DEFENDING AN INDIRECT NORMATIVITY OF BELIEF

_____

A Dissertation

presented to

the Faculty of the Graduate School

at the University of Missouri-Columbia

_____

In Partial Fulfillment

of the Requirements for the Degree

Doctor of Philosophy

_____

by

STEPHEN PERINCHERY-HERMAN

Dr. Peter Vallentyne, Dissertation Supervisor

DECEMBER 2019

The undersigned, appointed by the dean of the Graduate School, have examined the dissertation entitled

DEFENDING AN INDIRECT NORMATIVITY OF BELIEF

presented by Stephen Perinchery-Herman, a candidate for the degree of Doctor of Philosophy, and hereby certify that, in their opinion, it is worthy of acceptance.

_____
Professor Peter Vallentyne

_____
Professor Peter Markie

_____
Professor Robert Johnson

_____
Professor Paul Litton

_____
Professor Matthew McGrath

Dedication

For Anna, for her endless patience, belief and support. Without her this dissertation would not be possible.

And

For Mr. Furey, who taught me how to be a philosopher before I even knew what philosophy was

Acknowledgements

I would first like to thank my dissertation advisor, Dr. Peter Vallentyne, for his eternal patience in reading extensive drafts of my dissertation and always providing insightful commentary. When faced with a choice of what dissertation topic I ought to take, Peter encouraged me to pursue what I was passionate about, and I am endlessly grateful. I am especially indebted to Dr. Peter Markie for his advice when I felt lost in my dissertation. It is no exaggeration to say that my dissertation would be nowhere close to done if not for his help. I want to also thank Dr. Matt McGrath for his amazing mentorship, guidance, objections, and support. It was he who first encouraged me to really look closely at normative epistemology as a subject. I also would like to thank the other members of my committee, Dr. Robert Johnson, and Dr. Paul Litton, and the other members of my department.

I would like to thank the members of the University of Missouri Center for Health Ethics, especially Dr. Lea Brandt, and Dr. Melanie Johnson-Moxley. Both have always provided encouragement and support: for my dissertation and my career. Likewise, both have encouraged me to be the best ethicist I can be. They have always made this philosopher feel welcome in health ethics. I am grateful.

I would also like to thank several of my fellow graduate students. First, Doug Moore for patiently listening to me ramble about my dissertation and other philosophical subjects, for virtually our entire graduate degree. Second, Dr. Keith Harris, Fernando Alvear, Zeinab Rabii, and Sukhvinder Shahi for participating in a reading group for this dissertation. Third, Dr. Adam Kosezla, Dr. Richard Lauer, Wesley Chambers and Chetan Cetty for teaching me how to be a

Table of Contents

Abstract

In this dissertation, I seek to answer the following questions: is there such a thing as deontic epistemic normativity—obligations, permissions, and prohibitions to act in a certain way based on epistemic grounds—and if so what does it consist in, and is it important for determining what we ought to do in practical reasoning? I argue for an indirect account of epistemic normativity: epistemic obligations command believers to act in certain ways so as to affect beliefs downstream of their actions. Further, I argue that if an agent commits him/herself to epistemic normativity, then these epistemic obligations can matter for the purposes of practical reason.

Introduction and Background

In this dissertation, I seek to answer the following questions: is there such a thing as deontic epistemic normativity—obligations, permissions, and prohibitions to act in a certain way based on epistemic grounds—and if so what does it consist in, and is it important for determining what we ought to do in practical reasoning? Intuitively, deontic epistemic normativity exists. This is evidenced by ordinary language practices of criticizing others for their beliefs. For example, we say things like "you shouldn't have believed that" in circumstances where a believer lacks sufficient evidence for her belief. Likewise, epistemological discourse points in the direction of norms of belief. Certain problems in epistemology depend on what it is appropriate for an agent to believe. Take for example the clairvoyant who seems to do wrong when he believes some proposition that he lacks evidence for.[1]

Nevertheless, in the past few decades, there has been relatively little discourse on deontic epistemic normativity. The reason for this stems from the fact that the most obvious way of accounting for deontic epistemic normativity, epistemic obligations to believe certain propositions, seems false. Simply put, if there are epistemic obligations to believe, then believers have control over their beliefs. However, believers lack control over their beliefs. For example, it is impossible for the reader of this dissertation to change her belief that he/she is not reading this dissertation right now by a mere act of willing. Thus, there are no deontic epistemic obligations to believe.

---

[1] Bonjour, "Externalist Theories of Empirical Knowledge."

This dissertation will seek to draw upon contemporary literature and offer an account of deontic epistemic normativity that is not subject to this objection. The key consideration for this kind of epistemic normativity is that believers' epistemic obligations do not command them to alter their beliefs directly. Instead, epistemic obligations command believers to act in certain ways so as to affect beliefs downstream of their actions. To give a simple example, by reading a book—an action—I gain beliefs that I did not have previously. Epistemic obligations, then, are concerned not directly concerned with whether I should have the beliefs I have after I have read the book, but whether I should have read the book to begin with. They are *indirect.* The key question we must address, on the indirect epistemic obligation picture, is what such epistemic obligations would look like.

After answering the question of what epistemic obligation would look like, we must turn to another problem. Trent Dougherty has argued that any account of deontic epistemic normativity like the kind raised in the previous paragraph is subject to a fatal flaw: it will not *matter* to you.[2] Simply put, epistemic normativity will only matter if it helps to provide an answer to the question of what an agent ought to do. Suppose, for example, I told you that some action you were about to take would violate the norms of medieval chivalry. More likely than not, it would not matter to you, as you do not care about obeying the norms of medieval chivalry. Dougherty's main objection is that if any alleged epistemic obligation matters to an agent, then it must be because that alleged epistemic obligation is actually a moral or prudential obligation. Say, for example, we thought John had an epistemic obligation to correct his beliefs regarding Young Earth Creationism. According to Dougherty, if we thought John had this obligation because it was in his own best interest,

---

[2] Dougherty, "Reducing Responsibility," December 1, 2012.

that would suggest that his obligation is really a prudential obligation. If we thought that John had this obligation because it was in others' best interests, that would suggest his obligation is really a prudential obligation. If changing his own beliefs is not in John's or anyone else's best interests, Dougherty argues, John has no obligation to change his belief. This objection leads to the second major concern of this dissertation: proving that epistemic normativity matters in spite of this objection.

This dissertation is structured around answering the two preceding questions: (1) What does epistemic normativity look like? (2) Does it matter? The first two papers are dedicated to answering the first question, and the third paper is dedicated to answering the second. The structure of my dissertation is as follows.

In Chapter 1, I review the state of normative epistemic literature. I start with the historical view that *epistemic justification*, that property of belief that makes it appropriate to be held by a rational believer, and epistemic permissibility are coextensive. I elucidate the objection from doxastic involuntarism and evaluate why this historical view is clearly false. From here, I depart and argue that we should separate deontic epistemic normativity from epistemic justification. I argue that we have sufficient reason to be interested in epistemic normativity, even if it is not coextensive with epistemic justification. we should offer a theory of epistemically permissible action, where permissibility consists in some combination of doing one's best (or a sufficiently good job) to causally bring about good epistemic states and obeying certain epistemic deontological duties.

In Chapter 2, I respond to a theory that holds we can normatively assess beliefs directly, because believers have compatibilist control over their beliefs, and this is

sufficient to ground obligations to believe.[3] These doxastic compatibilists argue that we have no need to appeal to indirect assessment to give an account of epistemic normativity. In this paper, I argue that doxastic compatibilism fails to ground obligations to believe for two reasons. First, beliefs lack essential features that actions possess that make actions voluntary, for example the feeling of voluntariness. Hence, agents cannot have compatibilist control over their beliefs. Second, a compatibilist account of obligations to believe cannot account for why some beliefs are morally wrong, or why some actions done based on evidentially justified beliefs are morally wrong. Compatibilists will need to incorporate a theory of epistemically permissible action to account for this.

In Chapter 3, I argue for the claim that my version of deontic epistemic normativity, that epistemic duties consist in performing certain actions causally upstream of one's beliefs, is a normative standpoint that is important in its own right. I do this in part by rebutting Trent Dougherty argument that epistemic normativity fails to be an appropriate normative standpoint: one has no practical reason to abide by epistemic norms.[4], I object to Dougherty by arguing that (i) one can have a reason to abide by epistemic norms if one commits oneself to those norms, and (ii) Dougherty's objections fail to show that epistemic obligations do not matter universally.

---

[3] Steup, "Belief, Voluntariness and Intentionality"; Heller, "Hobartian Voluntarism"; Hieronymi, "Responsibility for Believing"; Ryan, "Doxastic Compatibilism and the Ethics of Belief."
[4] Dougherty, "Reducing Responsibility," December 1, 2012.

Chapter 1 - It is Justified, but Should You Believe It?

A deontic theory of epistemic justification makes use of deontic concepts such as permissibility and obligation, to analyze the concept of epistemic justification—the property of belief that makes it appropriate to be held by a believer. The crux of this position is that a belief is unjustified if and only if it is impermissible to believe it. According to a famous objection from William Alston, if Ought Implies Can is true, deontic theories render the verdict that even false beliefs unsupported by evidence can be justified, since it is impossible to believe otherwise.[5] Thus, it seems that deontic epistemic justification fails to be a good account of epistemic justification.

In this paper, I argue that, even though Alston's argument against deontic epistemic justification as a theory of epistemic justification is sound, epistemic permissibility is still important. Epistemic permissibility, properly construed, provides an action-guiding theory of how epistemic agents ought to act. I provide two different potential theories of epistemic permissibility that bear further analysis.

1.1 Two Theories of Epistemic Justification

A primary focus of epistemology is on the question of what beliefs are worthy of being held. Jaegwon Kim, for example, holds that modern epistemological discourse is dominated by two related questions: "What conditions must a belief meet if we are justified in accepting it as true? and What beliefs are we in fact justified in accepting?"[6] These two

---

[5] Alston, "The Deontological Conception of Epistemic Justification."
[6] Kim, "What Is "Naturalized Epistemology?," 381.

questions revolve around the concept of *epistemic justification*. Epistemic justification is that property of a belief that makes it appropriate to be held. For example, if John is justified in believing that 2+2=4, then it is appropriate for him to believe it.[7]

However, without an elucidation of what makes holding a belief appropriate, the concept of epistemic justification is incomplete. Traditionally, epistemologists have held that epistemic justification is an evaluative notion: a belief is epistemically justified based on whether it is held on adequate grounds. For example, some hold that a belief must fit with the believer's evidence in order to be justified; others hold that it must be formed by a reliable belief forming mechanism to be justified.[8] Evaluative justification has been the traditional notion of justification due to its intuitive appeal. Suppose that John believes that the square root of 9,604 is 98 based on a lucky guess. While John is correct about this, he does not know the square root of 9,604 is 98. He does not hold this belief on adequate grounds; he is not epistemically justified in believing it.

However, some epistemologists have tried to go further than that. "Justification," they argue, should be understood deontically, since this is the way the term is used in ethics and the law.[9] Call this type of justification "deontic justification." Proponents of deontic

---

[7] With his two questions, Kim actually makes a distinction between two related notions of epistemic justification: *propositional justification* and *doxastic justification*. A belief is propositionally justified for some believer just in case that believer has appropriate grounds for believing some proposition. One is doxastically justified in holding a belief just in case one's belief is propositionally justified, and one bases one's belief on those grounds which propositionally justify one's beliefs. For more on this see: Pollock and Cruz, *Contemporary Theories of Knowledge*, 35–36; Kvanvig, "Propositionalism and the Perspectival Character of Justification," 8. For this paper's purposes, we set aside the question of doxastic vs propositional justification.

[8] For examples of the evidence based justification, see: Earl Conee and Richard Feldman, *Evidentialism: Essays in Epistemology* (Oxford University Press, 2004). For examples of the reliable process account of justification see: Alvin I. Goldman, "What Is Justified Belief?," in *Epistemology. An Anthology*, ed. Ernest Sosa and Jaegwon Kim (Blackwell, 1979), 340–353.

[9] See, for example: Berman, "Justification and Excuse, Law and Morality."

justification argue that a belief's justificatory status is determined by whether an agent abides by her epistemic obligations in forming that belief. Every believer has certain epistemic duties. The precise content of a believer's epistemic duties on deontic justification is outside the scope of this paper. Some hold that a believer must conform her beliefs to her evidence or must not form beliefs on the basis of bad inference rules—like the fallacy of affirming the consequent. Whatever the correct account of epistemic duties is, for any belief a believer has, if she discharges her duties with respect to that belief, then her belief is epistemically justified. Deontic justification, then, consists in duty fulfillment.

In a sense, deontic epistemic justification is a kind of evaluative justification. It holds that a belief is justified only if it is held on adequate grounds. The primary difference rests in deontic justification's emphasis that adequate grounds consist in fulfilling one's epistemic duties. Deontic justification holds that epistemic justification concerns the rightness of behavior of the believer. An epistemically unjustified belief, on this model, is one held as a result of a believer's wrongful behavior. Further, the holder of an epistemically unjustified belief the appropriate object of criticism for that belief. Deontic justification gives epistemologists a starting point for understanding the normativity of belief, what epistemic wrongdoing consists in, how believers ought to guide their beliefs epistemically, how a believer is responsible for her beliefs, and under what conditions she is blameworthy for them. These are all attractive features of deontic justification for anyone who is interested in epistemic normativity.

Yet, problematically, deontic justification cannot be the correct account of epistemic justification. Following William Alston's famous objection, we can show the

faults of deontic epistemic justification in two steps.[10] First, there is the principle of Ought Implies Can. Standardly, ethicists hold that an action is obligatory only if it is feasible for the agent to perform the action.[11] Suppose that I originally have an obligation to run in a marathon. Yet, before the race I break my leg rendering me unable to compete in the marathon. It would make little sense to hold that I still have an obligation to run in the marathon. Obligations are action-guiding; one's obligations should provide reasons for performing a certain action. In other words, one's obligations should be able to guide one's choice in a specific matter. A theory would fail to be action-guiding if it constantly advised an agent to do the impossible.[12]

The second step of the objection holds that believers do not have any direct control over their beliefs. This is "doxastic involuntarism." One cannot, by a mere act of willing, decide to believe one way rather than another. Instead, beliefs seem to be things that happen to us. Take, for example, my belief that there is a computer in front of me. Try as I might, I cannot decide to believe differently than I do right now. On my evidence, it is clearly the case that there is a computer in front of me. Thus, I believe there is a computer in front of me. Nor is this an aberration. For most, if not all, of one's beliefs, one cannot affect what one will believe.

Together, doxastic involuntarism and Ought Implies Can create a problem for deontic epistemic justification. Let us assume, for simplicity, that believers have epistemic

---

[10] See: Alston, "The Deontological Conception of Epistemic Justification."

[11] See: Graham, "'Ought' and Ability"; Vranas, "I Ought, Therefore I Can."

[12] One might be tempted here to make the following kind of claim: It ought to be the case that I run in the marathon. Let us call this an "ideal ought": the world ought to be a certain way. An ideal ought, I take it, does not imply feasibility. Indeed, part of the motivation for asserting the ideal ought is that I cannot run in the race. But the ideal ought is surely different from the kind of obligation mentioned here. Ideal oughts merely claim that the state of the world ought to be different while the kinds of oughts discussed here hold that I ought to behave a certain way.

obligations to form beliefs based on their evidence. Suppose that Jane has a belief P that she forms contrary to her evidence. Intuitively, Jane's belief P is not epistemically justified. In order to fit with our intuitions, deontic epistemic justification should render the verdict that Jane's belief is unjustified. However, if doxastic voluntarism is right, Jane cannot help but believe P. If Jane cannot help but believe P, then, on Ought Implies Can, Jane has no obligation not to believe P. Recall that on Ought Implies Can Jane only has an obligation to not believe P, if not believing P is feasible. Yet, if Jane has no obligation not to believe P, then it is permissible for her to believe P.[13] Thus, actually, Jane's belief P is deontically epistemically justified.

Even worse however, this problem for deontic epistemic justification is generalizable. Due to doxastic involuntarism most, if not all, of any believer's beliefs will be deontically epistemically justified. Not only does deontic epistemic justification fail to cohere with our intuitions about what beliefs are epistemically justified, it fails to provide any meaningful judgements about which beliefs are justified and which are not.

1.2 Actions as Epistemically Obligated

If Alston's argument is sound, deontic justification cannot be the correct notion of epistemic justification. Some other type of evaluative justification must be the correct notion, since it does not rely upon believers having control over their beliefs. This has led many epistemologists to abandon deontic justification and epistemic permissibility, holding that its only usefulness is to provide an account of epistemic justification. This, I

---

[13] As stated previously, some choice is permissible if one violates no requirements.

argue, is too quick. Recall that there were other merits for deontic justification over and above its being an account of epistemic justification. It is a solid foundation for the project of epistemic normativity. It can provide an account of how believers ought to act epistemically. It can provide an account of when agents are criticizable for their behavior epistemically. For these reasons, it would be desirable to preserve deontic justification even if it is not the proper account of epistemic justification.

Before moving to this argument, however, I want to briefly touch on two other responses to Alston's argument. Firstly, *doxastic compatibilists* argue that all Alston has shown is that agents lack libertarian control over their beliefs. However, Alston has not shown that believers lack compatibilist free will. Roughly, believers beliefs are generally reasons-responsive, which is sufficient for compatibilist control.[14] This argument presents an interesting challenge for doxastic involuntarism, and I will address it more fully in Chapter 2. However, at this stage, I will set it aside. That believers have compatibilist control over their beliefs is controversial, and I think we can provide an account of epistemic normativity without appealing to such a controversial concept.

The second response to Alston is that Ought Implies Can is false. We do not need to be able to bring about some state of affairs in order to be obligated to do so. Richard Feldman, the most prominent proponent of this objection, argues that epistemic obligations are like obligations to pay one's taxes.[15] One is required to pay one's taxes regardless of whether one is able to do so. That is because paying one's taxes is a role-obligation: an obligation one has in virtue of the role that one occupies. According to Feldman, epistemic obligations are role-obligations one has in virtue of being a believer. Thus, it does not

---

[14] See: Steup, "Belief Control and Intentionality," 2012; Hieronymi, "Responsibility for Believing."
[15] Feldman, "The Ethics of Belief"; Feldman, "Epistemic Obligations."

matter whether beliefs are involuntary, since one has role-obligations regardless of whether one has control over fulfilling them.

I do not wish to object to Feldman's account here. I think there is something importantly right that we intuitively treat certain obligations as obligatory regardless of our control over them.[16] However, this paper is especially concerned with that aspect of normativity that is action-guiding. The role-obligation account of epistemic normativity that Feldman provides is incapable of doing this: Feldman cedes that believers do not have control over their beliefs. The principle Ought Implies Can seems to importantly safeguard the action-guidingness of one's obligations.[17] Thus, I shall attempt to provide an account of epistemic normativity that is consistent with Ought Implies Can.

Fortunately, we have the tools to give a theory of epistemic normativity that does not require doxastic compatibilism or denying Ought Implies Can. Believers can still be in control of their doxastic states even if those states are not under believers' direct control. Suppose that John is trying to determine what some scientific law is. He could choose to consult a scientific textbook or to do nothing. If he chooses to consult the textbook, he will develop belief based on good evidence about the scientific law, whereas if he does nothing, he will not develop any beliefs.[18] This example opens a new avenue for the control that

---

[16] I do think, however, that our intuitions are probably wrong here. Instead, any obligations we have when we are unable to meet some obligation are probably obligations to do the next best thing. See, for example: Vranas, "I Ought, Therefore I Can."

[17] See, for example: Mason, "Consequentialism and the 'Ought Implies Can' Principle." However, we should point out that there is some dissent here. Brian Talbot argues that theories that deny Ought Implies Can are better to suited to account for action-guidingness. Since this paper is not specifically concerned with this debate, we shall set it aside here. See: Talbot, "The Best Argument for 'Ought Implies Can' Is a Better Argument Against 'Ought Implies Can.'"

[18] Alston, himself, first raises this possibility. However, since he is concerned with epistemic justification, he quickly disregards it because it fails to provide a satisfying account of epistemic justification. More recently, Rik Peels and Nikolaj Nottelmann have defended this claim. See: Alston, "The Deontological Conception of Epistemic Justification"; Peels, "Responsible Belief and Epistemic Justification"; Nottelmann, "The Deontological Conception of Epistemic Justification."

believers have over their beliefs. Neither of the choices available to John are doxastic states. Instead, they are actions. Actions are the bearers of epistemic permissibility, not beliefs.

There are two good reasons to prefer an action-based account of epistemic normativity to a belief-based account. First, for those who hold that agents are free, it is uncontroversial that actions are generally free. Second, actions have epistemic consequences. As the example shows, actions are often causally upstream of beliefs. John's decision to consult the textbook or not, in this case, will determine whether he develops the belief about the scientific law or not. He cannot, by an act of the will, choose to believe the scientific law. However, by his actions, he can bring about that he believes the scientific law. Proponents of deontic epistemic justification, then, were looking in the wrong place for epistemic permissibility. Rather than epistemic permissibility being a property of non-voluntary doxastic states, epistemic permissibility should be a matter of voluntary actions that causally contribute to epistemic states.

In order to make this claim, we must provide an account of what a believer's epistemic obligations are, and when an action is epistemically permissible. The goal for the rest of this paper shall be to provide several accounts of the epistemic permissibility of actions that are not susceptible to the challenges raised for the deontic justification of beliefs above. My goal is not to defend a particular view here, but to show that there are a number of ways to build up an account of action-guided deontic justification.

Before doing this, however, I want to touch upon a potential objection that one might have to this project. The objection goes as follows. The only reason that deontic justification was important was because of its potential as an account of epistemic justification. If deontic justification is not the correct account of epistemic justification,

then there is no such thing as epistemic permissibility. Even if epistemic normativity does exist, we have no reason to care about it since, again, it is not an account of epistemic justification. This is an important objection since, as I said above, the reason not to abandon epistemic permissibility is due to its explanatory powers for epistemic normativity. If epistemic permissibility does not exist or we have no reason to care about it, then the rest of this paper would be incoherent or meaningless. Thus, before moving on, we must respond to this objection.

First, let us address the first prong of the objection that epistemic permissibility does not exist. As to this objection, there does not seem to be much going for it. There is no reason that we cannot meaningfully talk about epistemic permissibility of actions in regard to how they lead to beliefs or lack of beliefs. In fact, as Dale Dorsey has said, we can talk meaningfully about all sorts of different normative viewpoints: we can speak coherently about the demands of morality, self-interest, etiquette, even the demands of medieval chivalry.[19] The epistemic permissibility of actions provides us just one more normative viewpoint to consider. That the epistemic permissibility of actions does not provide an account of epistemic justification does not render it incoherent. The more important consideration is whether this normative viewpoint matters. Most will agree that even if medieval chivalry makes demands on us as agents, these demands do not matter. We are not, in any way, bad agents for having failed to abide by the obligations of medieval chivalry. Thus, we must turn to the second prong of the objection that epistemic permissibility does not matter.

---

[19] See Dorsey, *The Limits of Moral Authority*, 9.

There are two ways that one might argue that epistemic normativity matters. Firstly, one might attempt to show that epistemic normativity is an integral part of any agent's practical rationality. Given this, epistemic normativity obviously matters. Another way, and the one that I choose, would be to argue that epistemic normativity can matter to an agent, if she chooses for it to matter to her. Agents choose particular standpoints that they value all the time. For example, when an agent forms a friendship, norms of friendship matter to her in a way they don't beforehand. There is no reason to think that epistemic normativity cannot work in a similar way to friendship.[20] For now, we shall set aside this objection. However, we shall return to it in Chapter 2.

1.3 Two Models of Epistemic Permissibility

In this section, I shall offer two potential models of epistemic permissibility. However, before doing so, there are a few questions about the nature of epistemic permissibility that we must consider. My goal in raising these questions is not to provide a definitive answer to each of them. Rather, I raise them to point out that an answer to each of these questions will affect the account of epistemic permissibility offered and to provide avenues for future research.

The first question we must consider is whether obligation is belief-relative, evidence-relative, or fact-relative. On a *belief-relative* conception of epistemic obligation, what is permissible is determined by what an agent believes. For example, if John believes that he has an obligation to save 1,000 lives and dropping a coin in a well will save 1,000

---

[20] For a more in-depth version of this argument, see: Chapter 2.

lives, even though this belief is clearly false, it is permissible for John to drop the coin in a well. Generally, most reject the belief-relative conception of obligation as too permissively subjective. For example, John could have an irrational belief that burning a forest will bring prosperity to the world. If true, that might make burning the forest permissible. Yet, that John believes it seems insufficient to make it permissible. Due to its widespread rejection, I will set the belief-relative model aside. On the *evidence-relative* conception of obligation, whether something is permissible is determined by the agent's evidence. Suppose that Jane has good evidence that her donating a sum of money to charity will save 1,000 lives and she is obligated to save 1,000 lives. Now, even if it turns out to be false that donating this sum will save 1,000 lives, on the evidence-relative view she has an obligation to donate that money to charity. Finally, on a *fact-relative* conception of obligations, a believer's beliefs are irrelevant to determining what is obligatory for that agent. Instead, all that is important is what, in fact, obtains about the situation. Thus, it is obligatory for James to give money to charity to save 1,000 lives, given that he has an obligation to save 1,000 lives, only if that money will, in fact, save 1,000 lives.

The important question, then, is whether we should understand epistemic obligations on a fact-relative or evidence-relative conception. I shall not argue for one or the other here, but I shall note a few things. First, offering one account does not exclude offering the other. It is perfectly intelligible to think that there are (at least) two kinds of normative assessment: an evidence-relative and a fact-relative one. Picking one is not a repudiation of the other. Second, offering a fact-relative model is simpler than offering an evidence-relative model. We need not concern ourselves with the evidence that an agent has or what she takes her obligations to be. Rather, we simply need to concern ourselves

with what, as a matter of fact, a believer's epistemic obligations are, and how her actions actually affect her doxastic states. For this reason, as a simplifying assumption, I assume a fact-relative model of obligations. Do not take this to mean that the models I shall offer cannot be formulated on an evidence-relative model. It is merely that such a model would be more complicated and would be outside the scope of this paper: providing potential models of deontic epistemic obligations.

Next, we must touch upon a debate concerning the nature of our causal powers over our beliefs. Among those who claim that our indirect causal power over our doxastic states is what grounds epistemic permissibility, there are two camps. Some hold that believers have indirect control over their doxastic states. By "indirect control," I mean that a believer can *intentionally* bring about (or at least increase the objective chance of) some doxastic state by means of her actions and can refrain from bringing about that doxastic state by her actions. Others hold that believers only have *non-intentional* indirect influence over their doxastic states: S's actions have causal influence over her doxastic states, but she cannot intentionally form or maintain a doxastic state.[21]

Rik Peels holds that this debate is important for determining the source of doxastic responsibility.[22] Such a question is surely important, but it is tangential to the question of the nature of epistemic obligations. Doxastic responsibility concerns questions of whether one can be the appropriate object of praise or blame, whereas epistemic obligations detail precisely how one must act to behave permissibly. One need not be responsible for one's

---

[21] See Peels, *Responsible Belief*, 67.
[22] See Peels, 72.

beliefs if one fails to meet an epistemic obligation, nor does one necessarily need to be subject to an epistemic obligation to be responsible for one's beliefs.[23]

One way to make Peels' claim relevant to the topic at hand is to interpret his claim as that whether we bear indirect control over our doxastic states or merely indirect influence determines the content of our doxastic obligations. If a believer has indirect control over her beliefs, then she has control over whether she brings about a certain belief or not. If so, then her obligations can be spelled out in terms of the doxastic states that her actions bring about. For example, if her action will bring about three bad beliefs that may render it impermissible to perform. However, if believers only have indirect influence then they can at most be assessed for the way their actions affect things they can foresee, such as their dispositions to form beliefs. For example, Peter may not be able to foresee that reading a newspaper will cause him to have a justified belief that the President is in France, but he can foresee that reading a reputable newspaper will cause him to be better informed and more reliably produce good beliefs.

For the purposes of this paper, it does not matter much whether we endorse indirect control or indirect influence. We can offer our two theories of epistemic permissibility under either framework. There are a few important things to note here, however. First, if the correct account of epistemic obligations is evidence-relative, then we might have good reason to prefer indirect influence models to indirect control. Believers tend to be rather unclear about the causal influence of their actions on their doxastic states. Thus, the idea that a believer could reliably intentionally bring about some doxastic state as a result of her action is unconvincing. Second, specifying obligations in virtue of the beliefs they bring

---

[23] For an example of the latter see: Hieronymi, "Responsibility for Believing."

about makes the epistemic nature of epistemic permissibility clearer. That some action will give me better belief forming dispositions does not necessarily generate epistemic value, especially if that belief forming disposition never manifests. For example, there is nothing particularly valuable about being an expert on distinguishing a 15-foot human being from a 16-foot human being if there are no human beings of that size to identify. For this reason, since we are assuming fact-relative obligations, we shall assume that the indirect control model is correct.

Having considered these preliminary questions, I shall now offer two models of epistemic permissibility. I shall not construct these theories from scratch. Instead, I shall borrow heavily from the structure of moral theories in ethics. We can roughly divide theories of moral permissibility into two distinct categories. First, there are *consequentialist* theories that hold that permissibility is entirely a function of the (epistemic) goodness of the consequences of one's actions. An action is permissible if and only if its consequences are good enough relative to its alternatives. Second, there are *deontological* theories that hold that permissibility is not *fully* determined by the consequences of one's actions. In addition, there are certain rules that one must follow. Epistemic permissibility consists, at least in part, in following certain rules whose content is determined entirely independent of the consequences.[24]

To begin, let us start with consequentialist theories. On *epistemic consequentialism,* an action's permissibility depends on the epistemic value of its outcomes.[25] The motivation for epistemic consequentialism is twofold. First, many think that the only important

---

[24] Rule-utilitarians are consequentialists who think that the best way to maximize the consequences of one's actions is to follow rules that tend to maximize the goods. These are clearly not deontologists. This is why we must specify that the justification of the rules is not grounded in the consequences of obeying them.
[25] For the rest of this paper, I shall use "epistemic value" and "epistemic good" interchangeably.

consideration for the permissibility of an action on any normative viewpoint is how good the consequences are.[26] If this is the case, epistemic permissibility should be no different. Thus, we should endorse epistemic consequentialism. However, there is a special consideration for epistemic consequentialism over and above this. The only reason that actions are important to epistemic permissibility is because they affect which beliefs a believer has, and the grounds upon which one has a belief. If some action did not affect one's beliefs or the grounds for one's beliefs, it would be irrelevant to epistemic permissibility. Thus, it looks like, in a very real way, the doxastic consequences of an action are essential to epistemic permissibility. This is a good reason to be a consequentialist.

Since a theory of consequentialism is focused on promoting good consequences and avoiding bad consequences, in order to offer an account of epistemic consequentialism, we must determine what are good and bad epistemic states of affairs. Here, I take epistemically good states of affairs to be epistemically valuable states, and bad epistemic states of affairs to be epistemically disvaluable.

It certainly seems as if there are epistemically valuable states. For example, knowing some scientific fact seems more valuable than truly believing that fact on the basis of a lucky guess. Epistemologists have identified many potentially epistemically valuable states such as true belief, justified belief—in the sense required for knowledge—knowledge, rational belief, cognitive achievements, and good epistemic character.[27] Epistemic consequentialism's plausibility does not depend on any of these states actually

---

[26] See for example: Mill, *Utilitarianism*; Sidgwick, *The Methods of Ethics, 7th Edition*; Sinnott-Armstrong, "An Argument for Consequentialism."

[27] Over the course of this paper, I have held that actions are the proper objects of epistemic normative assessment because of their causal connection to doxastic states. If, however, good epistemic character or cognitive achievements is what is epistemically valuable, it will be actions causal connection to these things that determine whether the action is permissible.

being epistemically valuable. So long as something is epistemically valuable, and actions can produce and fail to produce these valuable states, we need not worry.

One worry that might arise here is if it turns out that there is more than one epistemic value and there is no master value that all other kinds of epistemic value can be reduced to. Suppose, for example, that justification and rationality are both of epistemic value but are incommensurable to one another. Here, the worry might be that if justification and rationality are incommensurable, then there is no way to compare states of affairs with differing amounts of both. For example, is it better to bring about three justified beliefs or three rational beliefs? This kind of value pluralism might make the job of determining the optimal state of affairs harder, but value pluralists have argued that it is possible in different domains. For example, Thomas Nagel has suggested that practical wisdom allows one to rationally prefer one state of affairs to another.[28] Since this issue is not central to this paper, we shall put it aside here.

One important question is what kinds of epistemic value the epistemic consequentialist should take into account. Should the consequentialist be an egoist, taking into account only the epistemic interests of the believer, or a universalist and take everyone's interests into account?[29] One's actions have the ability not only to affect one's own doxastic states but others as well. An example might better illustrate this concern. Suppose that Anna is an expert on the field of evolutionary genetics. She could give a lecture to an audience that greatly enhances the audience's knowledge of the field, but she

---

[28] Nagel, "The Fragmentation of Value."
[29] We can raise this question for epistemic deontology, but egoism has historically been more associated with consequentialism than deontology. See: Regis, "What Is Ethical Egoism?"

would gain no epistemic benefit from it. Could Anna have an epistemic obligation to give the lecture even though it has no possible epistemic import to her?

Most normative epistemology currently being done assumes an egoist model of permissibility.[30] This may be because of deontic epistemic justification history as an account of the epistemic justification required for knowledge—justifying one's own beliefs is inherently egoistic. Given that we have abandoned the connection between epistemic permissibility of actions and epistemic justification, it is plausible that one's epistemic obligations could be more universal. My guess is that, plausibly, believers do have universal epistemic obligations. For the sake of simplicity, however, we shall follow tradition and hold that deontic epistemic justification is egoistic.[31]

Let us return to epistemic consequentialism. As stated previously, consequentialism holds that an action's permissibility depends on promoting sufficiently good consequences relative to one's alternatives. On the epistemic consequentialism we have assumed, then, an action is permissible if and only if it promotes sufficiently good epistemic consequences. From this starting point, we shall now divide epistemic consequentialism into two types. First, we can assess whether certain action-types promotes sufficiently good consequences. Second, we can assess whether the action-token promotes sufficiently good consequences.

Let us call the first type "rule-consequentialism." On rule-consequentialism, whether an action-type is permissible depends on whether the rules that promote the best

---

[30] See, for example: Peels, *Responsible Belief*; Steup, "Belief Control and Intentionality," 2012; Nottelmann, *Blameworthy Belief*, 2007.

[31] If, however, you are concerned with universalist deontic epistemic justification, I do not think it would be difficult to convert the theories I shall offer into universalist theories. For consequentialist theories, all one needs to do is include the effects that one's actions have on other believers' doxastic states into the equation of what the right action is. For deontological theories one need offer other-regarding rules into the equation. For example, do not lie to others.

outcomes permit such action-types. While we cannot provide a full list of rules here, here are some commonsensical rules: always engage in thorough evidence collection, reflect on the coherence of your beliefs, and weed out bad belief forming behaviors. Each of these has a powerful effect on what one believes. Following these rules greatly increases the probability of producing more epistemic value than less epistemic value. For example, Sally has a much higher probability of developing justified beliefs if she reflects upon whether any of her beliefs are inconsistent than if she fails to do so.

Let us call the second type of consequentialism, the one that holds that permissibility is determined by assessing whether a particular action-token promotes sufficiently good consequences, "act-consequentialism." For act-consequentialism, the permissibility of an action is determined by comparison to other feasible actions available. How good an action must be relative to its peers in order to be permissible depends on whether the account in question is maximizing or satisficing. On a maximizing account, the action must produce the most net-good among all feasible actions. On a satisficing account it must be produce a sufficient net good in order to be permissible. To illustrate, let us consider the following example: suppose that Peter has three feasible options: a, b, and c. None of the three options produces any disvaluable beliefs. Option a will produce two valuable beliefs, option b will produce 4 valuable beliefs, and option c will produce six valuable beliefs. On maximizing consequentialism, the only permissible action is option c, since it produces the greatest good among all feasible options. On satisficing consequentialism, all, none, or some of a, b, and c will be permissible; it depends on what is good enough. Suppose that satisficing consequentialism requires that Peter's action

produces at least 3 valuable beliefs in this case. That being the case, b and c will be permissible, and a will be impermissible.[32]

Since both satisficing and maximizing theories are concerned with net epistemic value they will allow for permissible tradeoffs between valuable and disvaluable beliefs. For example, an action that produces a large amount of epistemic value at the cost of a little epistemic disvalue may be preferable to an action that produces a little epistemic value and no epistemic disvalue.

Let us turn to our second theory of epistemic permissibility: epistemic deontology. Deontology holds that the permissibility of an action is not entirely determined by the consequences. On a narrower conception of deontology, epistemic permissibility consists entirely in following the epistemically correct rules. This theory is different than rule-consequentialism because, unlike rule-consequentialism, these rules cannot be justified by their tendency to bring about good consequences. Instead they must be justified on some other grounds. Again, the question of egoism or universalism in epistemic permissibility appears again. Since we have assumed epistemic egoism, the only epistemically relevant features of an agent's actions are those that are related to that agent's epistemic state.[33]

---

[32] Some might wonder how satisficing act-consequentialism interacts with Ought Implies Can. Suppose that in some scenario, no action produces consequences that crosses the threshold of permissibility. In this case, it might seem that one's obligation to perform a good enough action cannot be satisfied. There are a couple of solutions here. One is to allow that no action is permissible because no action is good enough. This would be an example of what Peter Vallentyne calls a Prohibition Dilemma. Another is to hold that if no actions cross the threshold of permissibility one is obligated to pick the best among all feasible actions. Even though none of them is good enough, if one is clearly better than the others, that is the one an agent is obligated to do. See: Vallentyne, "Two Types of Moral Dilemmas."

[33] I think that epistemic deontology will have even more plausibility if we are universalists. Every epistemic patient is worthy of consideration. Epistemic patienthood consists in the capacity for rational thought. One must never act to epistemically harm another. Patients have rights not to be put into a worse epistemic situation, or rights against not being treated as epistemic agents. Thus, we have obligations not to lie to epistemic patients.

Impermissibility on epistemic deontology will consist in epistemically harming oneself. For the purposes of this account, what harming consists in will be determined by one's account of epistemic value. For example, suppose that it turns out that justified beliefs are epistemically valuable. If so, harming oneself would consist in setting back one's capacity to form justified beliefs. This might consist in setting back: a believer's epistemic rationality, her ability to form beliefs based on evidence, her ability to make inferences, or her ability to discover the truth, among others. Importantly, the lists of harms will change based upon what is epistemically valuable. Thus, epistemic deontology might hold that impermissibility consists in intending to epistemically harm oneself or actually epistemically harming oneself. Suppose I intend to refrain from gathering new evidence in a case. This would constitute epistemically harming myself since it would directly go against my capacity to pursue truth. Such reasoning might also prohibit one from getting drunk, since that dulls one's ability to behave epistemically rationality.

We can also understand these duties from a rights-based perspective. Epistemic agents have epistemic rights that correspond to the relevant epistemic values above. One has a right not to have her ability to make inferences based on her evidence be harmed. On this egoist perspective, these rights would only be held against oneself. One problem with an egoist rights-based perspective is that many deontologists think that agents can choose to forfeit many if not all of their own rights.[34] If one holds all one's rights against oneself, it seems that one could choose to forfeit those rights in order to do what one wishes. We might try to block this by holding that epistemic rights are non-forfeitable. Perhaps we might argue that epistemic rights are to properties that are constitutive of being a believer

---

[34] For an example of side-constraint consequentialism see: Nozick, *Anarchy, State, and Utopia*.

and these are properties that one cannot rationally give up. I doubt such an argument will convince those who think all rights are, in principle, forfeitable. Nevertheless, we can avoid this problem by sticking to conceptions of deontology that hold one has an obligation not to (intentionally) harm oneself, rather than obligations to refrain from impermissibly infringing one's own rights.

Moving from a narrower to a broader deontology, there are also plausible combinations of deontology and consequentialism that I shall briefly mention. Some have objected to the fact that deontology permits actions with bad consequences and forbids actions with good consequences. Suppose that one could take a pill that would grant one 1,000 justified beliefs, at the cost of robbing one of one's rationality for a day. Epistemic deontology, at least as I have construed it, would forbid this action since it would harm one's epistemic personhood. It might advance one's status as a former of justified beliefs, for example, but it does so at the cost of harming oneself. This is never permissible. Nevertheless, it certainly seems that the positive consequences vastly outweigh the costs. This has led some to abandon pure deontology and accept a mixed version of deontology in which deontological norms can be overridden by the consequences if the consequences are good enough.

Likewise, some are convinced by the fact that there are deontological prohibitions on certain kinds of actions but are also convinced that deontic epistemic justification requires maximizing or satisficing the good. For this kind of theory, which we can call "side-constraint epistemic consequentialism," in any decision scenario one goes through a two-tiered process. First, one weeds out all the actions that violate deontic side-constraints (no matter how epistemically good their consequences). Next one picks an (the) action

from the remaining actions that meets the maximizing or satisficing requirement of epistemic consequentialism. As an example, suppose that three choices are available to an agent, she can go to the movies, refuse to gather any evidence, and go to the library. Assume refraining from gathering evidence would produce the most justified beliefs, going to the library would produce the second most, and going to the movies would produce a net negative. In the first step, side-constraint epistemic consequentialism weeds out refusing to gather new evidence. Now the best (or set of good enough) action(s) is the permissible action. This is going to the library.

In this section I have offered a few different versions of epistemic consequentialism and epistemic deontology that provide possible accounts of epistemic permissibility. These have the added benefit of not being subject to the Alstonian objection, as those who endorse that agents are free hold that agents have control over their actions and can be action-guiding. In the next section I will address some problems for each of these theories.

1.4 Epistemic Consequentialism or Epistemic Deontology

Before concluding, I shall address problems for consequentialism and deontology in turn. Consider first epistemic rule consequentialism. Suppose that Sally is in a choice-scenario where she can either go to the movies or go to the library. Suppose that the set of rules that, if generally followed, would maximize epistemic value includes a rule that requires Sally to go to the library. Suppose, however, that, in her particular situation, going to the movies will net more justified beliefs for Sally than going to the library, and, even worse, going to the library will produce only unjustified beliefs. If the goal of consequentialism is to

produce more justified beliefs, to insist that Sally stick to the rules and go to the library seems like nothing more than rule worship. Regardless, this is a problem for rule-consequentialism *simpliciter* and not merely the epistemic variety. Thus, if there is a solution for the general problem then there is a solution for epistemic rule-consequentialism.

Objections to act-consequentialism are well known. However, one that is especially interesting for epistemic act-consequentialism is the demandingness objection. For act-consequentialism, we are obligated to pick the best choice in every scenario. Now, whatever a believer's epistemic obligations are, it seems plausible that some decisions should be optional for her. One's life is not just about epistemology and so there should be some situations where one is epistemically permitted to do more than the optimific action. There might be some situations where a believer would gain more justified beliefs if she, for example, went to the library, but she shouldn't have to always pick the option with the better consequences; an epistemic theory shouldn't be so demanding**.** Nevertheless, consider a choice scenario in which Sally is faced with two potential options. She could go to the movies, or she could go to the library: going to the library will produce more justified beliefs than going to the movies. Now let's assume that Sally has no pressing need to go to the library. Intuitively, Sally has no epistemic obligation to go to the library. We can admit that it would have been epistemically desirable, from an epistemic perspective, for Sally to go to the library, but it is surely optional. Nevertheless, act-consequentialism will render the verdict that going to the library is epistemically obligatory. In fact, for any scenario in which one option has slightly better consequences than all the others, a believer is required to pick that option. Act-consequentialism seems too demanding.

However, there is much to be said in epistemic act-consequentialism's defense. Firstly, this objection will not apply to certain satisficing forms of act-consequentialism. The threshold for an action's being good enough might be so low that it permits Sally both going to the movie and the library. Secondly, we should not be too concerned with this epistemic demandingness objection. The demandingness objection finds its origin in ethical theory; part of its bite in ethics comes from a view that the dictates of morality should override all other considerations: if morality obligates one to perform some action, one must follow morality's dictates no matter what. If morality crowds out every other aspect of one's life and obligates one to do exactly one thing in every choice scenario, there is little room for personal choice. This is the source of demandingness.[35] Yet, unlike morality, epistemology holds no such primacy over other considerations. The requirements of epistemology seem nowhere near as important as those of morality. Even if there is just one epistemically right thing to do in any situation, that does not suggest that a believer is obligated, given all her considerations (moral, prudential, practical), to perform that action. When one's epistemic obligations conflict with, for example, one's prudential interests, there is little reason to suppose that one's epistemic obligations will win. Sometimes, an agent's personal interests will be more important than what action will produce the most justified beliefs. Demandingness holds no sway over epistemic obligations. Thus, act-consequentialists need not worry about it.

---

[35] It should be noted that the viewpoint that morality overrides all other considerations is not at all uncontroversial. See: Dorsey, *The Limits of Moral Authority*.

There are problems for epistemic deontology as well.[36] We might object to epistemic deontology's connection to the principle of Ought Implies Can. Suppose there is a situation in which an agent is presented with two feasible choices: either take a pill that causes one to lose all of one's rational capacities for a limited period of time, or refrain from gathering new evidence. Epistemic deontology might hold that both actions violate epistemic norms. Yet, if both feasible options are impermissible, then it seems as if one is obligated to do that which is infeasible. However, this is just a denial of the principle Ought Implies Can. On Ought Implies Can one can only be obligated to do what is feasible. The question is how such situations are conceptually possible. One of our assumptions must go.

There are three ways we might try to respond to this worry. First, we could hold that every feasible action is permissible, since the agent has no obligation to perform some other action. One has an obligation only if one can discharge it. One cannot discharge one's obligations in this case. This solution is similar to the way that Alston's objection falsifies deontic epistemic justification as an account of epistemic justification. Since one cannot fulfill one's obligation, one's action is permissible. However, whereas this solution rendered deontic epistemic justification a bad theory, this solution is not nearly as problematic for epistemic deontology. Most of the time it will be feasible for an agent to fulfill her epistemic obligations. Thus, this solution is plausible for epistemic deontology.

Second, we can appeal to secondary obligations. We might hold that in these cases, one's initial obligations no longer obtain, since it is impossible to permissibly discharge

---

[36] Like act-consequentialism, some problems for ethical deontology do not arise for epistemic deontology. Consider the famous emergency objection where someone can do something intrinsically bad to save 1,000,000,000 lives. This seems morally permissible, yet moral deontology holds that it is actually impermissible. Epistemic deontology is not subject to this objection, since it seems that there are no epistemic emergencies that are so bad that would warrant doing something intrinsically bad to gain really good consequences.

those obligations. Instead, one gains secondary obligations to do the best one can in these circumstances.[37] For example, even if both actions violate one's deontological obligations one might be less bad than the other. In this case, one should pick the less bad option, or, if both are equally bad, both are permissible.

Third, we could allow for the existence of prohibition dilemmas.[22] Prohibition dilemmas are dilemmas in which every option in the choice scenario would be impermissible because it is not good enough to be permissible. In this case, since every action would go against a deontological norm, we can hold that every action is impermissible, even though the agent has no obligation to avoid any of these actions. That every feasible option is impermissible does not entail that there is an obligatory non-feasible option. Thus, Ought Implies Can is not threatened by this example.[38] Like the first solution, we need not worry about the prevalence of these kinds of cases. Most scenarios will not be instances of prohibition dilemmas. Thus, they need not concern us.

Another challenge to epistemic deontology is that since deontology is not entirely concerned with consequences, it will require that an agent, in some situations, bring about less valuable states than she could have. Suppose that Sally could take a pill that would produce many valuable epistemic states but prevent her from being able to reason for the next hour, and it is impermissible for her to take the pill, either because on narrow

---

[37] This solution follows Peter Vranas' solution to problems for Ought Implies Can. See: Vranas, "I Ought, Therefore I Can."

[38] We should stop to point out that neither the second nor the third solution saves the original conception of deontic epistemic justification from Alston's objection that we have no control over our beliefs. If we appealed to the second solution, the best I can do in the case of belief is to believe what I actually believe. Thus, the second solution would render all of my beliefs permissible. The third solution fails to provide an action-guiding account. Prohibition dilemmas would correctly label unjustified beliefs as impermissible and justified beliefs as permissible, yet agents cannot use the norms of prohibition dilemmas to alter their beliefs. Doxastic states are still involuntary. Thus, neither of these solutions will work for belief.

deontology the consequences do not matter or on broad deontology the consequences are not sufficient to override the rule not to harm oneself. This prohibition seems odd since the cost to herself is very low and the benefit is very high. Many see this as a strike against deontology. Yet, act-consequentialism has a similar problem. Since act-consequentialism is concerned with net-consequences it permits an agent to incur some epistemic costs if the epistemic benefits are worth it. This means that one could incur great costs if the net-benefit was sufficiently high, and this might be implausible. Nevertheless, these and other problems will need to be worked out in order to give a fully worked out theory of epistemic permissibility.

1.5 Direct Doxastic Normative Assessment as a Fiction

There is one last worry for the picture of epistemic normativity offered in this chapter. I have motivated the appeal to epistemic normativity by means of ordinary language. In ordinary language, we normatively assess others beliefs with phrases like "You shouldn't believe everything you hear." I have argued that it is preferable to have a deontic normative theory that captures these phrases of ordinary language, then rejecting them as clearly wrong out of hand.

However, if we follow this logic, then one might criticize my theory for failing to adequately capture the ways in which we assess others beliefs. When I assess someone's beliefs, I typically do so by means of direct assessment. I criticize the belief; I do not mean to criticize the actions that led to that belief. In other words, I say "you shouldn't believe that", rather than "you shouldn't have acted in a way to gain that belief" or "you should

take steps to get rid of that belief." All three of these, the objection might go, seem perfectly plausible, but the latter two do not seem substitutable for the direct assessment. Therefore, my account has not adequately captured ordinary language epistemic normativity.

One response that I shall briefly offer here, though I do not have time to fully defend, is to adopt a fictionalism about direct normative assessment. All direct doxastic normative assessments are false. When someone says, "you ought to believe that," if, indeed, they mean you ought to immediately alter your beliefs, then they are asserting something false. Ought Implies Can. Since one cannot directly alter their beliefs, they need not. However, ordinary language is still useful, even if false. I shall argue that it is much easier to directly assess one's beliefs then to assert the appropriate indirect assessment. For example, it is much easier to assert "You ought not to believe that vaccines cause autism" than it is to assert "You ought not to have taken those steps that caused you to believe vaccines cause autism."

There is another reason why it is useful to employ this fictional language, rather than using indirect doxastic assessment. On my account, it is relatively easy to figure out how one should act now to abide by one's present and future doxastic obligations. You should take steps to alter your beliefs to promote valuable epistemic states or conform with your deontological duties. However, when evaluating someone's current bad beliefs, it is difficult to determine precisely which actions causally brought about the doxastic state one now has, and to what extent each action is causally responsible for that doxastic state. Simply put, it is often impossible for me to determine precisely where someone who has a belief, they shouldn't went wrong. The substitution "you shouldn't believe that" is a useful generic assessment that covers these cases. Even if we cannot pinpoint which (set of)

action(s) you performed led to this bad doxastic state, I am confident there was a voluntary action that you took that led you to this moment that you shouldn't have taken. Direct doxastic assessment, I offer, is a fiction, but a useful one.

1.6 Conclusion

As we have seen, while belief-based models of epistemic permissibility fail because agents do not have control over their doxastic states, we can still provide action-based models of epistemic permissibility. Such models cannot provide an account of epistemic justification, but we have other reasons to be interested in them: namely their ability to provide an action-guiding account of how agents ought to behave epistemically. We can provide two different types of theories of permissibility: epistemic consequentialism and epistemic deontology. While neither is free from objections, they provide a valuable starting point for continued research into the nature of epistemic normativity.

Chapter 2: Doxastic Compatibilism's Normative Problem

On one fairly influential argument, there can be no obligations to believe or disbelieve anything. Believers cannot change or form new their beliefs at will. Thus, they cannot have any obligations to have certain beliefs. Epistemologists who want to account for obligations to believe generally do so by means of the indirect influence that agents have over their beliefs. An agent's actions causally influence her beliefs. Proponents of this *indirect influence account* argue that an agent can have obligations to act in certain ways so as to produce certain kinds of beliefs.

However, *doxastic compatibilists* hold that we need not countenance indirect influence to account for obligations to believe. The prior argument against the normativity of belief is unsound. Believers have compatibilist control over their beliefs. Roughly, they have control over their beliefs even if they do not have the ability to believe otherwise. This is sufficient for the normativity of belief. In this paper, I argue that doxastic compatibilism fails to be a strong challenge to indirect influence accounts for two reasons. First, if there are to be obligations to believe, then beliefs must be formed intentionally, regardless of whether believers seem to have compatibilist control over their beliefs. Beliefs are not, however, formed intentionally. Second, doxastic compatibilism offers an incomplete account of the normativity of belief in that it fails to capture anything but basic epistemic obligations that an agent has. Notably, doxastic compatibilism cannot account for the existence of moral obligations to believe. I shall argue that doxastic compatibilism will require an indirect influence account as a supplement in order to fill in these

deficiencies. But, if so, then doxastic compatibilism is no opponent to indirect influence accounts. It is, at best, a supplementary theory.

2.1. Doxastic Voluntarism and the Appeal of Doxastic Compatibilism

The idea that agents have obligations to believe certain propositions has intuitive support. In everyday life, we see examples of persons criticizing others for their beliefs in statements like "You shouldn't have believed that." Certain cases in epistemology even seem to suggest the existence of doxastic norms. For example, Laurence Bonjour presents the case of the clairvoyant that forms beliefs based on clairvoyant intuitions that he has no evidence are reliable.[39] Seemingly, part of the problem with such a case is that the clairvoyant seemingly violates his doxastic obligations in forming these beliefs. The most straightforward way of explaining these cases is that believers have obligations to have certain beliefs rather than others.

Commonly, obligations to believe are identified with epistemic obligations. Since epistemology is concerned with what makes a belief good, the connection between right belief and belief in accordance with these epistemic goods seems obvious. For example, one might think that one should believe only that which is true, or for which one has justification, or which is rational to believe, and so on. However, in principle, believers could have other kinds of obligations to believe. For example, one might have a *moral* obligation not to have sexist beliefs. A complete theory of the deontic normativity of belief

---

[39] Take, for example the case of Norman the Clairvoyant. See: Bonjour, "Externalist Theories of Empirical Knowledge."

should be able to account for every kind of obligation to believe.[40] Since the connection between epistemic norms and doxastic norms is so intuitive, for the rest of this paper, except when otherwise noted, I focus solely on epistemic doxastic norms.

Yet, for all its intuitive support, there is a problem for the concept of obligations to believe. Many think that an agent has an obligation only if that agent can satisfy that obligation. This is commonly known as the principle Ought Implies Can. Intuitively, one can be obligated to bring about some event or state of affairs only if one can voluntarily bring about that event or state of affairs. If, for example, Jane has broken her leg and cannot leave her home, it seems implausible that she has an obligation to run outside and save a drowning child. Since this action is not feasible, she cannot be obligated to do it. Note, however, that this has been challenged. Notably, in the doxastic obligation literature, Richard Feldman argues that Ought Implies Can is false. Persons can be obligated to do things outside of their control in virtue of the roles they occupy. For example, a parent has an obligation to be a good parent, even if he is incapable of doing so.[41] It is not the goal of this paper to defend the principle of Ought Implies Can. Since doxastic compatibilists seem to assume the truth of Ought Implies Can, I shall assume its truth as well.

The problem for obligations to believe is that beliefs are not the kinds of things that one can change at will. I cannot, for example, change my belief that 2+2=4 no matter how

---

[40] As far as I am aware, no one has attempted to give such an all-encompassing theory. Several have attempted to give a theory of doxastic ethical and epistemological duties. See, for example: Peels, "Responsible Belief and Epistemic Justification"; James, "The Will to Believe"; Clifford, *The Ethics of Belief and Other Essays*; Haack, "'The Ethics of Belief' Reconsidered."
[41] See: Feldman, "Subjective and Objective Justification in Ethics and Epistemology"; Feldman, "The Ethics of Belief."

hard I try.[42] Suppose, that John has an obligation to believe, now, that the earth is round, yet he believes that the earth is flat. On Ought Implies Can, John has this obligation only if he can meet it: only if he can believe, now, that the earth is round. Yet, John cannot, by a mere act of the will, change his belief right now.[43] He could, perhaps, take actions to change his later beliefs. He could educate himself or attempt to convince himself that the earth is round.[44] However, doing this would not satisfy an to believe that the earth is round now, since that is impossible. Instead, it would satisfy an obligation to take steps towards believing that the earth is round in the future. In short, John has done nothing wrong now by failing to believe that the earth is round. This leads to a reductio.

This realization has led those who want to preserve the normativity of belief to split into three camps: Proponents of *indirect influence accounts* argue that we should abandon obligations to believe, and, instead, locate the normativity of belief in voluntary acts that causally influence later doxastic states. Believers have obligations to act in certain ways because acting that way will bring about certain (kinds of) desirable doxastic states.[45] *Infeasabilism* denies that Ought Implies Can, arguing that doxastic normativity does not require this principle.[46] *Doxastic Compatibilism* holds that the argument from doxastic involuntarism fails to show that there is no normativity of belief. This argument does not

---

[42] For the origination of this argument see: Alston, "The Deontological Conception of Epistemic Justification."

[43] There are ways he could change his belief through an intermediary. Suppose, for example that he could take a pill that would cause him to believe that the earth is flat. While this would be an example of having voluntary control over his beliefs it would not be an instance of changing his beliefs at will.

[44] This suggestion is similar to what Pascal suggested for the one who wants to believe in God. See: Pascal, *Pensées*.

[45] There is some debate about whether one is responsible for the specific belief that one brings about or the kind of belief that one brings about: indirect control vs indirect influence. For an account of indirect control see Nottelmann, *Blameworthy Belief*, 2007.; for indirect influence see: Peels, *Responsible Belief.*. This debate is orthogonal to the paper topic and we can safely set it aside.

[46] Again, see: Feldman, "Epistemic Obligations"; Feldman, "The Ethics of Belief."

attempt to claim that believers can change their beliefs at will. Instead, the argument is that

believers do not require the kind of control over their beliefs that we have been discussing

in order for obligations to believe to exist.

This paper shall be concerned with objections to doxastic compatibilism. This

dissertation takes up a defense of an indirect influence account as the proper account of the

normativity of belief. In the first chapter of this dissertation, I offered some reasons to think

that regardless of the truth of Ought Implies Can, the account of the normativity of belief

offered by infeasabilism is inadequate.[47] This leaves doxastic compatibilism as the main

rival to indirect influence accounts. My goal in objecting to doxastic compatibilism is to

show that indirect influence accounts are the only plausible accounts of the normativity of

belief. Let us turn to our attention to doxastic compatibilism now.

According to doxastic compatibilists, the control mentioned in the argument against

voluntary belief is libertarian control: one has control over some belief if and only if one

has the ability to have that belief and not have that belief. However, there is another kind

of control that is posited in action theory: compatibilist control. Compatibilist control is

claimed to be the control that one has over one's mental states even if one could not do

otherwise. Frankfurt Cases provide intuitive inspiration for these compatibilist accounts.[48]

Take the hitman who is hired to assassinate a target, but who has a microchip implanted in

his brain so that he will be compelled to kill his target, if he decides not to. If the hitman

decides to kill his target without the microchip triggering, he seems to have control over

---

[47] See: Chapter 1.
[48] See: Frankfurt, "Alternate Possibilities and Moral Responsibility."

for this assassination, even though he could not have done otherwise. If compatibilism is plausible for actions, it may very well be plausible for doxastic states.

Doxastic compatibilism, as used in this paper, is the thesis that compatibilist control grounds obligations to believe.[49] One does not need libertarian control over one's beliefs in order to have obligations, even if Ought Implies Can is true. One need not be able to do otherwise given the exact same starting conditions and laws of nature, in order to have an obligation. Instead, one need just needs a weaker conception of alternatives: the ability to do otherwise, given distinct starting conditions. If one's beliefs proceed from one's agency, that is sufficient to ground obligations to believe.

Compatibilists offer three different ways of accounting for compatibilist control. I shall explicate each, here. The objections in this paper apply equally to each of these positions. The positions are simple doxastic compatibilism, higher-order doxastic compatibilism, and reasons-responsive doxastic compatibilism.[50]

Simple doxastic compatibilism holds that a believer has control over her doxastic states in the sense that, if she had decided to believe differently, she would have. An analogy with simple compatibilism in action theory is helpful here. A simple compatibilist might say that, even though an agent cannot refrain from doing some action, she still as control over that action, if it proceeds from her conscious willing as a result of a mixture of intentions, beliefs, and desires. If that mixture had been different, then she would have

---

[49] Some compatibilists offer accounts of compatibilist control to account for agent-responsibility: one is agent-responsible for some action if one had compatibilist control over its performance. Pamela Hieronymi, for example, offers her account of doxastic compatibilism specifically as an account of doxastic responsibility, denying the existence of doxastic obligations. However interesting these theories are, they are outside the scope of this paper. See: Hieronymi, "Responsibility for Believing.".
[50] These correspond with versions of compatibilism in action theory.

done differently. The action is hers in the sense that it flows from her will. Applied to doxastic states, a believer's beliefs are the product of her mental states: her evidence, her reasoning faculties, and so on. If those had been different, she would have believed differently. Thus, a believer has simple compatibilist control over her beliefs. The only example of simple doxastic compatibilism I can find in the literature comes from Matthias Steup (2008), who countenances this simple doxastic compatibilism, but quickly disregards it, thinking it fails to capture aspects about what we think is sufficient for free will. I agree and will set this aside.

The second version of doxastic compatibilism comes from Mark Heller (2000), who offers a *higher-order desire account* of doxastic compatibilist control. On this model, compatibilist control is a function of the coherence of one's higher order desires with one's will.[51] A higher-order desire is a desire to have certain desires. One's will amounts to one's lower-order desires that determine one's behavior. When one wills a certain outcome and the desires that produce one's willing cohere with one's higher-order desires, one has control over one's behavior. To give a simple example, the higher-order desire compatibilist would hold that a drug addict who sincerely wants to quit but who uses drugs because he is addicted lacks compatibilist control over his drug usage. While his lower-order desires may be to use drugs, he has a higher-order desire not to act upon those first order desires. The incongruity between orders of desires renders him not in control over his drug-usage. Yet, if someone flips a light-switch because they want to, and they endorse

---

[51] See: Heller, "Hobartian Voluntarism."

that desire, then they have control over that action, even if they could not have done otherwise.

According to Heller, beliefs fit neatly into a higher-order desire picture. There are various ways in which one forms beliefs. One may have a disposition to believe based on one's perceptions, for example. One may or may not want to believe in this way. For example, one might recognize that one's tendency to believe that one's lottery ticket will be the winning ticket is a faulty belief-forming process and want to change it. In the case in which one's belief forming mechanisms cohere with how one wants to form beliefs, one is in control of one's beliefs. In the case in which they do not, one is not in control.

The problem for higher-order doxastic compatibilism is that it does not seem to capture control correctly. Given the logic offered here, we could argue that digestion is under an agent's compatibilist control. A human's digestive system has different methods of digesting food. Amy sincerely wants to have these kinds of digestive processes. However, that does not show that Amy is in control over her digestive processes. As Peels (2017) points out, we can give a similarly unintuitive argument for beliefs. If Jane has the beliefs she does because she lives in an isolated community and endorses this kind of isolationism, her belief-forming behavior may be in accord with her higher-order desires, but that does not suggest that her belief-formation is under her control.

These two examples pick out a problem for a higher-order doxastic compatibilism. Part of the intuitive appeal for higher-order compatibilism for actions is that actions seem to ordinarily proceed from an agent's agency. When actions proceed from mental states that are not one's own—for example, one acts from an addiction that one does not want—

there are doubts that one's actions actually proceed from one's will. This is why higher-order compatibilists introduce second-order desires: to rule out cases in which one's actions proceed from sources other than one's will. The problem is, as we have already seen, one's belief forming processes do not ordinarily proceed from one's will, and thus they seem not to be one's own. They are more akin to digestion than action. But if they are not one's own, then the fact that they cohere with one's second-order desires is irrelevant. Thus, in order to make the higher-order desire account work, one needs to show that agents' beliefs flow from their wills. The last version of doxastic compatibilism purports to show just that.

The third, and perhaps most popular form of doxastic compatibilism is a *reason-responsive account* of compatibilist control: compatibilist control is a function of reason-responsiveness.[52] An agent has reason-responsive control over some behavior so long as that behavior is sensitive to reasons that that agent has.[53] As John Martin Fischer puts it, some behavior is reasons-responsive, if the reasons the agent has guide the action they (or he or she) takes. If I flip a light switch because I want to see the table, my action of flipping a light switch is guided by my reasons. Thus, reasons-responsive compatibilism would hold that I have compatibilist control over this action. However, I do not have compatibilist control over my digestion on this account. Regardless of whatever reasons I have, my body will digest food if I ingest it.

Normally, the reasons required for reasons-responsive compatibilism are practical reasons: the various considerations that count in favor of certain actions. For example, that

---

[52] See: Steup, "Belief Control and Intentionality," 2012, 2008; Ryan, "Doxastic Compatibilism and the Ethics of Belief"; Hieronymi, "Responsibility for Believing."
[53] See: Fischer, "Responsiveness and Moral Responsibility."

the ice cream is tasty is a practical reason that counts in favor of eating the ice cream. However, doxastic compatibilists hold that beliefs are sensitive to *intellectual reasons* rather than practical reasons.[54] Intellectual reasons are the reasons on which one forms or maintains one's beliefs. For Steup, this is one's evidence.[55] Hieronomyi holds that reasons for beliefs are the grounds upon which one settles for oneself whether some proposition is true, and which one offers when asked why one believes the way one does. These differ from those that bear on whether one should perform some action, on whether that actions is desirable: practical reasons.[56] If this more general understanding of reasons is the proper sense of 'reasons', most beliefs are reasons-responsive.[57] Believers tend to hold beliefs on the basis of what they take their evidence to be. For the rest of this paper, I shall take reasons-responsive doxastic compatibilism to be the strongest form of doxastic compatibilism.

Doxastic compatibilism fails to provide a sufficiently satisfying account of doxastic control. I shall show this with two different kinds of arguments. The first of these is that, contra the doxastic compatibilist, beliefs are not intentional. As a result, beliefs cannot be the object of deontic normative assessment. The second shall take a wider stance. Doxastic compatibilism may do a good job of accounting for epistemic obligations. However, it fails to be a compelling theory of the normativity of belief. It cannot, for example, account for moral obligations to believe. However, if doxastic compatibilism is incapable of

---

[54] See: Steup, "Doxastic Freedom." Steup uses the term "epistemic reason" rather than "intellectual reason," but it could turn out that there are practical epistemic reasons. For example, if the indirect influence account is correct, one can have an epistemic reason to read a book.

[55] See Hieronymi, "Responsibility for Believing," 358.

[56] Ibid.

[57] See: Steup, "Doxastic Freedom"; Steup, "Belief, Voluntariness and Intentionality."

accounting for any kind of obligation to believe other than epistemic obligations, then it will need to borrow from other theories to account for other kinds of obligations. As we shall see, this will prove that doxastic compatibilism should not be thought of as a rival to indirect influence theories, but, at best, an optional complement.

2.2. Feeling of Voluntariness

Let us begin with the first argument against doxastic compatibilism. In the literature, one of the main issues is whether doxastic states are intentional[58] In this paper, when I use the term "intentionality" I refer specifically to the intentionality manifested in voluntary actions.[59] Part of what separates an action from, say, digestion is that I can intentionally perform an action, whereas digestion is merely something that happens to me. If some mental state is unintentional, then one can have no obligations with respect to it. Intuitively, beliefs are unintentional. Regardless of whether I form an intention to or not, I cannot now change my belief that there is a table in front of me. Thus, the argument goes, even if one could show that beliefs are responsive to intellectual reasons, one could not show that beliefs are intentional.[60]

---

[58] See: Peels, "Against Doxastic Compatibilism.".

[59] See Searle, *Intentionality*. This use of intentionality is a species of more general uses of "intentionality" in philosophy of mind. Intentional mental states, in the latter usage, are those states that have representational content. Beliefs, desires, and intentions are intentional, in the latter sense, in that each content. Desires are not necessarily intentional in the former sense, because it is not clear that desires are voluntarily held.

[60] Again, it is possible for an agent to try to gather new evidence to change her belief. This would be an example of indirectly changing a belief.

However, this objection is not so cut and dried. For, as others have argued, there is a difference between an action being intentional and being performed on the basis of an intention.[61] Consider the following case:

> **BRAKE**: While driving one morning, a car swerves into your lane. You immediately brake to prevent from hitting the car.

Such an action certainly seems voluntary. However, it does not seem that you explicitly adopted an intention to brake and then acted on that intention. Doing so would seem to be one thought too many. The car is coming at you too fast for you to waste time debating whether to adopt an intention to brake. Instead, you immediately brake. The action of braking in this instance would be an example of an intentional action. John Searle calls these actions examples of intention in action without a prior intention.[62] Michael Bratman classifies these as examples of present-directed intentions.[63] According to Michael Bratman and Richard Holton, we label these actions as intentional due to the feeling of voluntariness felt in performing the action.[64]

If actions performed from present-direct intentions are intentional, then they can properly be said to be voluntarily performed by an agent. In BRAKE, Steup argues, the agent *could* perform differently than she does by adopting a different present-directed intention. Steup argues that a similar argument can be given for beliefs.[65] It may be impossible to adopt a belief based on a prior intention. However, beliefs may be instances

---

[61] See: Mele and Sverdlik, "Intention, Intentional Action, and Moral Responsibility."
[62] See: Searle, *Intentionality*.
[63] See: Bratman, *Intention, Plans, and Practical Reason*.
[64] See: Bratman; Holton, *Willing, Wanting, Waiting.*.
[65] Steup, "Belief Control and Intentionality," 2012, 155–56.

of present-directed intentionality. When one adopts a belief, one does so from a present-directed intention to believe a certain way. If so, then beliefs are voluntarily held, and the proper subject of normative assessment. Rik Peels, elsewhere, has argued that in order for some choice to be performed from a present-directed intention it must be capable, in principle, of being performed from a prior intention.[66] There are problems, I think, with this argument, but I do not want to assess it here. Instead, I shall offer two additional problems for the intentionality of belief that should cause us to doubt that beliefs are intentional.

First, it does not follow from the fact that some actions are performed based on present-directed intentions that every instantaneous action is performed based on them. Part of the thing about present-directed intentions is that they seem to require deliberation. Holton suggests that present-directed intentions are useful for overcoming certain psychological problems like indifference between two choices. Suppose, for example, that one needs to pick between two kinds of toothpaste to brush with, and one is indifferent between the two. An intention gets around the problem of indifference: one should just pick one so that one actually brushes one's teeth. This present-directed intention is the result of deliberation about one's reasons for performing an action. Present-directed intentions may also be manifestations of prior intentions. For example, I may form an intention to chop down a tree, and, when I go to do so, my intention becomes a present-directed intention. In this case, there is still deliberation in forming the intention. It is

---

[66] See: Peels, *Responsible Belief.*

merely that the deliberation occurred in forming the prior intention that turned into a present-directed intention.

The problem for BRAKE is that there need not be any deliberation involved in the action. There doesn't seem to be appropriate amount of time to respond to the reasons that one has. One cannot deliberate about one's reasons. Indeed, deliberation would be an example of one thought too many. By the time one deliberated on what the best course of action would be, one would hit the car one would be trying to avoid. Perhaps, one might have a standing intention to brake when a car swerves into one's lane. However, even though this is plausible, it is also plausible that one has not adopted a standing intention.

It is plausible, I argue, that pressing the brake could be an instance of an unintentional manifestation of a disposition. Compare the phenomenological experiences of first learning to drive with the ones of an experienced driver. When one first learns to drive, every decision is intentional. One wants to brake, and so one taps the brakes. One must be consciously aware of how much pressure one places on the brake: too much and it will be too abrupt, and too little and one will not brake enough. Of course, there are still conscious decisions for the experienced driver, but they are rarer than for the beginner. Unintentional actions are done all the time. For example, whenever I see a cop car while driving, I depress the brakes even if I am going the speed limit. I would not do such a thing if I were deliberating, but I have been habituated to the point where, if I see a cop car, I slow down.

However, if automatic behavior, as exhibited in BRAKE, is not an example of intentional action, as I argue, then one requires more than an analogy to this kind of

behavior to prove the intentionality of beliefs. Both are automatic seemingly unconscious behaviors, but this is evidence that the driving example is not an example of intentionality, not that beliefs are an example of intentionality. One needs more evidence to show that beliefs are intentional than that they are like these cases.

If the previous argument works, then Steup's argument for the intentionality of belief is unsound. However, I want to briefly consider a few objections to the above argument. First, one might think that my preceding argument must be unsound, simply because it proves too much. For, it may seem to imply that an agent does no wrong when her actions are manifestations of unconscious dispositions. As I have said, this would mean that her action was unintentional. Thus, instances that we normally describe as intentional, and, thus, are capable of being normatively assessed—tapping the brakes when a car swerves into one's lane, for example—would not be capable of being assessed. This would require a radical revision of which actions we normatively assess. The idea here, is that this revision must be wrong, and thus we have reason to doubt the soundness of my argument.

However, the aforementioned objection would simply read much into the argument. If one does something seemingly wrong as an unconscious behavior, she still may have violated an obligation in inculcating her unconscious behavior. Suppose that John has the habit of speeding when he drives. Such a habit is a result of how he taught himself how to drive. If the speeding seems wrong, we can point to the wrongness of how John inculcated behavior to explain its seeming wrongness. Of course, such an argument would not help the doxastic compatibilist. After all, this is just what the indirect influence position holds. One does not do wrong in holding some belief. Instead, one does wrong in inculcating the belief-forming behavior one has.

Second, one might argue that the preceding account gets it wrong because, as I mentioned above, the driver has a standing intention to brake whenever a car swerves into the driver's lane. Even if this were true—and again I'm not convinced—this would not help the doxastic compatibilist to argue that, in the case of BRAKE, one's present-directed intention is the result of some prior intention a driver has to brake when a car swerves into their lane. This is because, as we have noted, one cannot adopt belief based on prior intentions. Thus, even if the standing intention story was correct, it would do nothing to help the doxastic compatibilist. Since, all it would show is that braking could be the manifestation of a prior intention, not that beliefs can be.

Another objection to my first argument is that, while some instances of unconscious behavior are clearly unintentional, that does not show that BRAKE is unintentional. For, there is a feature of intentional action that picks BRAKE out as intentional: the feeling of voluntariness. Searle refers to the feeling of voluntariness of intentional action as the "experience of acting"[67] What picks out intentional actions from unintentional events or states of affairs, for Searle, is that the experience of acting brings about the event or state of affairs.[68] Holton has argued that it is the feeling of voluntariness that makes persons believe that their actions are free.[69] Thus, perhaps a defender of the intentionality of these kinds of actions can point to the feeling of voluntariness to save cases like BRAKE. This

---

[67] See: Searle, *Intentionality*, 87.
[68] For Searle, an action is intentional even if one does not successfully perform it. If someone tries, but fails, to raise their arms their action is intentional even if the intention is not satisfied. This is analogous to the case of the person who perceives some object that he is hallucinating. A perception is satisfied if one's perception accurately represents reality. The perception of the hallucination is intentional (in the philosophy of mind sense) even if it is not satisfied. See: Searle, *Intentionality.*
[69] See: Holton, *Willing, Wanting, Waiting*.

objection might hold that, given the circumstances, the feeling of voluntariness is sufficient to show that braking in BRAKE is intentional.

I find the argument from the feeling of voluntariness to be somewhat plausible and do not want to object to it here. Of course, we should point out that Searle and Holton have seemingly only given a necessary condition for intentionality, not a sufficient condition. Thus, it is not enough to show that braking feels voluntary to show that braking is intentional. Of course, opponents might argue that the circumstances of BRAKE along with the feeling of voluntariness are sufficient to show that braking is intentional. I do not want to belabor the point here.

However, recognition of the feeling of voluntariness as a necessary part of intentional action leads to my second argument against the intentionality of belief. The problem for doxastic compatibilist is that the feeling of voluntariness that accompanies actions is non-existent for beliefs. With actions, when one depresses the brake in an alleged instance of intentionality in action, one has a feeling of voluntariness.[70] Beliefs, however, do not possess this feeling of voluntariness. The argument that beliefs are involuntary is based on introspective analysis of one's beliefs. One considers some belief and determines that one cannot change one's beliefs at will. One also considers one's beliefs and realizes that the evidence that one appears to have determines which belief one forms. One's beliefs *feel* involuntary. If beliefs feel involuntary, then they lack this essential feature of intentionality. If so, then the doxastic compatibilist argument for obligations to believe falls

---

[70] See: Searle, *Intentionality*.

through. We have no reason to believe that beliefs are intentional in the sense required for obligations to believe.

Some have argued that there is a feeling of voluntariness in belief. Sharon Ryan (2008) has offered some examples to try to show the intentionality of beliefs. One example she gives is based on her experiences of 9/11:

> **9/11**: For example, on the morning of September 1 1, 2001, I was listening to NPR and I believed that a plane accidentally crashed into one of the World Trade Towers. After the second plane hit the second tower and the Pentagon was hit, I decided I that was mistaken. I decided that I had better face reality and disbelieve that it was an accident and that's what immediately happened.[71]

The feeling of deciding to believe that September 11[th] was not an accident is indeed a feeling of voluntariness. However, it is not clear that this is as voluntary as it seems. Ryan's evidence changed from before the second plane crashed to after. Her apparent evidence has determined that she disbelieve the plane crashes were accidents. Mere reflection on this evidence would necessarily produce a change in belief. From whence comes the feeling of voluntariness? Potentially, the voluntariness stems from her decision to open her belief up to evaluation. One can imagine examples in which one obstinately refuses to evaluate one's belief because one thinks that belief reevaluation would cause one to reject it. The freedom, here, is located in the fact that one could open up one's belief to reevaluation, if one wanted to do so. However, the argument from doxastic involuntarism does not hold that belief evaluation or reflection upon one's evidence are involuntary. Indeed, these processes seem intentional; a normal believer seems to have the power to decide to reflect upon her

---

[71] See Ryan, "Doxastic Compatibilism and the Ethics of Belief," 65..

evidence and current beliefs, or not to. Indeed, we can explain how this example actually fits with doxastic involuntarism by pointing out that, if one changes one's belief by recollecting one's evidence or opening up one's belief to evaluation, one has changed one's belief through intermediary mental states: evaluation and reflection. One still cannot change one's beliefs at will.

Why should we accept my story of belief change over Ryan's evaluation of having decided to change her belief? Part of my skepticism for Ryan's position is that this seeming example of voluntariness seems to be incongruent with virtually identical situations. For example, once Ryan believes that the 9/11 attack was not an accident, she cannot change her belief back to that it was an accident without further reflection or new evidence. Suppose she attempted, with no further deliberation and no new evidence, to change her mind back to 9/11 being an accident. This would be impossible. Ryan even cedes this: "It is true that I can't just decide, independently of how the evidence looks, to believe or disbelieve or suspend judgment on a proposition. My doxastic decisions are guided by what seems to me to be good evidence (or whatever it is that motivates me)."[72] If Ryan cannot change her belief except by changing her evidence, or what she takes her evidence to be, there is no reason to assume that Ryan's initial belief formation was voluntary. Given the apparent veracity of doxastic involuntarism, understanding Ryan's account as one of belief change through an intermediary is more plausible.

Another alleged example of apparent believing at will is offered by Peels.[73] Suppose that Peter will be offered a sum of money, if he believes that he will be offered

[72] Ibid.
[73] See Peels, "Believing at Will Is Possible.".

that sum of money. Further, Peter is aware of this arrangement. Peter, as a result, believes that he will be offered that sum of money. This seems to be an example of voluntary belief formation. Peter's belief obtaining makes the proposition believed true which is the opposite of normal beliefs. I am unconvinced that this is an instance of direct belief change, however. Surely, what happens here is that Peter reflects upon his evidence that if he believes then he will gain that sum of money. Upon this reflection, he may change his beliefs. However, this is still an instance of change in belief through a mediating mental state—reflection—and not direct belief change. Likewise, such a situation—one in which one's beliefs affect the truth of the proposition believed—would be incredibly rare, and not a suitable defense of doxastic compatibilism. If doxastic compatibilism was restricted to scenarios in which one's belief obtaining was a truthmaker for the content of one's belief, there would be very little that one had obligations to do. Thus, I do not think it is a very challenging objection.[74]

From the considerations of this section, we have good reason to think that Steup's argument for the intentionality of belief is unsound: namely, he has not proven that the kinds of actions that he thinks are similar to beliefs are actually intentional. Further, we have good reason to think that beliefs are not intentional. A necessary condition of intentionality is the feeling of voluntariness. Since beliefs lack this, beliefs are not intentional. Since beliefs are not intentional, there is no reason to think that believers have obligations to believe. In the next section, I will offer a different kind of argument against the theory of doxastic compatibilism.

---

[74] It is worth noting that Peels denies the truth of doxastic compatibilism. See: Peels, "Against Doxastic Compatibilism."

2.3. The Other Standpoints of Normativity

If the argument from the previous section is sound, then the core thesis of doxastic compatibilism, that we could have direct obligations to believe or refrain from believing certain propositions, is false. However, some may not be convinced by these considerations. That being the case, I want to shift my argumentative tactic against doxastic compatibilism for the rest of the paper. Instead of arguing that doxastic compatibilism is false, I will argue that doxastic compatibilism fails to meet an important desideratum for a theory of the normativity of belief. In order to meet these desiderata, I argue, doxastic compatibilists must countenance certain features of indirect influence theories. However, if doxastic compatibilism requires indirect influence theory in order to provide a complete theory of the normativity of belief, then doxastic compatibilism and indirect influence theories are not in tension. If not, as I shall argue, then proponents of indirect influence theories need not worry about doxastic compatibilism.

Before we begin looking at this argument, I want to point out that there is one way in which doxastic compatibilism could borrow features of indirect influence theory and still be in tension with it. If doxastic compatibilism entails that indirect influence theory is false, then even if it borrows certain features of indirect influence theory, it is still in tension with indirect influence theory. However, we should not worry that this is the case. At its core, doxastic compatibilism is a thesis that we can directly normatively assess beliefs. Indirect influence theory holds that we can indirectly normatively assess beliefs by means of intermediary actions. While proponents of these two theses are often at odds, there is no

reason they cannot both be true at the same time. Given that they are not, at their core, contradictory theories, I shall take the soundness of my argument to show that indirect influence theorists need not worry about doxastic compatibilism.

One desirable feature of a theory of the normativity of belief is that it be able to capture all the ways in which beliefs seem normatively assessable. This is, after all, what explains commonsense intuitions that not only are believers epistemically assessable for their beliefs, but they are also ethically assessable. For example, one way in which beliefs can be normatively assessed—indeed, the most obvious way—is epistemically. We can assess whether some belief is epistemically permissible by determining whether a believer abided by epistemic standards, either directly if doxastic compatibilism is right, or indirectly if indirect influence theory is right, in its formation. Here, doxastic compatibilism shines brightest. Steup, for example, thinks that a belief is epistemically permissible, if the believer fits that belief to her evidence.[75]

However, there are other standards by which we can normatively assess beliefs. For example, we can morally assess beliefs: it might be morally wrong to have certain beliefs about other persons. We can also prudentially assess beliefs: it is morally wrong to have beliefs, if it goes against one's self-interest. The desire to account for assessments of this kind is not ad hoc. In everyday life, persons make assessments of this kind. There seems to be a sense that not only are racist beliefs epistemically wrong—beliefs about the superiority

---

[75] See: Steup, "Doxastic Voluntarism and Epistemic Deontology."

of one race over another are not evidentially supported—but they are morally wrong as well. One has moral reason not to have racist beliefs.[76]

Can doxastic compatibilists show that racist beliefs are immoral? One might think that because the doxastic compatibilist has shown that agents have compatibilist control over their beliefs, compatibilism can account for every different kind of normative assessment for belief. This is, however, not the case. The problem rests in the argument that compatibilists give for the compatibilist control over beliefs. According to compatibilists, like Steup, beliefs seem to not be under an agent's control because they are not responsive to practical reasons but to theoretical reasons: reasons that concern the truth or falsity over some proposition. The lack of responsiveness to practical reasons is problematic for the doxastic compatibilist, here. For, I shall argue for a fairly plausible bridging principle that if beliefs are not sensitive to practical reasons, then there are no practical reasons for belief. A common principle for reasons among philosophers is that in order for something to be a reason for an agent, that reason must be capable, at least in principle, of motivating the agent.[77] That is, an agent must be capable of forming an intention from a reason in order for it to count as a reason. There are many complexities that this interpretation will run into. We might wonder, for example, whether the smoker has a reason to quit smoking (it is unhealthy for them), even if it is psychologically impossible for them to do so. I want to avoid these complications and operate under a very weak version of this principle:

---

[76] I use moral assessment here because it is the most poignant, yet we could run the following argument against doxastic compatibilism for virtually any normative standpoint.
[77] See, for example: Dancy, *Practical Reality*; Markovits, "Why Be an Internalist About Reasons?".

> *Motivation*: Φ is a reason for A to Ψ only if it is psychologically possible for some rational agent to Ψ having been motivated by Φ. [78]

Some agent somewhere must be capable of being motivated by a reason, in order for it to count as a reason. This is precisely why there can be no moral reasons to believe. According to Steup and others, beliefs are not sensitive to practical reasons. Try as hard as one might like, one cannot believe otherwise based on moral considerations. No agent with psychological capacities similar to our own could do so. But, if not, then moral reasons fail to be reasons to change one's beliefs.

By now, the argumentative structure should be apparent. Saying something is morally wrong, amounts to saying that one has conclusive moral reason not to do it. However, as shown above, there are no moral reasons to change one's beliefs *directly*, since beliefs are not sensitive to moral reasons. Thus, standard theories of doxastic compatibilism fail to account for the moral wrongness/rightness of certain beliefs. As stated above, this same argument could be run for virtually any normative standpoint besides the epistemic standpoint. Doxastic compatibilism, then, fails to meet this intuitive desideratum.

Below, I shall show how adopting features of an indirect influence theory can save doxastic compatibilism from this objection. However, before doing so, I want to briefly examine a way a doxastic compatibilist might resist this argument. That beliefs are only responsive to theoretical reasons is not a necessary premise in the argument for doxastic compatibilism. Indeed, *pragmatists* hold that beliefs are sensitive to practical reasons as

---

[78] I do not want to wade into the issue of whether reasons by themselves must motivate, or it is a combination of reasons and desires.

well as theoretical reasons. One such pragmatist, Stephanie Leary argues that practical

reasons can motivate by changing the threshold of evidence that one must reach in order to

believe some proposition.[79] She offers the following example:

> **MARY**: Mary is usually sceptical of other people's testimony and
> arguments. She didn't give religion much thought growing up, but at college
> Mary comes to believe that she would be happier if she were to believe that
> God exists, and that this is a strong reason to do so. While taking a
> philosophy class, she reads Aquinas's and Anselm's arguments for the
> existence of God and she befriends a student who tells her about his
> experiences of divine revelation. While this would usually not be enough to
> convince Mary (given her sceptical nature), because she recognizes the
> practical benefit of believing in God as a reason to do so, this causes her to
> be more swayed by those arguments, and she ends up believing that God
> exists.[80]

If this is what pragmatists have in mind for being swayed by practical reasons, it is not

problematic for the argument offered here. One way of interpreting MARY is that Mary

has practical reasons to attempt to adjust her belief-forming dispositions. Indeed, this is

how Leary interprets this example.[81] Mary does this by reading philosophical arguments in

support of God's existence, befriending those who believe in God and genuinely attempting

to believe. The argument against practical obligations to believe is that the recognition that

some belief is, for example, morally wrong cannot directly and immediately affect what

one believes. One may be able to take actions to change one's dispositions, but one cannot

---

[79] See: Leary, "In Defense of Practical Reasons for Belief."
[80] See: Leary, 538.
[81] See: Leary, 539.

immediately change one's beliefs based on purely practical reasons. Since, in addition, no doxastic compatibilist relies upon pragmatism for their argument, we shall, then, set aside the pragmatist challenge going forward.

As stated before, indirect influence theory accounts for this challenge well. We can account for the seeming moral wrongness of holding certain beliefs in the same way a proponent of indirect influence theory would account for epistemic wrongness. Certain actions are causally upstream of beliefs. Unlike with beliefs, actions are sensitive to practical reasons. If an agent recognizes that certain actions might lead to morally bad beliefs, that might provide her with a moral reason not to perform that action. This is, of course, indirect rather than direct assessment of the belief. However, if direct moral assessment of beliefs is impossible, and indirect influence theory can do a reasonable job of accounting for our intuitions—for example, why holding racist beliefs seems morally wrong—then that is preferable to being completely unable to explain the seeming moral wrongness of beliefs.

As stated before, I do not think that there is anything antithetical to doxastic compatibilism to adopting features of the indirect influence theory. The best way for a doxastic compatibilist to account for the moral wrongness of certain beliefs would be to adopt the strategy described in the previous paragraph. However, if one accepts that beliefs can be indirectly assessed by means of the actions that lead to them, then one accepts some version of indirect influence theory. The major source of contention is whether the doxastic

compatibilist would agree that epistemic assessment can also be assessed indirectly, or if they would dig in their heels and think it can only be assessed directly. [82]

However, I see no reason to think that epistemic assessment cannot be both direct and synchronic, if the doxastic compatibilist is correct, and indirect and diachronic, like the proponent of indirect influence theory argues. Take, for example, an agent who forms her beliefs based on her theoretical reasons but has engaged in bad evidence gathering practices. On the one hand, there is something right about believing based on one's theoretical reasons. When one considers one's beliefs, anything else but following what one has most reason to do seems irrational or wrong—remember, we are assuming for the purposes of this argument that doxastic compatibilism is correct. Nevertheless, it seems that there is something epistemically wrong with this agent's evidence-gathering practices. She should act better here. What explains this is that while from a synchronic perspective, the agent has believed correctly, from a diachronic perspective, she has not. There is nothing odd or incoherent about this two-tiered approach. And, if the doxastic compatibilist is right about the normativity of belief, then the indirect influence theorist should happily embrace it.

Might doxastic compatibilists object to this line of argument? It is not clear. At the very least, some compatibilists, Sharon Ryan for example, think that there are both synchronic and diachronic epistemic obligations and that both are incredibly important.[83] Of course, there is a difference between the proposition that doxastic compatibilism

---

[82] Although, I am not aware of any compatibilist who necessarily rejects indirect assessment as a part of epistemic normativity, very few actually accept it. Sharon Ryan, as far as I am aware, is the only one who does. Likewise, there is a prevailing opinion that there is no such thing as indirect epistemic normative assessment. See: Dougherty, "Reducing Responsibility," December 2012; Feldman, "The Ethics of Belief."
[83] See Ryan, "Doxastic Compatibilism and the Ethics of Belief," 47..

requires diachronic obligations, and the proposition that, doxastic compatibilism is no threat to indirect influence theory. A compatibilist might object that indirect influence theory, by failing to allow for synchronic obligations, fails to capture an essential aspect of epistemic normativity. There may be some times where indirect influence theory would hold that an agent held some belief permissibly, where doxastic compatibilism would hold that that agent held it impermissibly. For example, suppose that an agent abided by all of their diachronic obligations in gathering evidence and inculcating good reasoning methods. Nevertheless, as a matter of fact, that agent violates her compatibilist synchronic obligations in forming the belief. Here, an indirect influence theorist would say that the agent has done nothing wrong, while the compatibilist would say that the agent has done something wrong.

If doxastic compatibilism is true, an indirect influence theorist would be remiss in not allowing synchronic obligations of belief. However, that fails to bear on the thesis of this paper. Regardless of the truth of doxastic compatibilism, we have good reason to endorse an indirect influence account of obligations of belief. The same is not true of synchronic compatibilist obligations. We only have reason to endorse these synchronic obligations if doxastic compatibilism is true. If not, then there is no reason. Thus, I argue, the truth or falsity of doxastic compatibilism need be of no concern to indirect influence theory.

2.4. Conclusion

Where do these considerations leave doxastic compatibilism? If my first argument is sound, then doxastic compatibilism cannot account for obligations of belief, since beliefs are not intentional. If my second argument is sound, then doxastic compatibilism requires features of indirect influence accounts in order to account for all features of the normativity of belief. However, if this approach is taken, then doxastic compatibilism, understood this way, should not be viewed as a competitor to accounts of epistemically permissible actions, but as a complement. Of course, if doxastic compatibilism is true, any theory of indirect influence that does not permit doxastic compatibilism is incomplete. Nevertheless, doxastic compatibilism is no threat to indirect influence accounts of epistemic normativity.

Chapter 3 - Epistemic Normativity Matters: At Least to Some

Trent Dougherty (2011) argues that there is no distinctive standpoint of epistemic normativity—no assessments of epistemic obligations, permissions, praiseworthiness, blameworthiness—since any instance of epistemic normative assessment is, in fact, an instance of moral or prudential assessment. In this paper, I agree with Dougherty that whether an agent abides by an epistemic norm will coincide with whether that agent abides by a prudential or moral norm. However, I disagree that this eliminates the need for epistemic normativity.

3.1. Epistemic Normativity

This paper is concerned with an account of epistemic normativity. Normativity, as I shall use it here. Normativity is that branch of philosophy concerned with questions of whether some event or state of affairs is right or wrong, or good or bad. Thus, for example, morality makes deontic moral assessments about the moral rightness or wrongness and axiological moral assessments about the goodness or badness of events and states of affairs. There are many different normative *standpoints* by which one can make normative assessments. The term "standpoint" is one that I borrow from Dale Dorsey.[84] The term denotes a class of functions that take facts about a given choice scenario as inputs and output these normative assessments about that choice scenario: for example, that a given action is permissible or that a given agent is blameworthy. The prudential standpoint, for example, would take facts

---

[84] Dorsey, *The Limits of Moral Authority*, 9.

about a choice scenario as inputs and output prudential assessments: for example, taking a

high-stakes bet with a low probability of winning is prudentially impermissible; pleasure

is prudentially good.

Normative standpoints provide a useful architecture for discussing normativity.

Standpoints capture essential elements of the normative landscape one finds oneself in in

any choice scenario. Yet, not every normative standpoint is equally important. For

example, the rules of etiquette might tell me that at a formal dinner I should always eat

starting with the outside-most utensil and work my way inward. However, the norms of

etiquette may simply not *matter* to me. By "mattering," I mean that some mental state's

epistemic permissibility or impermissibility provides a *practical reason*, by itself, for or

against possessing that mental state. [85] My usage of the term "practical reason" here closely

resembles Sharon Stroud's term "reason for action" which she calls "a consideration

relevant to the generic practical question of what to do"[86] We can think of this as the most

general answer to the question of what one ought to do in some choice scenario. The term

"generic practical question of what to do" is slightly cumbersome, so we shall use the

shorthand "practical normativity" to capture it.

It follows from this discussion that if some normative standpoint does not matter,

then it is not important for practical reasoning; an agent has no need to include it in his

normative considerations.[87] Thus, for any potential normative standpoint, to determine

---

[85] Dorsey, 12–13.
[86] Stroud, "Moral Overridingness and Moral Theory," 172.
[87] I take no stand here on whether "mattering" is entirely contingent on the subjective choices of an agent:
that any normative standpoint, even morality, might not matter to an agent. Some versions of normativity
might require that agents have practical reason to abide by moral norms even if they don't think
normativity is important. Others might deny this. My analysis does not depend on answering this question.

whether that standpoint should be part of my practical reasoning, we must answer two distinct questions. First, does the standpoint exist? Second, does the standpoint matter?

In this paper, we shall be concerned with whether the alleged standpoint of epistemic normativity is a part of an agent's practical reasoning. This means determining the answer to the two questions above. To say that there is a distinctive standpoint of epistemic normativity is to hold that there are epistemic grounds upon which one can make normative assessments. That is, one can apply deontic or axiological assessments based on epistemic grounds to an agent's mental states or actions. I have in mind such states as an agent's occurrent mental states such as beliefs, desires, and intentions that she intentionally possesses that she forms or maintains voluntarily. In this paper, we are specifically concerned with the deontic component of epistemic normativity: assessments of the epistemic rightness, wrongness, permissibility, or blameworthiness of an agent's mental states or actions, and shall take "epistemic normativity" to be synonymous with "deontic epistemic normativity." To say, then, that epistemic normativity matters is to say that that some behavior is epistemically obligatory/permissible/impermissible counts as a practical reason for or against it. To give an intuitive example, that going to college is epistemically permissible for someone counts as a practical reason in favor of their going to college.

In the recent years, there has been a growing interest in the standpoint of epistemic normativity as a field of study.[88] However, some have challenged that there could be such a standpoint. In the next section, we shall turn to an objection that there can be no distinctive standpoint of epistemic normativity.

---

[88] See, for example: Booth and Peels, "Why Responsible Belief Is Blameless Belief"; Nottelmann, *Blameworthy Belief*, 2007; Peels, *Responsible Belief*.

3.2. A Dilemma for Epistemic Normativity

Contra those who defend it, Trent Dougherty argues that there is no distinctive standpoint of epistemic normativity.[89] Dougherty's argument is eliminitivist in nature: epistemic normativity, Dougherty claims, is reducible to morality and prudence. Therefore, it ought to be eliminated.[90] When Dougherty assesses epistemic normativity, he is specifically concerned with theories of epistemic normativity that indirectly assesses the normativity of beliefs by means of the actions that causally influence those beliefs—diachronic epistemic normativity.[91] For example, theories that hold that it is epistemically wrong to take an action that fails to maximize one's justified beliefs.[92] He puts aside those theories of epistemic normativity that directly normatively assess beliefs—synchronic epistemic normativity. Although he does not mention why, my guess is that he does so because he follows William Alston in holding that synchronic epistemic normativity isn't really normativity at all.[93] One cannot directly assess the beliefs of an agent, because agents lack control over their beliefs and Ought Implies Can.[94]

---

[89] Dougherty, "Reducing Responsibility," December 1, 2012.
[90] It is difficult to categorize Dougherty's project as reductionist or eliminitivist. Dougherty labels it as reductionist. However, one can reduce one concept into another and still think that it is useful to talk about the reduced concept. For example, even if biology is reducible to chemistry, it might still be useful to talk about biology. Dougherty wants to dispense with epistemic normativity altogether. Therefore, I side with eliminitivist. Dougherty, 537.
[91] See: Dougherty, 538.
[92] For a more detailed account of what diachronic theories might look like, see: Chapter 1.
[93] See: Alston, "The Deontological Conception of Epistemic Justification."
[94] For a more in-depth discussion of this point, see Chapter 1.

On diachronic epistemic normativity, a believer's beliefs are not directly normatively assessable, since they are not under her control. However, her beliefs are indirectly assessable, if they are the product of voluntarily causally upstream actions that she is agentially responsible for. Such an account of epistemic normativity does not face the issues of involuntary belief that synchronic normativity does. This is because, while beliefs may not be under an agent's control, actions normally are. Since actions are under an agent's control, they can be normatively assessed. However, Dougherty argues, diachronic epistemic normativity cannot account for epistemic normativity. Following Richard Feldman, he argues that any alleged instance of diachronic epistemic permissibility is really an instance of moral or prudential permissibility.[95] Call this thesis "Identity"

> **Identity**: Each instance of epistemic impermissibility is just an instance of purely non-epistemic impermissibility (either moral or prudential).[96]

This is not the original formulation of Dougherty's thesis. Dougherty's original thesis is that any alleged instance of epistemic *irresponsibility* is just an instance of non-epistemic irresponsibility. The problem, here, is that it is not clear what "irresponsibility" is supposed to mean. The term may come from epistemic responsibilists who think that epistemic justification depends upon behaving epistemically responsibly.[97] One behaves epistemically responsibly when one manifests certain positive character traits. However,

---

[95] See: Feldman, "The Ethics of Belief," 689; Dougherty, "Reducing Responsibility," December 1, 2012.
[96] The use of the term "prudential" is mine. Dougherty uses the term "instrumental irrationality" but this term seems ambiguous between prudential normativity and practical normativity. Given what he says elsewhere, I take prudential normativity to be what he means by "instrumental rationality." See: Dougherty, "Reducing Responsibility," December 1, 2012, 537; Dougherty, "Internalist Evidentialism and Epistemic Virtue," 286–87.
[97] See, for example: Axtell, "From Internalist Evidentialism to Virtue Responsibilism."

one need not be a responsibilist to think that epistemic normativity is important. I, thus, prefer the use of the terms 'permissible' and 'impermissible' to keep the argument clearer.

One need not worry that we have inappropriately altered Dougherty's argument, for Dougherty himself thinks his argument applies to more than epistemic irresponsibility. Consider this claim: "I want to validate Feldman's claim that *any normativity concerning belief that goes beyond fitting the evidence*, and in particular epistemic responsibility, is either moral or instrumental. [emphasis mine]".[98] Epistemic impermissibility is an assessment of epistemic normativity and, thus, challenged by Dougherty's argument. There is no reason not to change Dougherty's thesis in the way I have. Although, if one is really committed to Dougherty's version, any argument that I forward for the rest of this paper would work equally well as an argument against Dougherty's thesis.

Dougherty's argument for Identity is, at its core, an argument from parsimony.[99] Dougherty contends that if we can fully explain any alleged instance of epistemic normativity in terms of moral or prudential normativity, then for reasons of parsimony we need not posit the category of epistemic normativity in our ontology. Dougherty contends that we can reduce any alleged instance of epistemic normativity to moral or instrumental. To see his argument for this, let us consider some alleged instance of epistemic normativity: John has an epistemic obligation to read a book to improve his knowledge. According to one who endorses diachronic epistemic normativity, this obligation is determined entirely by epistemic factors. Dougherty, however, pushes this thought. He offers a trilemma for

---

[98] Dougherty, "Reducing Responsibility," December 1, 2012, 534.
[99] Dougherty outright states this. "Reductionist theories are to be preferred to non-reductionist theories due to their parsimony." See: Dougherty, 537.

any alleged instance of epistemic normativity. John's epistemic obligation to read a book gives him a practical reason to read a book only if one of the following three conditions is met. First, no one's interests are at stake. Second, only others' interests are at stake. Third, only one's own interests are at stake.

Each of the three potential conditions are problematic for a defender of epistemic normativity, however. Dougherty argues that if no one's interests are at stake, one has no practical reason to perform an action. This makes some intuitive sense. For example, imagine that I had an epistemic obligation to know whether the number of blades of grass in my lawn is even or odd, even though no one's interests are at stake. Here, it seems that my epistemic obligation confers no practical reason on me to count the number of blades of grass in my lawn. There are some potential solutions to this challenge for epistemic normativity.[100] However, for the purposes of this paper, I shall concede this point.

Suppose, instead, that John has this obligation to read this book in virtue of other-regarding interests. John should read the book because, for example, it will improve his literacy on vaccines. If he reads it, then he will vaccinate his children, preventing certain diseases from affecting those that his children come in contact with, a moral good. However, note that the defense appeals to the moral goodness of the consequences of reading the book to defend the obligation. John has a moral reason to improve his epistemic states, because of the positive moral downstream effects. John's obligation is really moral, not epistemic. We can posit just such an explanation for any instance in which one's alleged

---

[100] See: Chapter 1. There is another problem that Dougherty's claim leaves out aesthetic obligations, since one can have an aesthetic obligation even if no one's interests are at stake. For example, I have an obligation to save a Monet by securing it in a vault underground, even if no one will ever be able to see that Monet again ever. However, my discussion is not related to aesthetic obligations, so this challenge is outside the scope of this paper.

epistemic obligations depend upon other-regarding considerations. There is, then, no need to posit the existence of such a normative standpoint as an explanation, and our ontology would be better without it.

Unfortunately, by parity of reasoning, if the content of an obligation depends upon one's own interests, one's obligation is also not epistemic. Suppose that the reason John needs to read the book is because it is prudentially valuable for him to get the vaccine. This goal, we shall suppose for simplicity, will not affect anyone else.[101] If John fails to read the book, he has failed to meet his goal. This prudential failure explains the wrongness of not reading the book, and there is no need to appeal to epistemic reasons to explain it. Again, for reasons of parsimony, Dougherty argues that we should not posit the existence of an epistemic obligation when we already have a complete explanation from prudence.

Thus, Dougherty's argument for parsimony is complete. Any alleged diachronic epistemic obligation either affects someone's interests or it doesn't. If it doesn't, then there is no obligation. If it does, then it either affects one's own interests or it affects someone else's. If it affects one's own interests, it is really a prudential obligation not an epistemic obligation. If it affects someone else's, it is really a moral obligation not an epistemic obligation. Any alleged diachronic epistemic obligation is actually not an epistemic obligation. If epistemic normativity exists, then there is at least one epistemic obligation. Since there are no epistemic obligations, epistemic normativity does not exist.

---

[101] We could discharge this assumption without much problem. Reading the book might be both prudentially and morally obligated. If so, the alleged epistemic obligation can be understood both as a prudential and moral obligation. This is no threat to Dougherty's argument.

3.3. A Brief Comment on Normative Standpoints

By Dougherty's lights, if his argument goes through, then there is no distinct standpoint of epistemic normativity, and there is no such thing as normative epistemology. The spirit of this objection is powerful, and we shall deal with it carefully for the rest of the paper. However, I would like to address one point about it immediately. At least as characterized, Dougherty's objection is false. Dougherty posits that there is no standpoint of epistemic normativity. However, at least in the way we have been describing it, there must be a standpoint of epistemic normativity. After all, epistemic normativity is just that normative standpoint that assesses objects based upon epistemic criteria. Regardless of whether it is important to any agent, the fact is that it is perfectly cogent to assert a standpoint that assesses actions based on epistemic criteria, in the same way that one might assess actions based on certain club norms. The existence of a normative standpoint does not depend on whether it matters practically.

Indeed, the very objection from parsimony does not seem to apply to normative standpoints. There are all kinds of ways that we normatively assess actions, and many more ways that we could. We could assess a person's behavior based upon the norms of medieval chivalry. Clearly, such an assessment is not important now. That does not mean the assessment could not be made. We could also normatively assess agents in ad-hoc ways. For example, John's behavior may be impermissible on the grounds that it violates not-reading-books norms. It would be odd to deny that any of these are normative standpoints because we are able to cogently make normative assessments like this. Nor, necessarily, do these assessments need to be based on advancing one's own interests or someone else's.

Some of these, I take it, are completely unconnected to that. About the only way one could show that a certain normative standpoint did not exist would be to show that the grounds of assessment on which it is based are incoherent or non-extant.

Dougherty has not claimed that epistemic grounds of assessment do not exist. Indeed, Dougherty's trilemma relies upon explaining away the epistemic grounds. But Dougherty does think that changing one's beliefs so that one believes truly *can* be important, and, indeed, he takes this to be what epistemic assessment amounts to. His contention is, merely, that this concern for the truth only appears in instances in which one's own interests or another's are at stake. So, it would go against his argument to claim epistemic grounds of assessment are incoherent or non-extent. Thus, Dougherty cannot and should not attempt to argue that epistemic normativity does not exist.

However, as I said before, this does not show that Dougherty's objection is weak. It merely shows that it is misplaced. Rather than targeting the existence of epistemic normativity as a normative standpoint, Dougherty should target the claim that it matters for practical reason. Consider Dougherty's trilemma again. The main force of the trilemma is that the reason-giving force of epistemic obligations either doesn't exist, if no one's interests are at stake, or the force is completely captured by moral and prudential norms. If morality and prudence can entirely explain the reason-giving force of any alleged epistemic norm, then epistemic normativity is superfluous for explaining why that alleged epistemic norm is a practical reason. In other words, because we can explain John's practical reasons for reading a book entirely based on moral and prudential norms, there is no need to appeal to epistemic norms to show that he has a practical reason to do so. Thus, epistemic norms being unnecessary, we should abandon the thought that they matter for practical reason.

This updated version of Dougherty's objection has many virtues. First, the core intuition and thought experiment behind the objection remain the same. The seemingly irrelevant nature of epistemic norms is still the main motivation for the objection. We are merely positing a different explanans for this intuition and thought experiment. Second, we do not need to deny the existence of the standpoint of epistemic normativity to make this argument. People can still make judgments about whether an action is epistemically permissible, nor is there anything incoherent about doing so. Thus, all things considered, it is better not to do deny its existence. Third, the argument is no longer an argument from parsimony. The claim is that epistemic normativity is irrelevant to considerations of practical reason, and we need not countenance it. This argument is much stronger for any person who is suspicious of parsimony as an argumentative tool.

Nevertheless, as we shall see Dougherty's argument is still unsound. Before we get to that, however, I first want to deal with a few potential objections to Dougherty that are instructive.

3.4. The Additional Reason-Giving Force Objection

Whichever version of the Dougherty's argument one considers, there is one particular move that one might object to. Just because one can explain the seeming reason-giving force of epistemic normativity in terms of moral or prudential normativity, that does not mean that epistemic normativity has no reason-giving force. Epistemic normativity may have some reason-giving force in addition to the moral or prudential force. Thus, just

because we have explained some of the reason giving force in terms of moral and prudential norms does not mean that we have explained away all of the reason giving force. Another way of putting this, is that Dougherty has failed to explain away all of the reason-giving force of epistemic normativity. I take it that this is Guy Axtell's position when he holds:

> His [Dougherty's] points about how practical interests and values may inform issues of attention, etc. serve as a corrective to any veritists [normative epistemologists] who would claim 'a general, impersonal duty to seek truth as such.' But these points I see going nowhere in regards to showing that all choices of attention and all investigative strategies in pursuit of true belief, knowledge, and understanding, etc. are nothing but moral directives or else matters of prudential/instrumental rationality.[102]

For Axtell, the additional reason giving force is a matter of rational confidence. One should abide by epistemic norms because they increase one's confidence in one's belief's being true, increase the arsenal of reasons one's have to defend one's beliefs if challenged on them, and will provide the agent with rational grounds to abandon or maintain her beliefs in light of counter-evidence.[103]

There is clearly something to Axtell's objection that abiding by epistemic norms is an aid to one's rationality in the maintenance and updating of beliefs. Yet, I think that Dougherty's objection can and does come back again in the following way. Suppose that we accept that there is this matter of rational confidence that abiding by epistemic norms will affect. Why, then, should we abide by these epistemic norms? Humans have scarce resources. They cannot pursue all the ends that might be worthy of pursuit. It is simply not practically rational to pursue some item of knowledge, if it has no importance to oneself or

---

[102] Axtell, "From Internalist Evidentialism to Virtue Responsibilism," 447.
[103] Axtell, 447.

others. Imagine a person who counts the number of blades of grass in her yard simply because knowing that fact is an additional bit of knowledge. In light of any other countervailing consideration, this course of action seems practically impermissible. Why, then, would they pursue any bit of knowledge? Either because it is important for something else, or because they want to. The knowledge is valuable because of its instrumental role in pursuing things that are of value for oneself or others. But if we pursue this line of thought then we find ourselves right back in Dougherty's trilemma. If it is of value for oneself, then it is a matter of prudence. If for others, then it is a matter of morality.

Someone who defends this objection could, potentially, resist this move. What they might say is that the need for knowledge to be important is entirely a function of our scarce resources. This does not show that epistemic norms have no reason giving force by themselves, merely that their reason-giving force rarely makes a difference. After all, a perfectly rational being with unlimited time might have reason to collect all the knowledge that she could. Thus, perhaps epistemic norms matter, they just matter relatively little compared to other norms. This is an interesting line of thought. However, I shall posit an objection to Dougherty that is not capable of being targeted by the Doughertian trilemma. For now, however, let us move to another objection to Dougherty.

3.5. The Doxastic Compatibilist Objection

A large part of Dougherty's argument against epistemic normativity rests upon his setting aside synchronic epistemic normativity. Dougherty does not even think of this as truly normative. For Dougherty, synchronic epistemic permissibility just concerns whether the appropriate relation exists between one's evidence and one's beliefs. For Dougherty, to say that one has a synchronic epistemic obligation is to hold one to what Richard Feldman calls a "role obligation," an obligation one has regardless of one's ability to satisfy them.[104] According to Dougherty, role oughts that one is unable to satisfy do not matter for the purposes of practical normativity.[105] And, indeed, it is commonly held that since believers do not have the right kind of control over their beliefs, whenever they fail to meet their synchronic epistemic obligations, they are not able to meet them.[106] Thus, Dougherty holds, a proponent of epistemic normativity cannot appeal to synchronic epistemic obligations to save it.

Yet, it should be noted that this move to deny the normativity of synchronic epistemic obligations is not uncontroversial. The second objection one could raise against Dougherty's argument is to push back against this decoupling of normative epistemology from synchronic epistemic normativity. Some epistemologists, known as "doxastic compatibilists," hold that a believer is subject to an epistemic obligation if she is responsible for them in a compatibilist sense—take compatibilist responsibility to be responsibility that one has even if one could not have done otherwise.[107]

---

[104] Feldman, "The Ethics of Belief."
[105] See: Dougherty (2012: 537-538). Here he follows Richard Feldman and Hilary Kornblith. See: Feldman; Kornblith, "Epistemic Obligation and the Possibility of Internalism."
[106] For the most notable case, see: Alston, "The Deontological Conception of Epistemic Justification."
[107] See, for example: Steup, "Belief, Voluntariness and Intentionality"; Hieronymi, "Responsibility for Believing"; Heller, "Hobartian Voluntarism"; Ryan, "Doxastic Compatibilism and the Ethics of Belief."

One compatibilist, Matthias Steup, holds that agential-responsibility is a function of reasons-responsiveness: an agent has control over some behavior so long as that behavior is sensitive to reasons that that agent has.[108] Normally, the reasons, here, are cashed out as practical reasons. So, an agent does not have control over digesting food, because, no matter what reasons an agent has, that agent's body will always digest food, if there is food to digest. However, normally an agent does have control over flipping a light switch. That agent flips a light switch, for some reason, and if that agent had different reasons, she would not have flipped the light switch. Steup holds that beliefs are reasons-responsive but are sensitive to *intellectual reasons* rather than practical reasons.[109] Intellectual reasons are the reasons on which one bases one's beliefs. They are also the reasons that one could offer another, if pressed on why one believes the way they do.[110] As it turns out, then, most beliefs are reasons-responsive.[111] Believers tend to hold beliefs on the basis of what they take their evidence to be.

Another doxastic compatibilist, Mark Heller, holds that responsibility is a function of the coherence of one's higher order desires with one's will.[112] Will, in this sense, refers to the lower-order desires that one has that determine one's behavior. When one acts based on certain desires, and desires to act on those desires, then one has control over one's behavior. To give a simple example, the higher-order desire compatibilist would hold that a drug addict who wants to quit is not responsible for his drug usage. While his lower-order

---

[108] See: Fischer, "Responsiveness and Moral Responsibility."

[109] Steup uses the term "epistemic reason" to describe these reasons, but that would unnecessarily confuse them with the epistemic D-reasons I talk about epistemic normativity. Thus, the terminological change.

[110] See: Hieronymi, "Responsibility for Believing," 358.

[111] See Steup, "Doxastic Freedom"; Steup, "Belief, Voluntariness and Intentionality."

[112] "…a person is free to the extent that her highest-order desires determine her will, *her will being those of her lower-order desires which determine her actions*. [emphasis mine]"Heller, "Hobartian Voluntarism," 135.

desires may be to use drugs, he has a higher-order desire not to act upon those first order desires. The incongruity between orders of desires makes it so that he is not in control for drug-usage. Yet, if someone flips a light-switch because they want to, and they endorse that desire, then they are in control over that action, even if they could not have done otherwise.

Regardless of whether the doxastic compatibilist endorses the reasons-responsive or higher-order desire account of compatibilist responsibility, doxastic compatibilism presents a problem for Dougherty. For, according to doxastic compatibilism, if a believer has compatibilist control over her beliefs, then she is subject to doxastic obligations. This means, that synchronic epistemic normativity *is* actually a kind of normativity, and, except in rare circumstances, believers have an obligation to conform their beliefs to their evidence. This is problematic is because, unlike the case with diachronic epistemic normativity, we cannot simply appeal to some other branch of normativity to explain away the seeming normativity of a synchronic epistemic norm. For, failing to conform one's beliefs to one's evidence seems to be a failure of theoretical rationality. And the wrongness of failing to abide by the standards of theoretical rationality seems to need no other explanation. It just seems to be a bad thing not to abide by the standards of theoretical rationality. Any attempts to deny this, would entail a denying of the importance of theoretical rationality. On this account, epistemic normativity needs no appeal to morality or prudence to render it important.

Note, however, that, while there may be a normativity attached to synchronic epistemic normativity, it is not of the same kind as moral or prudential normativity. Moral normativity, for example, is a proper part of practical rationality. It purports to give an

answer to the question of what to do, from a moral perspective: it provides practical reasons for actions. The same can be said of prudential normativity. Notice, however, that synchronic epistemic normativity does not give these kinds of answers. It does not tell an agent how she should act. Instead, it tells an agent what she should believe. Further if she fails to abide by those norms, her failure is theoretical, not practical. For her failure is one of rational coherence between her beliefs and her evidence. However, if epistemic normativity derives its normative force from theoretical rationality, it is hard to see how it is supposed to matter from the practical perspective. How is the question of what an agent should believe, related to the question of how an agent should act?

The obvious answer is that one should inform the other. An agent should believe a certain way because it is valuable for practical rationality. After all, agents are better at achieving their practical ends if they have justified beliefs about what the facts of a scenario are. Thus, maybe abiding by the dictates of theoretical rationality is important, in addition to the fact that agents should increase their theoretical rationality because it aids in practical rationality. However, there is something interesting to note here. Not every belief matters for practical rationality. For most agents, a belief such as whether the number of stars in the universe is even or odd will not ever affect their lives. There are no stakes involved in it, and it cannot matter to practical rationality. Thus, only a subset of a believer's beliefs will be important for practical rationality. But what about the ones that do? The important question for these is why they are important for practical rationality. There are two possible answers, either they are important for oneself or they are important for another, or both. We have encountered Dougherty's trilemma again. If they are important for oneself/ they are entirely explainable by prudential norms. If they are important for another, they are

entirely explainable by moral norms. Either way, positing a practical dimension of synchronic epistemic normativity is unnecessary and we can do away with it. Until we show that Dougherty's trilemma is unsound, there is no way to show the existence of epistemic normativity.

A proponent of Identity should not be at all worried about this kind of objection. Regardless of whether compatibilism is true, one can concede the existence of synchronic epistemic normativity without giving up anything regarding the importance of epistemic normativity. If there is no practical dimension to epistemic normativity, epistemic normativity still does not extend beyond the coherence of evidence with belief. If one does try to connect theoretical rationality to practical rationality, then one will encounter another instance of Dougherty's trilemma. We have no other choice, then, but to object to Dougherty's trilemma directly.

3.6. A First Pass: Other Standpoints of Normativity

Before responding to Dougherty's argument, we must notice an ambiguity in it. As I noted before, there are different ways in which a standpoint matters for an agent. A particular normative standpoint might only matter for a particular agent, or particular group of agents. For example, agents might take on commitments to the norms of tennis, and thus tennis normativity matters for them. Importantly, tennis does not necessarily matter for every agent. Alternatively, it might matter universally. This means that that normative standpoint matters for everyone. Typically, many think that morality matters universally.

Dougherty's Identity thesis is ambiguous between these two claims. It is not clear if he is arguing that epistemic normativity does not matter for any agent or the weaker claim that epistemic normativity does not matter universally. If the former, then Dougherty's Identity should be construed as:

> **Strong Identity**: Necessarily, for every agent, each instance of epistemic impermissibility is just an instance of purely non-epistemic impermissibility (either moral or prudential).

I use the modal term "necessarily" because there is a conceptually possible world in which no one takes on a commitment to epistemic normativity. This would be similar to a world in which no one takes themselves to tennis norms. This would not show that it is impossible for agents to commit themselves to epistemic normativity, as it is merely a contingent feature of the world that no one takes on this commitment. If the latter, then Dougherty's thesis is really:

> **Weak Identity**: Possibly, there is an agent for whom each instance of epistemic impermissibility is just an instance of purely non-epistemic impermissibility (either moral or prudential).

I use the modal term "possibly" here because there is a conceptually possible world in which every agent takes on a commitment to epistemic normativity. This, however, would not show that epistemic normativity matters universally. In order for it to matter universally, it must matter for every agent in every conceptually possible world. As we shall see below, I am more interested in the falsity Strong Identity than Weak Identity, so I shall begin with the stronger thesis. However, once I have shown Strong Identity to be false, I shall evaluate the strength of Dougherty's argument for Weak Identity.

The negation of Strong Identity is that it is conceptually possible that there is some agent for whom epistemic normativity is non-reducible to morality or prudence. Thus, we must determine under what circumstances some normative standpoint might matter for an agent. One answer is that only standpoints that matter universally are capable of mattering for an agent. So, for example, only morality or prudence—if, indeed, these standpoints matter—or some similar standpoint can matter for an agent. While this would certainly aid the Doughertian argument, it seems obviously false. Many agents abide by the norms of clubs that they belong to. If one is a member of the Society for Creative Anachronism, then one may have a practical reason to bow to the group's monarch. It seems that an agent can take on a *commitment*—an adoption by an agent of a normative standpoint such that it matters to that agent—to virtually any normative standpoint. This, is, then how I propose we understand the circumstances under which normative standpoints matter. First, any normative standpoint that is universal matters for every agent. Second, any normative standpoint that an agent takes on as a commitment to matters for that agent. I do not mean these two ways to be an exhaustive list of the ways in which a standpoint might matter to an agent. However, these two will be sufficient for my argument against Dougherty.

If agents can take on commitments to other normative standpoints, then here is a quick response to Dougherty's argument. It is, at least, conceptually possible that an agent takes on a commitment to epistemic normativity. If it is possible that an agent can take on a commitment to epistemic normativity, then it is false that necessarily, epistemic normativity does not matter for any agent. Strong Identity is false. However, I think this argument is too quick. There is an obvious way that Dougherty could resist this argument. Agents cannot take on commitments to epistemic normativity, because any commitment to

epistemic normativity would be reducible to a commitment to morality or prudence. Any commitment to a normative standpoint whose mattering is reducible to some other normative standpoint is, itself, reducible to a commitment to that other normative standpoint. That is, if normative standpoint B is reducible to A, then my commitment to B must itself be reducible to A.

However, this reducibility argument fails to capture how normative commitments actually work. Indeed, if an agent commits herself to some standpoint, then that standpoint matters independently of other standpoints. Suppose, for example, that John commits himself to his friendship with Sam. Friendship norms towards Sam—for example, an obligation to spend time with Sam—now matter for John. Suppose we offered a similar trilemma for friendship as Dougherty did for epistemic normativity.[113] Friendship obligations either don't matter, they matter due to other-regarding reasons and are really moral obligations, or they matter for self-regarding reasons and are really prudential obligations, or some combination of the two. Friendship is reducible to morality and prudence.

The preceding argument is unsound. I shall show how the standpoint of friendship cannot be fully reduced to morality and prudence. Consider the following scenario. Suppose that Sam has done something morally wrong and John is considering whether to report his friend to the authorities. John has, let us suppose, conclusive moral reason to do so. Let us also say that he would receive a monetary reward for doing so, large enough that

---

[113] Interestingly, given his argument against the existence of epistemic normativity, Dougherty seems committed to this argument. If every self-regarding reason is prudential and every other-regarding interest is moral, then it follows that any normative standpoint that matters for self-regarding or other-regarding reasons is reducible to prudence and/or morality.

it would vastly outweigh whatever negatives he would incur for reporting his friend. He thus also has conclusive prudential reason to report Sam. Thus, John has conclusive moral and prudential reasons to report Sam. Sam, even though John's action is both morally and prudentially correct, blames John for turning him in. John may even feel blameworthy for turning Sam in. This would be the case, even if both John and Sam accepted that John acted morally and prudentially rightly.

This blameworthiness of John provides us with the tools we need to defend the standpoint of friendship. One commonly accepted feature of deontic normativity is that one is blameworthy for some action only if one's action is wrong. The problem for the Doughertian argument is that, simply put, one cannot point to the moral or prudential wrongness of John's action to explain his blameworthiness. John hasn't behaved morally or prudentially wrongly. The only thing that can explain John's blameworthiness is that he has done something wrong from the standpoint of friendship. But, if this is correct, then there is some aspect of friendship that is not reducible to morality or prudence. Friendship must somehow provide reasons to act that are normatively independent of morality or prudence.

The argument we have just provided for the independent reasons provided by friendship has returned us to the additional reasons-giving force argument from above. It is true that self-regarding interests tend to be prudential, and other-regarding interests tend to be moral. The problem for Dougherty's trilemma is that this general truth does not speak against the possibility that other normative standpoints might provide additional reasons for action based on the same considerations. We can see this in the case of John and Sam. It is plausible that John has a pro-tanto moral reason not to report Sam, as doing so would

negatively affect Sam's interests: an other-regarding reason. However, John's other moral reasons outweigh this pro-tanto reason, and so he has conclusive moral reason to report Sam. Nevertheless, John also has a conclusive friendship reason not to report Sam because doing so would negatively affect Sam's interests. This would violate norms of friendship. The same consideration provides both a pro tanto moral and a conclusive friendship reason not to report Sam. If one commits oneself to friendship, friendship is independent of morality or prudence.

Some may argue that my argument for the independence of friendship is unsound. The reason that John has seemed to do wrong is that the presence of prudential reasons has morally tainted John's actions: if John reports Sam because of the reward, he is doing something morally wrong. However, John need not be motivated by his prudential reasons in order for those reasons to prudentially justify him in turning Sam in. Instead, we might say that what motivates John is that turning in Sam is the morally right thing to do. All that is required for John to do the prudentially right this is that he is prudentially justified in doing so. This is merely the difference between a normative reason, which justifies an action, and a motivating reason, which is the reason that motivates an agent to act.

By parity of reasoning, the same reasoning that guarantees the independence of friendship guarantees the independence of epistemic normativity. So long as a commitment to epistemic normativity provides an instance in which one will have an epistemic obligation that goes against one's moral and prudential obligations, and one can be appropriately blamed for violating an epistemic obligation, then epistemic obligations are not entirely reducible to morality and prudence. The burden of proof is on the defender of

reducibility to show how epistemic normativity is different than friendship. I, personally, do not think such an argument can be offered.

Thus, Dougherty has failed to show that Strong Identity is true. Furthermore, my argument against Dougherty should be taken as a positive argument that it is at least conceptually possible that epistemic normativity matters for an agent. This is a significant blow against the Doughertian argument. If epistemic normativity does matter for certain agents, then we will want to be able to explain precisely what epistemic norms entail. This would be the same for an explanation of any other normative standpoint, although it is important to provide a complete description of epistemic normativity, specifically. Many people seem, at least implicitly, committed to epistemic normativity. Indeed, that agents have epistemic obligations would seem to be a useful tool in answering certain epistemic questions.[114] This is enough to safeguard epistemic normativity from Doughertian eliminitivism. However, in the interests of completeness, I shall turn to Dougherty's argument for Weak Identity in the next section.

3.7. Universal Normativity

Let us examine where we are in the scope of the paper so far. Dougherty has argued that epistemic normativity's mattering is reducible to prudence and morality. I have argued that this is too quick. A myriad number of standpoints exist, and, in principle, so long as an

---

[114] I have in mind here cases like Norman the Clairvoyant, or the New Evil Demon Case. See: Bonjour, "Externalist Theories of Empirical Knowledge"; Cohen, "Justification and Truth."

agent commits themselves to a particular standpoint, that standpoint matters for that agent. It is conceptually possible that epistemic normativity matters for all agents. Further, I think many agents do take on commitments to epistemic normativity. Some persons take epistemic obligations very seriously. Scientific research, for example, depends upon scientists taking their epistemic obligations seriously.

Given this, the question remains of whether there is anything left to Dougherty's reductionist project. There is one aspect left. Namely, we might still question whether epistemic normativity matters *universally* for agents. A standpoint matters universally when every agent has a practical reason to abide by the norms of that standpoint. Even if we show that epistemic normativity can matter, that does not show that it necessarily matters.

Before even examining the question of whether epistemic normativity matters universally, I want to highlight several reasons I am skeptical of the importance of this question. As I have said previously, I think showing that Strong Identity is false is sufficient to motivate interest in epistemic normativity. Another reason to be skeptical of the importance of determining whether epistemic normativity matters universally is that it is controversial whether any normative standpoint matters for every agent. One of the questions of metaethics is whether everyone has practical reason to be moral. Many metaethicists have attempted to prove that everyone does have a reason to be moral.[115] Other metaethicists, the normative pluralists, such as David Copp and Evan Tiffany, argue that there is no sense in which morality—or any standpoint for that matter—matters

---

[115] See: Korsgaard, *The Sources of Normativity*.

universally.[116] Normative pluralists would have reason to be suspicious of Dougherty's argument from the start. Sure, epistemic normativity does not matter universally, but neither do morality or prudence.

Thus, while it might be interesting to provide an answer to the question of whether epistemic normativity matters universally, it is certainly not necessary to defending epistemic normativity and, further, answering the question presupposes certain controversial metaethical commitments that we need not commit ourselves to. Nevertheless, I take it that some defenders of epistemic normativity might want to prove that epistemic normativity matters universally. I can certainly see the appeal of this position and am sympathetic to it. For that reason, I want, before concluding this paper, to briefly consider whether Dougherty's argument proves that epistemic normativity does not matter universally.

I think that in the context of the question of whether epistemic normativity matters universally, Dougherty's trilemma is a powerful argument. If one tries to argue that epistemic normativity is universal, but every reason one provides for its universality is reducible to moral or prudential reasons, that would be incredibly problematic for the defender of the universality of epistemic norms. After all, if morality and prudence is doing the work of explaining our seeming universal epistemic obligations, there is no work left for epistemic obligations to do.

---

[116] See: Tiffany, "Deflationary Normative Pluralism"; Copp, "Toward a Pluralist and Teleological Theory of Normativity."

Note that we cannot appeal to commitments to defend universal epistemic obligations. That is because agents only take on commitments to standpoints that do not already matter for them. The thesis Weak Identity is attacking is that epistemic normativity already matters for every agent. Thus, if we are to show that Dougherty's argument is unsound, we must look elsewhere than voluntary commitments of agents.

For all its strength, Dougherty's argument is unsound. Reducibility of one normative standpoint to another does not entail that we should eliminate the reducible normative standpoint. Suppose that Dougherty is right that the mattering of epistemic normativity is reducible to morality and prudence, and, on this basis, we should eliminate the category of epistemic normativity. Problematically, the reducibility of normative standpoints does not end here. We can posit an isomorphic argument for the reducibility of morality and prudence to practical reason. Moral (prudential) obligations only matter if one has a practical reason to abide by morality (prudence). Recall that, on our definition, a standpoint's "mattering" means that one has practical reason to abide by the norms of that standpoint. If having a moral (prudential) reason does not entail having a practical reason, then moral (prudential) reasons simply do not matter.[117] If Dougherty's argument is sound, however, then this reducibility entails elimination. We ought to abandon these standpoints—and, really, all standpoints, since if a normative standpoint matters, it matters in virtue of practical rationality—and, instead, speak of normativity only in terms of practical rationality.

---

[117] I take this to be similar to the standpoint of the normative pluralists. There is no normative standpoint that matters universally.

My version of the reducibility argument elucidates, helpfully, the strengths and weaknesses of Dougherty's argument. Positively, my argument for the reducibility of morality and prudence to practical rationality seems sound. Morality's mattering is entirely dependent on practical reason. This is a feature of practical rationality, not a bug. Every standpoint that matters matters because it provides an important—albeit, incomplete—answer to the question of the generic question of what an agent ought to do. If one had no practical reason to act morally, one would have no reason to do so.

The trouble for my argument is not in the reducibility, but in the eliminitivism. We should not give up talk of morality and prudence simply because they are reducible to practical rationality. Practical rationality is simply too coarse-grained a standpoint of practical normativity to provide any meaningful insight into the ways in which different considerations shape our decision-making. Consider why we focus on morality and prudence so much in our study of practical normativity. The normative standpoint of morality picks out the ways in which all of our other-regarding reasons bear on what we should do. For example, that my action will hurt my neighbor, help a bystander, and do nothing for a stranger are all important considerations for morality. These are importantly different than self-regarding reasons such as how some action will affect my well-being. Of course, these moral considerations—hurting my neighbor, and so on—are all practical reasons. But labeling them only as practical reasons fails to appreciate the way in which the analysis of our actions from the perspective of how it will help others is essential to our practical decision-making. A similar point can be made for prudential reasoning. At the very least, it is important to be able to talk about these reasons—prudential and moral—as

meaningfully different from one another, and to recognize the ways in which one set of reasons might yield a different verdict on what to do than another set of verdicts.

Further, practical reasons from one normative standpoint might be more or less weighty than practical reasons from another. For example, some think that moral reasons hold *supremacy* over all other kinds of reasons: if one has conclusive moral reason—a moral obligation—to perform some action, then one practically ought—one has conclusive practical reason—to perform that action.[118] This need not be the only consideration. There are even combinatorial possibilities: some action is practically permissible only if it is both morally permissible and prudentially permissible. These considerations provide good reasons to think that we ought not to eliminate the categories of morality and prudence simply because they are reducible to practical rationality.

The reasoning, then, is clear. Mere reducibility of one standpoint to another does not entail elimination of the standpoint. There might be many different reasons to keep the standpoint, regardless of reducibility. With epistemic normativity, it might be important to maintain a standpoint that deals exclusively with the obligations agents have to affect their beliefs. There might even be special normative force that epistemic obligations have like the supremacy of morality that requires that we maintain the category. For this reason, Dougherty's reducibility argument has failed to show that we should abandon the category of epistemic normativity: even at a foundational level.

---

[118] As far as I am aware, Dale Dorsey most clearly elaborates this position, though ultimately rejects it. See: Dorsey, *The Limits of Moral Authority*.

3.8. Conclusion

It would be problematic if we had to eliminate epistemic normativity from our discourse. It is, at the very least, useful to be able to discuss the epistemic obligations that agents have independent of moral or prudential considerations. Dougherty's trilemma seems, on first face, like a challenging objection to the thesis of epistemic normativity. However, it fails to show either (i) that epistemic normativity cannot matter or (ii) that epistemic normativity does not universally matter. For my own part, while I think that epistemic norms can matter, I am not at all convinced that they matter universally. I am skeptical that many standpoints beyond prudence or morality, could matter universally. Nevertheless, such an analysis is beyond the scope of this paper. It is enough for our purposes that epistemic normativity does matter to some.

Conclusion

I have argued that we can provide an account of epistemic normativity by appealing to the actions that agents take that causally influence their doxastic states. In addition, I have argued that these epistemic duties cannot be reduced to mere moral and prudential duties—at least for those who have committed themselves to epistemic normativity—and that the truth of my account does not depend upon whether we are able to direct assess the permissibility of beliefs; we have reason to indirectly assess beliefs regardless. Epistemic normativity exists, and it is important (at least to some). Before concluding, I want to first look at one potential application of this theory of epistemic normativity.

One interesting upshot of this project that came out in Chapter 2 is that, even though this dissertation has been focused on epistemic obligations to act in certain ways so as to produce certain beliefs, we might also have moral or prudential or other obligations to act in certain ways so as to produce certain beliefs. For example, I may have a moral obligation not to possess a racist belief.

The potential applications of these epistemic, moral, prudential and other obligations are numerous. However, I want to offer one potential application from the medical field. Consider the following example. Suppose that a medical patient is unconscious and unable to make decisions for himself. His family and the medical team treating him are at odds about what to do in this situation. The family wants to disable the patient's Implantable Cardioverter Defibrillator ICD. However, the company that provides the ICD will only do so if the patient has expressly asked for it. Now, if the patient has

never expressed the circumstances under which he would want his ICD disabled, then there may be nothing to be done.[119]

However, this is exactly where indirect duties to believe can come into play. The physician who installed the ICD in the patient, foreseeing that there might come a point in the future where knowing the patient's preferences will be important, may have an obligation (epistemic or moral) to have a conversation with the patient ahead of time about what his preferences are regarding when, if ever, his ICD should be disabled.

In fact, the diachronic reasoning in epistemic normativity can be used to justify all kinds of ways in which good communication can be prioritized in healthcare. Information gathered by clinicians in the present may be useful in the future. Likewise, if we take seriously the idea from Chapter 1 that our epistemic duties might be universalist—concerned with everyone's interests—rather than egoist—concerned with only our own interests—we might hold that doctors have duties to inform their patients about their prognosis and the goals of healthcare so as to avoid conflicts in beliefs in the future. One example of this would be in the case of *futile* treatment; there is no chance that some medical treatment will bring about any beneficial consequences. It is generally accepted that doctors have no ethical duty to provide treatment to a patient for whom that treatment would be completely ineffective. Indeed, they have a duty not to provide such treatment. Often, however, patients want some treatment when a doctor knows it's medically futile. This creates intractable conflicts. Since, generally speaking, medical ethics holds that

---

[119] This example is actually based on several real-world cases.

patient's preferences should be respected but that doctors shouldn't perform medically futile treatments.

However, duties to affect beliefs could be helpful here. For example, if the patients are properly informed that some treatment won't be effective there is a good chance there won't be a conflict in the first place. This suggests that doctors might have moral obligations, over and above the obligation not to perform medically futile treatment, to appropriately inform patients about various different treatment options. Such reasoning would helpfully elaborate moral duties in the medical field, among others.

The potential applications of this theory of the indirect normativity of belief are intriguing. To get there, we must be do additional work to answer certain key questions. For example, are indirect duties to believe egoist or universal? This question will be essential for answering whether or not I have any duty to inform others, or merely have a duty to be informed. As we continue research into epistemic normativity specifically, but the normativity of belief generally, we may find it helpful for resolving all sorts of ethical problems, like the problems in medical ethics countenanced in this section. This is where I propose we direct our attention from here on.

Bibliography

Alston, William P. 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257–99.

———. 1989. "Concepts of Epistemic Justification." In *Epistemic Justification : Essays in the Theory of Knowledge*, 81–114. Ithaca: Cornell University Press.

Audi, Robert. 2001. *The Architecture of Reason: The Structure and Substance of Rationality*. 1 edition. Oxford ; New York: Oxford University Press.

———. 2015. *Rational Belief: Structure, Grounds, and Intellectual Virtue*. Oxford University Press.

Axtell, Guy. 2011. "Recovering Responsibility." *Logos and Episteme*, no. 3: 429–454.

———. 2012. "(More) Springs of My Discontent." *Logos and Episteme* 3 (1): 131–137.

Baehr, Jason. 2009. "Evidentialism, Vice, and Virtue." *Philosophy and Phenomenological Research* 78 (3): 545–67.

Berman, Mitchell N. 2003. "Justification and Excuse, Law and Morality." *Duke Law Journal* 53 (1): 1–77.

Bonjour, Laurence. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5 (1): 53–73.

BonJour, Laurence. 1988. *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.

Booth, Anthony Robert, and Rik Peels. 2010. "Why Responsible Belief Is Blameless Belief." *Journal of Philosophy* 107 (5): 257–265.

Bratman, Michael E. 1999. *Intention, Plans, and Practical Reason*. Cambridge University Press.

Chisholm, Roderick M. 1989. *Theory of Knowledge*. Prentice Hall.

Clifford, William Kingdon. 1947. *The Ethics of Belief and Other Essays*. Prometheus Books.

Cohen, Stewart. 1984. "Justification and Truth." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 46 (3): 279–95.

Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford University Press.

Copp, David. 1990. "Explanation and Justification in Ethics." *Ethics* 100 (2): 237–58.

———. 2009. "Toward a Pluralist and Teleological Theory of Normativity." *Philosophical Issues* 19: 21–37.

Dancy, Jonathan. 2000. *Practical Reality*. Oxford University Press.

Derose, Keith. 2000. "Ought We to Follow Our Evidence?" *Philosophy and Phenomenological Research* 60 (3): 697–706.

Dorsey, Dale. 2016. *The Limits of Moral Authority*. United Kingdom: Oxford University Press.

Dougherty, Trent. 2011. "Re-Reducing Responsibility: Reply to Axtell." *Logos & Episteme* 2 (4).

———. 2012a. "Internalist Evidentialism and Epistemic Virtue: Re-Reply to Axtell." *Logos and Episteme* 3 (2): 281–289.

———. 2012b. "Reducing Responsibility: An Evidentialist Account of Epistemic Blame." *European Journal of Philosophy* 20 (4): 534–47.

———. 2014. *The 'Ethics of Belief' Is Ethics (Period): Reassigning Responsibilism*. Oxford University Press.

Fantl, Jeremy, and Matthew McGrath. 2012. *Knowledge in an Uncertain World*. Reprint edition. Oxford ; New York: Oxford University Press.

Feldman, Richard. 1988. "Epistemic Obligations." *Philosophical Perspectives* 2: 235–56.

———. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60 (3): 667–95.

Fischer, John Martin. 1987. "Responsiveness and Moral Responsibility." In *Responsibility, Character, and the Emotions: New Essays in Moral Psychology*. Cambridge: Cambridge University Press.

Frankfurt, Harry G. 1969. "Alternate Possibilities and Moral Responsibility." *The Journal of Philosophy* 66 (23): 829–39.

Goldman, Alvin I. 1979. "What Is Justified Belief?" In *Epistemology. An Anthology*, edited by Ernest Sosa and Jaegwon Kim, 340–353. Blackwell.

Graham, Peter A. 2011. "'Ought' and Ability." *Philosophical Review* 120 (3): 337–82.

Haack, Susan. 2001. "'The Ethics of Belief' Reconsidered." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 21–33. Oxford University Press.

Haji, Ishtiyaque. 2010. "Incompatibilism and Prudential Obligation." *Canadian Journal of Philosophy* 40 (3): 385–410.

———. 2012. "Reason, Responsibility, and Free Will: Reply to My Critics." *The Journal of Ethics* 16 (2): 175–209.

Harman, Gilbert. 1986. *Change in View*. Cambridge, Mass: MIT Press.

Heller, Mark. 2000. "Hobartian Voluntarism: Grounding a Deontological Conception of Epistemic Justification." *Pacific Philosophical Quarterly* 81 (2): 130–41.

Hieronymi, Pamela. 2008. "Responsibility for Believing." *Synthese* 161 (3): 357–73.

Holton, Richard. 2009. *Willing, Wanting, Waiting*. Oxford University Press UK.

Huss, Brian. 2009. "Three Challenges (and Three Replies) to the Ethics of Belief." *Synthese* 168 (2): 249–71.

James, William. 1979. "The Will to Believe." In *The Will to Believe and Other Essays in Popular Philosophy*, 13–33. Cambridge, Mass: Harvard University Press.

Kim, Jaegwon. 1988. "What Is 'Naturalized Epistemology?'" *Philosophical Perspectives* 2: 381–405.

Kornblith, Hilary. 2001. "Epistemic Obligation and the Possibility of Internalism." In *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, edited by Abrol Fairweather and Linda Trinkaus Zagzebski, 231–248. Oxford University Press.

Korsgaard, Christine M. 1996. *The Sources of Normativity*. Cambridge University Press.

Kvanvig, Jonathan L. 2003. "Propositionalism and the Perspectival Character of Justification." *American Philosophical Quarterly* 40 (1): 3–17.

Leary, Stephanie. 2017. "In Defense of Practical Reasons for Belief." *Australasian Journal of Philosophy* 95 (3): 529–542.

Markovits, Julia. 2011. "Why Be an Internalist About Reasons?" *Oxford Studies in Metaethics* 6: 255.

Mason, Elinor. 2003. "Consequentialism and the 'Ought Implies Can' Principle." *American Philosophical Quarterly* 40 (4): 319–31.

Mele, Alfred, and Steven Sverdlik. 1996. "Intention, Intentional Action, and Moral Responsibility." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 82 (3): 265–87.

Mill, John Stuart. 2002. *Utilitarianism*. Edited by George Sher. Second Edition,2 edition. Indianapolis: Hackett Publishing Company, Inc.

Moran, Richard A. 2001. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton University Press.

Nagel, Thomas. 1979. "The Fragmentation of Value." In *Mortal Questions*, edited by Thomas Nagel. Cambridge University Press.

Nottelmann, Nikolaj. 2007. *Blameworthy Belief: A Study in Epistemic Deontologism*. Dordecht: Springer Science & Business Media.

———. 2013. "The Deontological Conception of Epistemic Justification: A Reassessment." *Synthese* 190 (12): 2219–41.

Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.

Pascal, Blaise. 2017. *Pensées*. GIANLUCA.

Peels, Rik. 2014. "Against Doxastic Compatibilism." *Philosophy and Phenomenological Research* 89 (3): 679–702.

———. 2015. "Believing at Will Is Possible." *Australasian Journal of Philosophy* 93 (3): 524–41.

———. 2016. *Responsible Belief: A Theory in Ethics and Epistemology*. 1 edition. New York, NY: Oxford University Press.

———. 2017. "Responsible Belief and Epistemic Justification." *Synthese* 194 (August): 2895–2915.

Pollock, John L., and Joseph Cruz. 1999. *Contemporary Theories of Knowledge*. Second edition. Lanham, Md: Rowman & Littlefield Publishers, Inc.

Regis, Edward. 1980. "What Is Ethical Egoism?" *Ethics* 91 (October): 50–62.

Ryan, Sharon. 2003. "Doxastic Compatibilism and the Ethics of Belief." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 114 (1/2): 47–79.

Searle, John R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.

Sidgwick, Henry. 1981. *The Methods of Ethics, 7th Edition*. 7th edition. Indianapolis: Hackett Publishing Company.

Sinnott-Armstrong, Walter. 1992. "An Argument for Consequentialism." *Philosophical Perspectives* 6: 399–421.

Steup, Matthias. 2000. "Doxastic Voluntarism and Epistemic Deontology." *Acta Analytica* 15 (1): 25–56.

———. 2008. "Doxastic Freedom." *Synthese* 161 (3): 375–392.

———. 2011. "Belief, Voluntariness and Intentionality." *Dialectica* 65 (4): 537–559.

———. 2012. "Belief Control and Intentionality." *Synthese* 188 (2): 145–163.

Stroud, Sarah. 1998. "Moral Overridingness and Moral Theory." *Pacific Philosophical Quarterly* 79 (2): 170–89.

Talbot, Brian. 2016. "The Best Argument for 'Ought Implies Can' Is a Better Argument Against 'Ought Implies Can.'" *Ergo, an Open Access Journal of Philosophy* 3.

Tiffany, Evan. 2007. "Deflationary Normative Pluralism." *Canadian Journal of Philosophy* 37 (sup1): 231–62.

Vallentyne, Peter. 1987. "Prohibition Dilemmas and Deontic Logic." *Logique et Analyse* 30 (117/118): 113–22.

———. 1989. "Two Types of Moral Dilemmas." *Erkenntnis (1975-)* 30 (3): 301–18.

Vranas, Peter B. M. 2007. "I Ought, Therefore I Can." *Philosophical Studies* 136 (2): 167–216.

William, James. 1896. *The Will to Believe*. New York: Longmans, Green and Co.

Yudell, Michael, Dorothy Roberts, Rob DeSalle, and Sarah Tishkoff. 2016. "Taking Race out of Human Genetics." *Science* 351 (6273): 564–65.

Vita

Stephen Perinchery-Herman was born on April 18th, 1990 in San Antonio, Texas. He grew up in Yorktown, NY and graduated from the Montfort Academy in Katonah, NY in 2008. He attended Binghamton University where he received a B.A. in Philosophy and Latin in 2012. He received a M.A. in Philosophy from Georgia State University in 2014. He then travelled to Missouri where he received an M.A. in Philosophy in 2016 and a PhD in Philosophy in 2019. He is married to Anna Perinchery-Herman, whom he met at the University of Missouri. When not researching, he enjoys hiking, reading, gardening, cooking, and drawing.