevant intellectual options. However, Battaly thinks that factors other than obliviousness can prevent one from being open-minded—indeed, even make one *closed*-minded—even when one is willing to engage with relevant intellectual options. For example, one might be in an epistemically hostile environment—i.e., an epistemic environment in which intellectual options are restricted, controlled, or worse, disappeared—and thus be ignorant of certain relevant intellectual options. Being in such environments, she argues, can actually make one closed-minded, even when one is willing to engage with the relevant intellectual options she cannot access. Of such a person, she writes:

[W]e can say that he is not *disposed* to be closed-minded, provided that he *would* engage with intellectual options were he in a hospitable environment. But ... we can also say that he *is* closed-minded relative to his current environment. This point is important. Our environments can make us closed-minded.²² (2018b, 275-6, original emphasis)

Here is one place where I must part ways with Battaly. For one, I think that—all else being equal—merely being ignorant of relevant intellectual options does not suffice to make one closed-minded. As I will argue, one can actually count as being open-minded to relevant intellectual options of which one is currently ignorant. And, more generally, failing to engage seriously with relevant intellectual options does not suffice to make one closed-minded. I will address each point in turn.

²²Battaly notes that someone who is prevented through no fault of her own from engaging with relevant intellectual options may not be blameworthy for failing to be open-minded (assuming failing to be open-minded is something a person deserves blame for in the first place) (2018b, 276).

Merely being ignorant of relevant intellectual options

Sheltered Shelly

Shelly lives in a secluded religious community, one which censors what kind of material she learns in school. As far as she knows, creationism is the only viable explanation for how biological life came to be as it is. She thus fails to engage seriously with Darwin's theory of evolution in her biology class—or anywhere else, for that matter—for she does not know that this relevant intellectual option exists. However *were* Shelly to learn about Darwin's theory, she *would* engage seriously with it.

Battaly would say that Shelly is disposed to be open-minded, but, relative to her current epistemic environment, she is closed-minded. I resist this characterization, again, because I take receptivity to be explanatorily central to open-mindedness. When we think of open-mindedness in terms of receptivity, we see that being open-minded to certain relevant intellectual options is compatible with being ignorant of their existence. To repeat: on my view, one is open-minded when one's mind is (critically) *open*. This means that what makes one open-minded is partially agent-determined. If she is critically receptive to relevant intellectual options she is open-minded. This is so even if the relevant intellectual options to which she has access are few.

To see the point, it might be useful to distinguish between different ways in which one can be prevented from being receptive to relevant intellectual options. As I see it, there are three ways in which relevant intellectual options might be blocked from one's mind, so to speak. First (1) they may be externally blocked. This can happen when an agent's epistemic environment is controlled, polluted, or hostile, and so

the relevant intellectual options she has access to are limited. Second (2) they may be internally blocked. This happens when implicit biases, intellectual vices, and the like prevent one from being critically receptive to relevant intellectual options. Third (3) they may be both internally and externally blocked. This occurs when both an agent's environment is controlled, polluted, or hostile, and where biases or vices prevent her from being critically receptive to both (i) relevant intellectual options she still can encounter in her epistemic environment, and (ii) relevant intellectual options she would encounter were her environment not controlled, hostile, or polluted.

Cases of (2) are the most obvious. Oblivia, for example, counts as someone who cannot seriously engage with intellectual options because the intellectual options are blocked by internal factors. Her environment is certainly not to blame. There are plenty of relevant intellectual options with which to engage, and there's no indication that any outside factors are preventing her from engaging with them. Instead, her underdeveloped intellectual character has made it such that she does not recognize relevant intellectual options (or "distinct cognitive standpoints") as such when she encounters them.

Cases of (3), though less obvious, are sadly not unfamiliar. To borrow Battaly's examples (2018b, 276), Hitler's *Jugend*, as well as children raised by the KKK or ISIS, are prime examples of agents who are both internally and externally blocked from being critically receptive to relevant intellectual options. On the external side of things, their exposure to relevant intellectual options is very limited and controlled. On the internal side, the values and intellectual dispositions that have been ingrained in them—such as conformity, obedience, loyalty, and distrust of out-group

members—prevent them from being receptive to many of the relevant intellectual options they do happen to encounter.²³

Cases of (1), though even less obvious than cases of (3), do exist. To illustrate, let's return to Shelly. She is, by hypothesis, not internally blocked, for we've stipulated that, if she were to encounter Darwin's theory of evolution, she would seriously engage with it. What's preventing her from engaging with Darwin's theory is that her epistemic environment is controlled; in her community, creationism is the only intellectual option concerning the origins of biological life to which she has been given access. However—and this is important—while she cannot engage with it since she does not know of its existence, she is critically receptive to it, for she *would* engage with it were she to encounter it.

At this point, the reader might wonder how Shelly could not be internally blocked as well. After all, communities that control what intellectual options their members have access to surely aim to instill certain values in their community members. For example, community leaders might aim to instill in community members values and dispositions such as conformity, a reluctance to question authorities, and more. If this is so, how could it be the case that someone like Shelly is externally blocked, but fails to be internally blocked (at least with respect to intellectual options like

²³It is worth noting that, when one's epistemic environment is controlled, polluted, or hostile, which intellectual options will be difficult to access will depend on the context. For example, children raised in secluded religious communities might never be exposed to Darwin's theory of evolution because it competes with creationism. However, because learning chemistry poses no obvious threat to religious doctrine, they will likely be given the same basic introduction to chemistry that we might expect any children their age would get. Which relevant intellectual options are accessible or not in a controlled, polluted, or hostile epistemic environment will often depend on (1) the reasons for why the environment is the way it is in the first place, and (2) who, if anyone, has control over the environment.

Darwin's theory)?

It's possible that cases like Shelly's are somewhat unlikely. Nevertheless, I think they are not implausible. To see why, let's build some more details into the case. Suppose Shelly is naturally a curious and independent thinker. Suppose also that her educators encourage her curiosity and critical tendencies in some domains. For example, when Shelly gets excited after learning about microorganisms in school, her biology teacher convinces her parents to purchase a microscope for her. Despite the fact that her environment is controlled, Shelly develops into someone who is disposed to be critical and curious, not just with respect to certain domains, but in general. This is what makes it the case that if Shelly *were* to learn about Darwin's theory, she *would* engage seriously with it, even though it is the very sort of intellectual option she is meant to reject.

Perhaps the reader is still unconvinced. One might worry that, even if Shelly is to some degree critical and curious, it's implausible that she would be entirely unaffected by her environment. Given the way she was raised, she will at least be *somewhat* reluctant to question religious teaching. So, it's not the case, the worry goes, that she really would be receptive to Darwin's theory of evolution were she to encounter it. I have two things to offer in response to this worry. First, I would like to point out that open-mindedness comes in degrees. So, even if Shelly would not be *as* critically receptive to Darwin's theory of evolution (were she to encounter it) as she might be had she been raised in a different epistemic environment, this does not mean that she would not be critically receptive to it at all. Second, we can simply change the example:

Julie the Journalist

Julie is an American journalist investigating the Syrian conflict. She was raised in a normal epistemic environment, and is a paragon of open-mindedness. Due to a very unfortunate turn of events, Julie finds herself in a hostile epistemic environment—while in Syria, she is captured as a political prisoner, and the information she is fed while a prisoner is both misleading and restricted. When she inquires about her release, her captives do not always provide her with accurate information, and they fail to mention certain information that is relevant to her inquiries.

Does the mere fact that Julie is in a hostile epistemic environment make it impossible for her to be open-minded? I submit that it does not. Allow me to elaborate. When it comes to the intellectual options she does not have access to, she can still remain critically *receptive* to them, in the sense that were she to be given access to them, she would critically evaluate them. Placing the emphasis on the “receptivity” aspect of critical receptivity helps us understand why Julie can still count as being open-minded, even with respect to the relevant intellectual options she is currently being prevented from accessing.²⁴

When it comes to the intellectual options to which she does have access—namely, the information her captors provide—she can be critically receptive to them as well. One caveat to mention upfront is that depending on what the intellectual options in question are, they may or may not count as relevant. Suppose some of them are quite far-fetched. Julie could reasonably dismiss them and not fail to count as open-minded. Open-mindedness only requires being critically receptive to relevant intellectual options. Being critically receptive to irrelevant intellectual options is not

²⁴Baehr (2011, 153) would appear to agree.

to be open-minded, but rather to be fanciful or gullible. Ignoring these far-fetched, irrelevant intellectual options is thus consistent with being open-minded.

Suppose some other intellectual options presented to her are not so far-fetched, but can reasonably be considered relevant to her inquiry. They are, at least to some degree, evidentially relevant to her. It seems that Julie should be critically receptive to these intellectual options. However, a qualification is in order, for it also seems to me that, while Julie should be critically receptive to the intellectual options which she deems to be relevant to her inquiry, she should also take the information she's given with a grain of salt. After all, she knows she's in a hostile epistemic environment, and that even information which seems genuinely relevant to her inquiries may nevertheless be misleading. So, when she is critically receptive to these intellectual options, she should be more critical toward them than she would be were she in a normal epistemic environment. Placing the emphasis here on the "critical" aspect of critical receptivity helps us understand why Julie counts as being open-minded to the relevant intellectual options to which she has access, despite the fact that she's in an epistemically hostile environment.

In short, it is still possible for Julie to exercise the virtue of open-mindedness in this hostile environment. Moreover, what this discussion shows us is that, while hostile environments do not necessarily prevent one from being critically receptive to relevant intellectual options, what they can do is affect the degree to which one should be critical versus receptive. In other words, there is a sense in which the "critical" and "receptive" aspects of critical receptivity are pulling in opposite directions. Part of what's involved in the exercise of virtuous open-mindedness is finding the right

balance between being critical and being receptive. One effect of being in hostile epistemic environments is that the point on the spectrum between criticism and receptivity that one ought to aim for can shift. In more hostile environments, one should err more on the side of being more critical than receptive.

Importantly, in arguing that being ignorant of relevant intellectual options does not suffice to make one closed-minded, I do not mean to deny that epistemically hostile environments are so dangerous in part because they *foster* closed-mindedness (and other intellectual dispositions or traits that are on the whole bad). However, fostering closed-mindedness is different from automatically rendering one closed-minded. Contra Battaly, I think hostile or unhealthy epistemic environments do the former, but not necessarily (or not always) the latter.

Failure to engage with relevant intellectual options: laziness or lack of curiosity

Just now, we discussed cases in which someone does not seriously engage with relevant intellectual options because she is ignorant of them. We might also consider cases in which one fails to engage with relevant intellectual out of mere laziness:

Lazy Lucy

Lucy wants to learn more about US involvement in Iraq. At first, she entertains the idea of seeking out technical, rigorously researched reports. After some thought, however, she decides to only read Huffington Post articles about the topic in an effort to save time and mental energy.

Importantly, it is not as though Lucy only reads Huffington Post pieces because she does not want to discover the truth about the matter. On the contrary, it is *because* she

wants to know the truth that she investigates. Still, as she is too lazy to conduct a thorough investigation, she takes the easy route by only reading un-nuanced and accessible synopses of the findings to date.

Is Lucy closed-minded because she fails to seriously engage with more rigorous or technical reports (or even synopses of said reports)? Perhaps so, but not necessarily. Certainly, Lucy is being lazy in her investigation, but I do not think that suffices to make her closed-minded. The reason for this, I think, is that it is not clear that she's entirely *unwilling* to read more detailed, technical reports. Suppose that, were a friend were to call Lucy out for taking the easy route, Lucy would buckle down and read the more serious, technical reports. If this were the case, then I think the best way to describe Lucy is to say that she is neither open-minded nor closed-minded, but rather something in between. To reiterate a point made above: failing to seriously engage with relevant intellectual options does not make one closed-minded. One can fail to seriously engage with relevant intellectual options and still be open-minded, as Julie is (and, I think, Shelly). Or, one can fail to seriously engage with relevant intellectual options and be neither closed-minded nor open-minded, but rather somewhere in between, as I submit Lucy is.

5.5.2 Engaging in an unserious manner

The second general way in which one can fail to be critically receptive to a relevant intellectual option is to engage with it, but to not do so seriously. Following Battaly (2018a), I will identify two ways in which this might happen. First, one can *dismiss*, rather than ignore, relevant intellectual options (2018a, 24). To illustrate, let's return

to Elitist Ellen. Suppose that instead of flat out refusing to read the paper her friend recommends, Ellen agrees to take a look at it. However, she does not put any effort into appreciating the merits of the author's contributions, or to temper her biases against non-philosophical scholarship. Instead, Ellen merely skims through the paper, judges that it is not insightful, and moves on to something else. While Ellen in some sense engages with the paper, she does so very superficially and uncharitably.

Second, one might *divert* attention to something else (Battaly 2018a, 22). For example, one can superficially engage with an intellectual option long enough to realize that one does not want to continue engaging with it. Rather than dismissing it, one might divert one's attention (or one's audience's attention) somewhere else. As Battaly points out, politicians are quite fond of this move. As was discussed just above, what the underlying reasons are for why one engages with relevant intellectual options in a merely cursory manner will make a difference when it comes to whether or not one counts as closed-minded, or merely fails to be open-minded.

5.6 Conclusion

To sum up: On my view, to be open-minded is to be critically receptive to relevant intellectual options. Capturing open-mindedness this way means that being open-minded is compatible with failing to seriously engage with relevant intellectual options.

Chapter 6

Conclusion

The purpose of this dissertation has been to clarify the role that open-mindedness ought to play in our lives, particularly when it comes morally and politically significant issues. The view I've presented is in one sense traditional, and in another sense unorthodox. It is traditional insofar as it vindicates the value of open-minded engagement. After all, most are happy to grant that we generally ought to be open-minded. My view is unorthodox in that I insist that what supplies epistemic obligations with most (if not all) of their normative force are moral, social, and political obligations. This distinctly pragmatic approach to epistemic normativity is a decisive break from traditional lines of thought in epistemology, according to which epistemic normativity ought to be both treated as distinct from, as well as theorized separately from, non-epistemic varieties of normativity.

One important consequence of this view is that our epistemic obligations—including the obligation to be open-minded—are both stronger, and come into play more of-

ten, than we might ordinarily think. Not only ought epistemic agents generally be open-minded, but also, it will typically hold that the more significant a topic is non-epistemically, the stronger the epistemic obligation to open-mindedly engage with intellectual options relevant to that topic.

This final chapter is dedicated to bringing together the various themes of the dissertation. I begin by highlighting three important takeaways, namely: we have epistemic obligations to be open-minded; these obligations are underpinned by non-epistemic obligations; and, while the obligation to engage open-mindedly is agent-universal, it nevertheless holds that when any given individual ought to be open-minded is something that can only be determined on a case-by-case basis. I conclude by identifying two additional values of open-minded engagement which serve to further highlight the role it can and should play in democratic life.

6.1 On epistemic normativity

Earlier, we discussed how it is typical for accounts of epistemic normativity to treat so-called purely epistemic obligations as being of central theoretical importance, and to treat morally grounded and socially grounded epistemic obligations as being at best peripheral concerns. For instance, we have Rik Peels' observation that "philosophers are likely to be especially interested in *epistemic* obligations" (2017, 101, original emphasis), where epistemic obligations are understood as (1) non-contingent, and (2) as not concerned with avoiding any kind of "badness" other than epistemic badness (2017, 101-02). We also encountered Richard Feldman's conception of epis-

temic obligation, according to which they are “obligations to believe to which the practical benefits of beliefs are not relevant. They are obligations that arise from a purely impartial and disinterested perspective” (1988, 236).

I have shown this way of theorizing about epistemic obligations to be misguided. While there may be such a thing as purely epistemic obligations—I will remain neutral on this point, though for what it’s worth, I’m skeptical they do exist—I maintain that any account which does not give due weight to non-epistemic considerations will wind up with a distorted conception of the nature and extent of our epistemic obligations. The first important takeaway, then, is that the most promising accounts of epistemic normativity will be heavily social in nature. In other words, insofar as an account of epistemic normativity is meant to be informative and provide action-guidance for epistemic agents living in the actual world, it must treat moral, social, and political considerations as being of central theoretical concern.

I have also established that, while others are entitled to expect in general that the beliefs we have are justified, the more important a certain topic is morally, socially, or politically, the stronger in general the entitlement to expect that beliefs relevant to these topics are justified. More precisely: epistemic agents are entitled to expect of each other that beliefs relevant to their moral, social, and political obligations are justified. Moreover, the the stronger one’s non-epistemic obligation, the stronger others’ entitlement to expect that one’s beliefs about topics relevant to that obligation are justified. In short: non-epistemic considerations shape both what epistemic obligations we have, as well as how strong those epistemic obligations are.

6.2 On the obligation to open-mindedly engage

Second, it is a consequence of my view of epistemic normativity that moral, social, and political agents are entitled to expect that others be virtuously open-minded—i.e., critically receptive—to intellectual options that are relevant to their non-epistemic obligations. For when others are entitled to expect that one has justified beliefs or knowledge about a certain topic, they are thereby generally entitled to expect that the beliefs one has (or will come to have) about that topic are formed through intellectually virtuous acts, or through the exercise of intellectually virtuous skills, including open-mindedness. It also follows that the stronger one's moral, political, or social obligations, the stronger others' entitlement to expect that one be critically receptive to intellectual options relevant to those obligations.

As far as the topic of relevancy goes, I'd like to take this opportunity to offer some tentative thoughts. Earlier, we distinguished between three types of relevancy: evidential, objective, and social. I argued that all three senses are legitimate in that an intellectual option can count as relevant full-stop if it's socially relevant, but not evidentially or objectively relevant; objectively relevant, but not evidentially or socially relevant; and so on.

While further work is necessary for determining how this plays out in practice, my prediction is this: although epistemic agents are entitled to expect that one be critically receptive to intellectual options that are either socially, evidentially, or objectively relevant to one's non-epistemic obligations, and although the three senses of relevancy will often overlap (i.e., what's socially relevant will often also be evidentially and objectively relevant), the type of relevancy which will often be of

special interest is social relevancy. Here is why. The theory I've developed concerns what epistemic agents are entitled to expect of each other. It stands to reason, then, that in most cases, the intellectual options which others will be entitled to expect one to be critically receptive to are intellectual options that are widely believed (and with good reason) by members of one's epistemic community to be likely to aid in inquiry.

To illustrate the point, let's return Ignorant Ian, who failed to meet his moral obligation to not promote or prop up unjust wage disparity in his company. The explanation on offer for why he failed to meet his moral obligation is that he failed to be critically receptive to intellectual options that were socially relevant to his obligation, such as: the fact that race, gender, and gender expression impact one's chances of moving up social and professional ladders; that being socialized in virtue of having particular social identities can impact whether one is a victim of, or beneficiary of, wage disparity; that stereotypes and implicit biases can distort one's perception of who is more or less qualified, and who has a better or worse work ethic.

As I argued, others are entitled to expect more from Ian. For one, his investigation was superficial. More to the point, those who fight against social inequality regularly point out that (1) differences in qualification and work ethic are often incorrectly cited as a means to justify wage disparity, and (2) that implicit biases, networking, socialization, and other factors are often causes of unjust wage disparity. In other words, Ian failed to be critically receptive to intellectual options that were socially relevant to his moral obligation. In doing so, he failed to meet others' legitimate expectations.

Again, this is not to say that the other two sense of relevancy are not in play in this case, or that they are not worth paying attention to, here or elsewhere. It's simply that, insofar as meeting one's epistemic obligations involves meeting others' legitimate expectations, when we look at particular cases in which one has an epistemic obligation to be critically receptive to relevant intellectual options, it stands to reason that social relevancy will typically be of special interest.

6.3 On the context-sensitivity of the obligation to open-mindedly engage

The third important takeaway is that, while the obligation to engage open-mindedly is agent-universal, it nevertheless holds that whether any given individual ought to be open-minded in a particular instance is something to be determined on a case-by-case basis. To repeat what I emphasized at the beginning of this project: I do not pretend to be able to provide the final say on the question of when to be open-minded. In fact, I think that would be a mistake to assume that this is even realistic or desirable. Rather, my intent has been to provide a solid foundation for being able to properly answer the question at all. In part to ensure that there are no misunderstandings about my view, I will now take the opportunity to lay out both what I have argued, as well as what I have not argued.

What I have not attempted to provide is all-encompassing verdict on when to be open-minded. In other words, while I have argued that others generally are entitled to expect that we be open-minded, especially when it comes to matters of moral,

social, or political significance, I am not suggesting that all epistemic agents have the same obligations (or same strength of obligation) to open-mindedly engage with the same kinds of intellectual options. This is because what an individual's epistemic obligations are is a function of that individual's non-epistemic obligations. And, non-epistemic obligations can vary quite a bit from person to person.

In more detail: what epistemic obligations an individual has will be determined by a number of factors which include, but are not limited to: the non-epistemic significance of the topic—in particular, whether it is relevant to one's non-epistemic obligations; whether or not the intellectual option to be considered is significantly morally problematic;¹ one's positionality, and how that affects what one's non-epistemic obligations are in the first place, as well as what others are entitled to expect of one; and more. Determining what any individual ought to do in a given situation, then, is not something that can be settled in advance. A prime benefit of providing a proper framework for answering the question of whether one ought to be open-minded is that when we do examine specific cases, we will be in a much better position to accurately determine whether open-minded engagement is called for.

What I have argued is that, despite the sense in which the epistemic obligation to engage open-mindedly does vary person-to-person, there is nevertheless another sense in which the obligation is agent-universal. Here is what I mean by this. In the first place, epistemic obligations underpinned by moral obligations are agent-universal since all epistemic agents (at least the non-idealized ones about whom I'm

¹Just as morality can dictate that we ought to investigate and have justified beliefs about certain topics, so too can morality dictate that we don't open-mindedly engage (or engage at all) with certain intellectual options.

theorizing) are also moral agents. And moral agents can't well decide to simply stop being moral agents. Thus, even though what moral obligations various individuals are bound by will vary, all moral agents have moral obligations, and by extension, corresponding epistemic obligations. This is the sense in which morally grounded epistemic obligations are agent-universal. Second, epistemic obligations underpinned by social obligations are also agent-universal since, although social agents can often decide to exit particular roles in order to relieve themselves of obligations specific to those roles, they cannot exist all social roles without going entirely off the grid. Thus, even though what social obligations various individuals are bound by will vary, all social agents have social obligations, and by extension, corresponding epistemic obligations. This is the sense in which socially grounded epistemic obligations are agent-universal.

6.4 My view compared to Fantl's

To shed further light on what my view recommends, it would be useful to return to the starting point of our discussion and compare my view to Fantl's. He warns that it can often be risky for laypersons to open-mindedly engage with arguments they lack the expertise to reliably evaluate. I agree that being critically receptive to relevant intellectual options can sometimes produce bad epistemic effects. This can happen, for instance, when one is not intellectually virtuous enough on the whole to be able to inquire responsibly. To take an example, open-mindedly engaging with evidence suggesting that human activity is not a major contributor to climate

change can produce bad epistemic effects if a person is quite gullible, if she is inept at distinguishing reliable sources from unreliable sources, or if she is unable to determine which experts to trust.

Fantl is right that if a person lacks the ability to reliably evaluate an argument, she ought not attempt to assess the merits of it on her own. What we disagree on is what the layperson ought to do, given her lack of expertise. He argues that a laypersons' inability to reliably evaluate misleading evidence which conflicts with what she knows means that she should be closed-minded toward that evidence. More specifically, she ought to bracket any evidence the counterargument provides and only count as relevant the evidence she had prior to engaging with it. She should not try to determine what experts might have to say about such evidence. Indeed, she should give the matter no further thought.

I've demonstrated that Fantl's bracketing argument fails, meaning that an inability to reliably evaluate a sophisticated counterargument does not license a layperson to closed-mindedly dismiss it. Rather, depending on a number of factors (including, but not limited to: what one's non-epistemic obligations are, how morally, politically, or socially significant the disputed matter is, how much time one reasonably has to engage, and so forth), one ought to seek out expert opinion on the matter. Of course, engaging with expert opinion can lead to bad epistemic effects if one lacks the ability to reliably determine which experts to trust. However, the fact that open-mindedly engaging with expert opinion can be epistemically risky does not mean that we should steer clear of open-mindedly engaging with expert opinion. It simply means that we ought to become better at determining which experts to trust.

Returning to the swimming analogy from Chapter 5, the solution Fantl advocates is akin to telling someone she should give up on learning to swim if she lacks good form. But this is obviously misguided. Although trying to make it down the length of a pool while she has poor form is likely to lead to bad results, this does not mean she should give up on swimming. Rather, she should, if she wants to learn how to swim, work on improving her form. Similarly, even if a person is so lacking in intellectually virtuous skill—including the skill of being able to determine which experts to trust—such that open-mindedly engaging with certain intellectual options would likely lead to epistemically undesirable results in the short term, this does not mean that she need not engage with such intellectual options. Rather, she should, if she is to be a good inquirer, work on cultivating the requisite virtuous skills in order to be a responsible inquirer—i.e., in order to be able to open-mindedly engage with these intellectual options responsibly or virtuously.

One key difference, of course, between swimming and inquiry is that whereas it is often good to learn how to swim, one typically isn't obligated to learn how to swim. In contrast, epistemic agents have epistemic obligations to have justified beliefs on topics that are relevant to their non-epistemic obligations. Being able to virtuously open-mindedly engage is necessary for meeting those obligations. In other words, the fact that controversial topics often matter so much, morally or politically, means that it matters that our beliefs about these topics are justified, as well as that we can justify them to others. Moral and political agents are entitled to expect this much of one another. If the problem is that many people are inept at distinguishing good evidence from bad, or at determining which experts to trust, the

answer is not more closed-mindedness. Rather, the answer is to promote and nurture skills such as critical thinking, scientific literacy, and the ability to determine which experts and sources are reliable and trustworthy, such that more virtuous open-minded engagement is possible.

6.5 Open-mindedness and public discourse

As a final consideration in favor of open-minded engagement, I'd like to highlight further benefits of virtuous open-minded engagement which I believe hold special value in the context of public discourse. Consider a case in which a person is faced with evidence which challenges a proposition she takes herself to know. Implicit in many discussions of the value of open-minded engagement (Fantl's included) is the idea that the most interesting way in which that person can be open-minded is to be open-minded to the proposition itself. But this, I think, is an incredibly limited way of thinking about the epistemic (and political) value of open-mindedness. Often times what one has to gain, epistemically and otherwise, from engaging with conflicting views or opinions in moral or political disagreements is not necessarily that one will come to change one's view, at least to any significant degree, about the topic at hand. Instead, one might walk away with a better understanding of why one's interlocutor's hold the views that they do.

Suppose I open-mindedly engage with parties who disagree with me on some morally or politically charged issue—e.g., the travel ban, the benefits and/or drawbacks of nuclear energy, or abortion. In the first place, I may end up increasing or

decreasing my confidence in my view to some degree or other. This may be because, as a result of the engagement, I end up changing my beliefs about how much (or little) I understand the relevant issues, or about how complex and nuanced the issue really is. And yet, even if I do not change my beliefs to any significant degree, either about the topic itself or about how well I understand some of the relevant issues, the result of open-minded engagement can be that I have a better idea of why those who disagree with me disagree in the first place. In other words, I might go into the exchange believing I have a solid understanding of why those who disagree with me hold the views they do. One result of the engagement, however, may very well be that I come out with more accurate beliefs about what the actual bases of their beliefs are.

This is, I think, an often overlooked point. Yet it is an incredibly important point to highlight, especially when considering the value of open-mindedness in the context of political discourse. For not only is there epistemic benefit (i.e., one can come to have more accurate beliefs about the reasons one's interlocutors have for believing what they do), but also, there is significant potential political benefit. In particular, it can potentially help in bridging the seemingly ever-increasing divide between opposing political ideologies. This, in turn, has the potential to reduce or mitigate the effects of political polarization, as well as the tendency to dislike (or, more strongly, contempt) those who disagree with one on some political or moral issue.

Worth pointing out is that this way of understanding the value of open-mindedness highlights the virtues to Baehr's conception of open-mindedness. Recall that on his

view open-mindedness involves a willingness and ability to take up a distinct cognitive standpoint. Although my conception of open-mindedness is more overtly similar to Battaly's or Hare's, I am quite sympathetic to Baehr's view for this very reason. It highlights the importance of being open-minded to where others are coming from, so to speak. In short: there seems to be a tendency in discussions about open-mindedness to focus on the benefits or harms that might come from changing one's view on the issue at hand. However, especially in the context of public discourse, I think serious attention ought to be paid to the benefits that might come from simply changing one's views on why those who disagree with one hold the views that they do.

The second point to highlight is that open-minded engagement may also be valuable because the process of engaging is itself valuable. In particular, open-mindedly engaging with those with whom you disagree on some moral or political issue is a way to model good political behavior. In addition to, or even separate from, any immediate good epistemic effects that might come from open-minded engagement, there can also be good long-term or downstream political (and likely epistemic) effects produced by such engagement. In short: the more open-minded engagement is normalized, the better off public discourse is likely to be as a whole.

6.6 Closing remarks

We have come a long way from where we started. Moving forward, we have a much better idea of what considerations make the key differences as to whether one ought

to be open-minded in a given instance—where these considerations include, but are not limited to: what others are entitled to expect of one, *qua* moral, political, and social agent; how what others are entitled to expect is determined in large part by one’s positionality; and how the obligation to open-mindedly engage with certain intellectual options (including both the strength of the obligation, as well as whether one is obligated to engage at all) relates to one’s non-epistemic obligations. We also have a better understanding of the additional value that open-minded engagement promises to hold in the context of public discourse. Taken together, what this project has demonstrated is that open-mindedness ought to play a much larger role in our lives—particularly when it comes to morally or politically significant issues—than we may have otherwise thought.

Bibliography

- [1] Allen, T. & Lynch, M. P. (2020). Can We Be Reasonable? Bias, Skepticism, and Public Discourse. Forthcoming in N. Ballantyne & D. Dunning (Eds.), *Reason, Bias, and Inquiry: New Perspectives from the Crossroads of Epistemology and Psychology*. Oxford: Oxford University Press.
- [2] Alston, W. P. (1980). Level-Confusions in Epistemology. *Midwest Studies in Philosophy* 5(1), 135-50.
- [3] Alston, W. P. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives* 2, 257-299.
- [4] Baehr, J. (2011). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford: Oxford University Press.
- [5] Baldwin, D. A., Markman, E. M., & Melartin, R. L. (1993). Infants' Ability to Draw Inferences about Nonobvious Object Properties: Evidence from Exploratory Play. *Child Development* 64 (3), 711–28.
- [6] Banaji, M. R. (2002). Social Psychology of Stereotypes. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences* (pp. 15100-04). New York: Pergamon.
- [7] Battaly, H. D. (2018a). Can Closed-mindedness be an Intellectual Virtue? *Royal Institute of Philosophy Supplement* 82, 23-45.
- [8] Battaly, H. D. (2018b). Closed-mindedness and Dogmatism. *Episteme* 15(3): 261-282.
- [9] Christensen, D. (2010a). Higher-Order Evidence. *Philosophy and Phenomenological Research* 81(1), 185-215.

- [10] Christensen, D. (2010b). Rational Reflection. *Philosophical Perspectives* 24, 121-140.
- [11] Clifford, W. K. (1886). The Ethics of Belief. In L. Stephen & F. Polluck (Eds.), *Lectures and Essays by the Late William Kingdon Clifford, F.R.S.* (pp. 339-363). London: MacMillan and Co. (Original work published in 1877)
- [12] Coady, D. (2012). *What to Believe Now: Applying Epistemology to Contemporary Issues*. West Sussex, UK: John Wiley Sons, Ltd.
- [13] Cohen, G. L. (2003). Party over Policy: The Dominating Impact of Group Influence on Political Beliefs. *Journal of Personality and Social Psychology* 85, 808-22.
- [14] C-SPAN. (2020, February 28). *President Trump Campaign Event in North Charleston, South Carolina*. <https://www.c-span.org/video/?469663-1/president-trump-campaign-event-north-charleston-south-carolinastart=1337>
- [15] Dana, J., & Loewentstein, G. (2003). A Social Science Perspective on Gifts to Physicians from Industry. *Journal of the American Medical Association* 290, 252-55.
- [16] Davidson, D. (2005). *Truth, Language and History*. Oxford: Clarendon Press.
- [17] Devine, P. G., & Sharp, L. B. (2009). Automatic and Controlled Processes in Stereotyping and Prejudice. In T. D. Nelson (Ed.), *Handbook of Prejudice, Stereotyping, and Discrimination* (pp. 61-82). New York: Psychology Press.
- [18] Dewey, J. (1933). *How We Think: A Restatement of the Relation of Reflective Thinking to the Educative Process*. Lexington, MA: D.C. Heath and Company.
- [19] Dovidio, J. F., & Gaertner, S. L. (1991). Changes in the Expression and Assessment of Racial Prejudice. In H. J. Knopke, R. J. Norrell, & R. W. Rogers (Eds.), *Opening Doors: Perspectives of Race Relations in Contemporary America* (pp. 119-48). Tuscaloosa: University of Alabama Press.
- [20] Dovidio, J. F., & Gaertner, S. L. (2000). Aversive Racism and Selection Decisions: 1989 and 1999. *Psychological Science* 11(4), 315-19.
- [21] Dovidio, J. F., & Gaertner, S. L. (2004). Averse Racism. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (vol. 36, pp. 1-52). San Diego: Academic Press.

- [22] Du Bois, W. E. B. (1903). *The Souls of Black Folk*. Chicago: A. C. McClurg Co.
- [23] Dunton, B. C., & Fazio, R. H. (1997). An Individual Difference Measure of Motivation to Control Prejudiced Reactions. *Personality and Social Psychology Bulletin* 23, 316-26.
- [24] Epley, N., & Dunning, D. (2000). ‘Feeling ‘Holier Than Thou’: Are Self-serving Assessments Produced by Errors in Self- or Social Prediction? *Journal of Personality and Social Psychology* 79, 861-75.
- [25] Fanon, F. (2008). *Black Skin, White Masks*. (R. Philcox, Trans). New York: Grove Books. (Original work published in 1952)
- [26] Fantl, J. (2018). *The Limitations of the Open Mind*. Oxford: Oxford University Press.
- [27] Fantl, J. & McGrath, M. (2002). Evidence, Pragmatics, and Justification. *The Philosophical Review* 111(1), 67-94.
- [28] Fantl, J. & McGrath, M. (2007). On Pragmatic Encroachment in Epistemology. *Philosophy and Phenomenological Research* 75(3), 558-89.
- [29] Fantl, J. & McGrath, M. (2012). Pragmatic Encroachment: It’s not Just About Knowledge. *Episteme* 9(1), 27-42.
- [30] Fantl, J. & McGrath, M. (2019). Clarifying Pragmatic Encroachment: a Reply to Charity Anderson and John Hawthorne on Knowledge, Practical Adequacy, and Stakes. *Oxford Studies in Epistemology* 6, 258-66.
- [31] Faust, J. (2008). Can Religious Arguments Persuade? *International Journal for Philosophy of Religion* 63, 71-86.
- [32] Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in Automatic Activation as an Unobtrusive Measure of Racial Attitudes: A Bona Fide Pipeline? *Journal of Personality and Social Psychology* 69, 1013-27.
- [33] Feldman, R. (1988). Epistemic Obligations. *Philosophical Perspectives* 2, 235-256.
- [34] Feldman, R. (2000). The Ethics of Belief. *Philosophy and Phenomenological Research* 60, 667-95.

- [35] Fricker, M. (2007). *Epistemic Injustice: The Power and Ethics of Knowing*. New York, NY: Oxford University Press.
- [36] Gendler, T. S. (2011). On the Epistemic Costs of Implicit Bias. *Philosophical Studies* 156, 33-63.
- [37] Gert, B. (1998). *Morality*. New York: Oxford University Press.
- [38] Gilbert, M. (2006). Rationality in Collective Action. *Philosophy of the Social Sciences* 36(1), 3-17.
- [39] Gilbert, M. (2013). *Joint Commitment: How We Make the Social World*. New York: Oxford University Press.
- [40] Goldberg, S. C. (Forthcoming). Social Epistemic Normativity: The Programme.
- [41] Goldberg, S. C. (2017). Should Have Known. *Synthese* 194(8), 2863-2894.
- [42] Goldman, A. (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 73(20), 771-91.
- [43] Goldman, A. (2001). Experts: Which Ones Should You Trust? *Philosophy and Phenomenological Research* 63(1), 85-110.
- [44] Govier, T. (1999). *The Philosophy of Argument*. Newport News, VA: Vale Press.
- [45] Greenawalt, K. (1989). Free Speech Justifications. *Columbia Law Review* 89(1), 119-55.
- [46] Greenwald, A. G., McGhee, D. E., & Schwartz, J. K. L. (1998). Measuring Individual Differences in Implicit Cognition: The Implicit Association Test. *Journal of Personality and Social Psychology* 74, 1464-80.
- [47] Grimm, S. R. (2009). Epistemic Normativity. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic Value* (pp. 243-264). Oxford: Oxford University Press.
- [48] Gunn, H. (2018). *Social-Epistemic Agency* (Publication No. 1875) [Doctoral Dissertation, University of Connecticut]. Retrieved from OpenCommons. (<https://opencommons.uconn.edu/dissertations/1875/>)
- [49] Gutmann, A., & Thompson, D. (1990). Moral Conflict and Political Consensus. *Ethics* 101(1), 64-88.

- [50] Haidt, J., Björklund, F., & Murphy, S. (2000). *Moral Dumbfounding: When Intuition Finds No Reason*. (Lund psychological reports; Vol 1 No 2). Department of Psychology, Lund University.
- [51] Haidt, J. (2001). The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review* 108(4), 814-34.
- [52] Hare, W. (2001). Bertrand Russell and the Ideal of Critical Receptiveness. *Skeptical Inquirer* 25(3), 40-44.
- [53] Hare, W. (2003). Is it Good to be Open-Minded? *International Journal of Applied Philosophy*, 17(1), 73-87.
- [54] Hare, W. (2009). Socratic Open-mindedness. *Paideusis*, 18(1), 5-16.
- [55] Harman, G. (1973). *Thought*. Princeton, NJ: Princeton University Press.
- [56] Heath, C. (1999). On the Social Psychology of Agency Relationships: Lay Theories of Motivation Overemphasize Extrinsic Incentives. *Organizational Behavior and Human Decision Processes* 78, 25-62.
- [57] Huang, P. (2020, March 20). *How The Novel Coronavirus And The Flu Are Alike ... And Different*. NPR. <https://www.npr.org/sections/goatsandsoda/2020/03/20/815408287/how-the-novel-coronavirus-and-the-flu-are-alike-and-different>
- [58] Johnson, C. (2018). For the Sake of Argument: The Nature and Extent of our Obligation to Voice Disagreement. In C. Johnson (Ed.), *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public* (pp. 97-108). New York, NY: Routledge.
- [59] Johnson, R. (2014). *The Rise of Informal Logic: Essays on Argumentation, Critical Thinking, Reasoning and Politics*. Windsor, ON: University of Windsor.
- [60] Jones, K. (2017). But I was Counting on You! In P. Faulkner & T. Simpson, (Eds.), *The Philosophy of Trust* (pp. 90-109). Oxford: Oxford University Press.
- [61] Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus, Giroux.
- [62] Kelly, T. (2005). Moorean Facts and Belief Revision, or Can the Skeptic Win? *Philosophical Perspectives* 19(1), 179-205.

- [63] Kelly, T. (2010). Peer Disagreement and Higher-Order Evidence. In R. Feldman & T. A. Warfield (Eds.), *Disagreement* (pp. 111-74). Oxford, Oxford University Press.
- [64] Kelp, C. (2014). Two for the Knowledge Goal of Inquiry. *American Philosophical Quarterly* 51, 227–32.
- [65] Kessler, G., Rizzo, S., & Kelly, M. (2019, October 14). *President Trump Has Made 13,435 False or Misleading Claims Over 993 Days*. The Washington Post. <https://www.washingtonpost.com/politics/2019/10/14/president-trump-has-made-false-or-misleading-claims-over-days/>
- [66] Kornblith, H. (1993). Epistemic Normativity. *Synthese* 94(3), 357-376.
- [67] Korsgaard, C. M. (1996). *The Sources of Normativity*. (O. O'Neill, Ed.). Oxford: Oxford University Press.
- [68] Kruger, J., & Dunning, D. (1999). Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments. *Journal of Personality and Social Psychology* 77(6), 1121-34.
- [69] Kujawski, M. (2019, September 5). *Misinformation vs. Disinformation vs. Mal-information*. Medium. <https://medium.com/@mikekujawski/misinformation-vs-disinformation-vs-mal-information-a2b741410736>
- [70] Kunda, Z. (1999). *Social Cognition: Making Sense of People*. Cambridge, MA: MIT Press.
- [71] Lackey, J. (2018a). The Duty to Object. *Philosophy and Phenomenological Research*. doi: <https://doi.org/10.1111/phpr.12563>
- [72] Lackey, J. (2018b). Silence and Objecting. In C. Johnson (Ed.), *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public* (pp. 82-96). New York: Routledge.
- [73] Lemon, D. (2020, March 18). *President Trump is Gaslighting You* [Video]. CNN. <https://www.cnn.com/videos/us/2020/03/18/don-lemon-president-trump-politicizing-coronavirus-ctn-vpx.cnn>
- [74] Leslie, S-J. (2017). The Original Sin of Cognition: Fear, Prejudice, and Generalization. *The Journal of Philosophy* 118 (8), 393-421.

- [75] Lynch, M. P. (2009). The Values of Truth and the Truth of Values. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic Value* (pp. 225-242). Oxford: Oxford University Press.
- [76] Lynch, M. P. (2019). *Know-It-All Society: Truth and Arrogance in Political Culture*. New York, NY: Liveright Publishing Corporation.
- [77] Mill, J. S. (1998). *Utilitarianism*. (R. Crisp, Ed.). New York, NY: Oxford University Press. (Original work published in 1861)
- [78] Mill, J. S. (2003). *On Liberty*. New Haven, CT: Yale University Press. (Original work published in 1859)
- [79] Millar, A. (2011). Why Knowledge Matters. *Proceedings of the Aristotelian Society, Supplementary Volume, 85*, 63–81.
- [80] Miller, D. T., & Ratner, R. K. (1998). The Disparity Between the Actual and Assumed Power of Self-Interest. *Journal of Personality and Social Psychology* 74, 53-62.
- [81] Mills, C. W. (2017). Black Rights/White Wrongs: The Critique of Racial Liberalism. Oxford: Oxford University Press.
- [82] Nelson, M. (2010). We Have No Positive Epistemic Duties. *Mind* 119, 83-102.
- [83] Nguyen, T. (2018). Echo Chambers and Epistemic Bubbles. *Episteme*, 1-21. <https://doi.org/10.1017/epi.2018.32>
- [84] Nickerson, R. S. (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology* 2(2), 175-220.
- [85] Paxson, T., & Lehrer, K. (1969). Knowledge: Undefeated Justified True Belief. *Journal of Philosophy* 66, 225-37.
- [86] Payne, B. K. (2001). Prejudice and Perception: The Role of Automatic and Controlled Processes in Misperceiving a Weapon. *Journal of Personality and Social Psychology* 81(2), 181-92.
- [87] Peels, R. (2017). *Responsible Belief: A Theory in Ethics and Epistemology*. New York: Oxford University Press.
- [88] Plato. (2004). *Republic*. (C. D. C. Reeve, Trans.). Indianapolis, IN: Hackett Publishing Company, Inc.

- [89] Pritchard, D. (2008). Knowing the Answer, Understanding and Epistemic Value. *Grazer Philosophische Studien* 77, 325-39.
- [90] Pronin, E., Berger, J. A., & Molouki, S. (2007). Alone in a Crowd of Sheep: Asymmetric Perceptions of Conformity and Their Roots in an Introspection Illusion. *Journal of Personality and Social Psychology* 92, 585-95.
- [91] Pronin, E., & Schmidt, K. (2013). Claims and Denials of Bias and Their Implications for Policy. In E. Shafir (Ed.), *The Behavioral Foundations of Public Policy* (pp. 195-216). Princeton, NJ: Princeton University Press.
- [92] Rawls, J. (1993). *Political Liberalism*. New York: Columbia University Press.
- [93] Riggs, W. (2016). Open-mindedness, Insight, and Understanding. In J. Baehr (Ed.), *Intellectual Virtues and Education* (pp. 18-37). New York: Routledge.
- [94] Robinson, R. J., Keltner, D., Ward, A., & Ross, L. (1995). Actual Versus Assumed Differences in Construal: 'Naïve Realism' in Intergroup Perception and Conflict. *Journal of Personality and Social Psychology* 68, 404-17.
- [95] Rogers, K. (2020, March 20). *Trump Now Claims He Always Knew the Coronavirus Would Be a Pandemic*. The New York Times. <https://www.nytimes.com/2020/03/17/us/politics/trump-coronavirus.html>
- [96] Rorty, R. (1995). Is Truth a Goal of Enquiry? Davidson vs. Wright. *The Philosophical Quarterly* 45, 281-300.
- [97] Rorty, R. (2000). Universality and Truth. In R. Brandom (Ed.), *Rorty and His Critics*. (pp. 1-30). Oxford: Blackwell.
- [98] Rosen, G. (2002). Culpability and ignorance. *Proceedings of the Aristotelian Society* 103(1), 61-84.
- [99] Russell, B. (2004). *Sceptical Essays*. New York, NY: Routedledge. (Originally published in 1928. London: George Allen and Unwin, LTD.)
- [100] Scriven, M. (1984). Critical for Survival. *National Forum* 65, 9-12.
- [101] Sinnott-Armstrong, W. (2005). You Ought to Be Ashamed of Yourself (When You Violate an Imperfect Moral Obligation). *Philosophical Issues* 15, 193-208.
- [102] Smith, H. (1983). Culpable ignorance. *Philosophical Review* 92(4), 543-571.

- [103] Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. New York, NY: Cambridge University Press.
- [104] Stanley, J. (2005). *Knowledge and Practical Interests*. New York: Oxford University Press.
- [105] Stapleford, S. (2013). Imperfect Epistemic Duties and the Justificational Fecundity of Evidence.” *Synthese* 190(18), 4065-4075.
- [106] Uhlmann, E. L., & Cohen, G. L. (2007). ‘I Think It, Therefore It’s True’: Effects of Self-Perceived Objectivity on Hiring Discrimination. *Organizational Behavior and Human Decision Process* 104, 207-23.
- [107] Vazquez, M. (2020, March 3). *Task Force Health Expert Contradicts Trump about Coronavirus Vaccine Timing*. CNN. <https://www.cnn.com/2020/03/02/politics/donald-trump-coronavirus-vaccine-push-back/index.html>
- [108] Vivian, J. E., & Berkowitz, N. H. (1992). Anticipated Bias from an Outgroup: An Attributional Analysis. *European Journal of Social Psychology* 22, 415-24.
- [109] Weiland, N. (2020, March 7). *Anyone Who Wants a Coronavirus Test Can Have One, Trump Says. Not Quite, Says His Administration*. The New York Times. <https://www.nytimes.com/2020/03/07/us/politics/trump-coronavirus-messaging.html>
- [110] Williams, B. (1973). Deciding to Believe. In B. Williams (Ed.), *Problems of the Self* (pp. 136–51). Cambridge, MA: Cambridge University Press.
- [111] Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- [112] Wittgenstein, L. (1969). *On Certainty*. (G. E. M. Anscombe & G. H. von Wright, Eds.). (D. Paul & G. E. M. Anscombe, Trans.). Oxford: Blackwell.
- [113] Zagzebski, L. T. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. New York, NY: Cambridge University Press.