

AGENCY AND REASONS IN EPISTEMOLOGY

A Dissertation Presented

by

LUIS R.G. OLIVEIRA

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

May 2016

Philosophy

© Copyright by Luis R.G. Oliveira 2016

All Rights Reserved

AGENCY AND REASONS IN EPISTEMOLOGY

A Dissertation Presented

by

LUIS R.G. OLIVEIRA

Approved as to style and content by:

Hilary Kornblith, Chair

Ernesto V. Garcia, Member

Christopher J.G. Meacham, Member

Seth Cable, Member

Joseph Levine, Department Head
Philosophy

DEDICATION

To my nephew and my niece, Daniel and Julia.

We need a method if we are to investigate the truth of things.

— *René Descartes, 1701*

You don't need a weatherman to know which way the wind blows.

— *Bob Dylan, 1963*

ACKNOWLEDGMENTS

Far too many people deserve credit for this dissertation. Members of my immediate and extended family; fellow philosophers at Calvin College and at UMass Amherst; friends scattered around this country and the world. I am very grateful to them all.

I begin by thanking the elder members of my family (I have too many cousins to mention them by name): Hernani and Marcia Gouvêa, Fernando and Mari Gouvêa, Lisias and Liliana Ransan, Ricardo and Lilia Gouvêa, Lisianas and Waléria Ransan, Célia de Gouvêa Franco and Celso Pinto, Edin and Helda Abumanssur, Heron and Suzi Oliveira, Helton Oliveira, Kerstin Westerståhl, and Márcio and Nelma Esteves de Andrade. They have done much for me—much more than was their duty—throughout my life. I would also like to thank my sisters and my brother-in-law: Ana Raquel and Milton Santos, and Viviane Oliveira. They are my closest and dearest friends; no one knows me like they do.

Now I would like to thank some of my Calvin College and UMass Amherst professors: Kevin Corcoran, Kelly Clark, Rebecca DeYoung, Gregory Mellema, Terence Cuneo, Matt Halteman, Ruth Groenhout, Jamie Smith, Fred Feldman, Joe Levine, Pete Graham, Eileen O’Neill, Louise Antony, Alejandro Pérez Carballo, Vanessa de Harven, Chris Meacham, and Ernesto Garcia. Each of

them has benefited me with either their friendship or their guidance, and some of them have benefited me with both. No doubt they will notice that portions of their views and styles are sprinkled throughout this dissertation. I stole from them as much as I could.

I would also like to thank a legion of friends: Mário Ransan, Marcos Gouvêa, Luis Pinto, Clésio Geraldo, Márcio Júnior, Felipe San Giovanni, Paul Freston, Nathan Ward, Elizabeth Damstra, Michael Boogaard, Ross Thelen, C.J. Majeski, Cedric Parsels, Timothy Perrine, Mateus Macedo, Jordan Kohl, Chris Mangione, Brandon Hahn, Shannon Moselle, Rachelle Velgersdyk, Renee Meadors, Reid Coe, Nevin Schaller, Ken Bekendam, Erin Breiner, Ted Koopmans, Tricia Magalotti, Dennis Kavlakoglu, Edward Ferrier, Heidi and Mac Furey, Josh and Tara DiPaolo, Liz Fox and Dan McGloin, Bob Gruber, Miles Tucker, Scott Hill, Kristian Olsen, and Jeffrey Spooner. These are folks from different times and different lives. In a variety ways, I have benefited from knowing every single one of them much more than they have benefitted from knowing me.

Three monumental influences, however, stand alone above all else. I am specially grateful to my uncle Ricardo Gouvêa, far and away the most responsible for my early interest in the life of the mind. I have learned more from him in friendly conversation than in serious study from anyone else; I could not hope to measure how much of him there is in me. Ironically, there is little in this dissertation that will tingle his philosophical predilections—except, perhaps, this little bit of irony itself. I am also specially grateful to my undergraduate advisor Stephen Wykstra. Steve taught me how to properly read a philosophi-

cal text and taught me how to think carefully about a philosophical idea. To this day, his devotion to the yeoman labor of a well-written essay continues to inspire and embarrass my own efforts towards *le mot juste*. Whatever modest measure of philosophical style I may have, I certainly owe it to Steve. Finally, I am specially grateful to my dissertation advisor Hilary Kornblith. I have been influenced in equal measures by his philosophical naturalism, his epistemological externalism, and his career-long effort to make philosophy into more than the mere musings of an empirically detached elite. Throughout graduate school, he has spoiled me with generous portions of his time and with the encouraging portions of his thoughts. In his work and in his ways, Hilary has given me the finest example of the philosopher that I wish to be. I am deeply in his debt.

Yet none of these various people would have mattered without the sacrifices made by my parents, Hilton and Lenora, for my education. Looking back, I see that I failed to meet even the lowest standards of gratitude over and over and over. I am embarrassed by that fact. I could not see, back then, what my privileged education was costing them, and what it would one day allow me to afford. I see it now. Thank you. I hope this dissertation stands as proof that your efforts were not in vain.

Luis R.G. Oliveira

UMass Amherst, 2016

ABSTRACT

AGENCY AND REASONS IN EPISTEMOLOGY

MAY 2016

LUIS R.G. OLIVEIRA

B.A., CALVIN COLLEGE

M.A., UNIVERSITY OF MASSACHUSETTS AMHERST

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Hilary Kornblith

Ever since John Locke, philosophers have discussed the possibility of a *normative epistemology*: are there epistemic obligations binding the cognitive economy of belief and disbelief? Locke's influential answer was *evidentialist*: we have an epistemic obligation to believe in accordance with our evidence. In this dissertation, I place the contemporary literature on *agency* and *reasons* at the service of some such normative epistemology. I discuss the semantics of obligations, the connection between obligations and reasons to believe, the implausibility of Lockean evidentialism, and some of the alleged connections between agency and justification.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	vi
ABSTRACT	ix
 CHAPTER	
1. INTRODUCTION	1
1.1 The Meta-Normative Big Picture	1
1.2 The Meta-Normative Extensions	3
1.3 The Epistemological Focus	4
1.4 Brief Summary of Each Chapter	6
2. EXPRESSING THE ‘OUGHT’ OF MOST REASON	9
2.1 Introduction	9
2.2 Schroeder Against the Uniformity View	12
2.3 The Normative Challenge	15
2.4 Privileged Standards of Evaluation	20
2.5 Super Standards of Evaluation	28
2.6 The Contextual Standards Dilemma	34
2.7 Conclusion	41
3. ON EVADING THE DEONTIC PUZZLE	42
3.1 Introduction	42
3.2 Evaluations and Prescriptions	44
3.3 The Evasive Strategy	49
3.4 Epistemic Responsibility	53
3.5 Undefeated Normative Reasons	59

3.6	Letting Reasons Go	65
3.7	Conclusion	70
4.	DEONTOLOGICAL EVIDENTIALISM, WIDE-SCOPE REQUIREMENTS, AND PRIVILEGED VALUES	72
4.1	Introduction	72
4.2	The Consequentialist Moral Argument	75
4.3	The Grimm-Inspired Moral Argument	80
4.4	The Epistemic Point-Of-View Argument	87
4.5	Kornblith's Incomplete Reply to Premise (4)	93
4.6	Normative Requirements and Privileged Values	98
4.7	Privileged Values and the Nature of Normativity	105
4.8	Conclusion	107
5.	DEONTOLOGICAL EVIDENTIALISM AND THE PRINCIPLE THAT <i>OUGHT</i> IMPLIES <i>CAN</i>	109
5.1	Introduction	109
5.2	Reasons, Deliberation, and Transparency	112
5.3	What if OIC is False?	118
5.4	What if OIC is True?	124
5.5	Conclusion	129
6.	AMPLIATIVE TRANSMISSION AND DEONTOLOGICAL INTERNALISM	132
6.1	Introduction	132
6.2	Clarifying (BAT)	134
6.3	The Intuitive Defense of (BAT)	136
6.4	(BAT) and Transmission Failure	138
6.5	Extending (BAT)	141
6.6	Trouble for Deontological Internalism	143
6.7	Conclusion	148
7.	NON-AGENTIAL PERMISSIBILITY IN EPISTEMOLOGY	150
7.1	Introduction	150
7.2	Non-Agential Permissibility	153

7.3 Praise and Blame	156
7.4 Conclusion	158
BIBLIOGRAPHY.....	159

CHAPTER 1

INTRODUCTION

1.1 The Meta-Normative Big Picture

My dissertation centers around a familiar experience and a hypothesized relation underlying it. The familiar experience is that of being *required* to perform an action or to take on an attitude. We experience the authority of this requirement in feeling, in belief, and in discourse. The hypothesized relation underlying it is what I call a *normative requirement*. The hypothesis, more exactly, is that our varied feelings, beliefs, and discourse about what is required of us and others are at bottom feelings, beliefs, and discourse about an authoritative relation of normative requirement that holds between individuals and acts or attitudes. I attempt to characterize this experience and this relation in various ways throughout the following chapters, but I do not defend directly the hypothesis that the latter truly underlies the former.

My dissertation also accepts that this hypothesized relation between individuals and acts or attitudes can be characterized by two substantive principles. One of these principles is what I call *normative rationalism*. The other is what I call *normative consequentialism*:

(NR) There is a normative requirement for S to ϕ if and only if (and because) there is most normative reason for S to ϕ .¹

(NC) There is a normative reason for S to ϕ only if (and because) S ϕ -ing promotes some value v .²

The first principle hypothesizes that our multi-faceted experience of normative requirements is, fundamentally, a *rational* experience. The existence of the relation of normative requirement, that is, depends on the prior existence of normative reasons. The second principle hypothesizes that *value*, whatever that turns out to be, is the most fundamental feature of the normative aspect of the world. The existence of normative reasons, that is, depends on the prior existence of values. These two principles are neither idiosyncratic nor uncontroversial. I attempt to draw modest support for these principles here and there, but I do not offer a sustained defense of them. Instead, (NR) and (NC) form the starting position from which I launch into my discussions.³

¹I take this to be one way of expressing the view made famous by W.D. Ross (1930). For more recent developments, see Broome (1999), Raz (2002), Skorupski (2010), Parfit (2011), and Smith (2013). See Schroeder (2007) for the contrast between normative, explanatory, and motivating reasons.

²See Raz (2001).

³My dissertation thereby enjoins the reader in a way familiar from Peter Railton (1999, 341): “for now, let us cheerfully assume that we don’t mind mystery, as long as its name is rationality.” Various aspects of the mystery of rationality, however, are here prodded and pulled and stretched—for verification, that is, of their cost and plausibility. What my chapters below reveal is what is and what is not compatible with it. The reader is thus invited to make her own appraisals of what is worth retaining and what is worth abandoning in these packages of incompatibilities.

1.2 The Meta-Normative Extensions

(NR) and (NC) are my background assumptions throughout. From them, my dissertation defends a further principle:

(NC*) There is a normative reason for S to ϕ only if S ϕ -ing promotes some privileged value pv .

The key motivation behind (NC*) is the fact that ‘value’ is very often used in a *normatively thin* sense, where it only indicates what is guiding a certain class of evaluations. I argue that the kind of value that is at the foundation of our experience of normative requirements is somehow privileged over these more ordinary kinds.

One of my chapters, moreover, rejects an influential argument against the converse of (NC*). Since I think this argument is in fact one of the main arguments against the converse of (NC*), I take its rejection as indirect support for it. If I am right about this, then the following principle is true:

(NC**) There is a normative reason for S to ϕ if and only if (and because) S ϕ -ing promotes some privileged value pv .⁴

Together, (NR) and (NC**) form the position that I call *Reasons Consequentialism*. This dissertation is by no means a full defense of this position. But it inches towards it at every step.

⁴Drawing on Feldman (1986), I take it that ϕ -ing promotes some privileged value pv if and only if there are no accessible possible worlds where I do not ϕ that are ranked higher by pv than accessible possible worlds where I ϕ

Many of my chapters focus on the English ‘ought’. In every case, I am assuming two things about it. First, that ‘ought’ is ambiguous between a variety of different meanings. Second, that ‘ought’ sometimes expresses the hypothesized relation of normative requirement. When used in this way, ‘ought’ is used as the *ought of most reason*. The need for being careful about which sense of ‘ought’ is at play—and being careful about which senses of ‘ought’ could or could not be at play—is a major theme throughout.

1.3 The Epistemological Focus

My dissertation avoids the assumption that there is a “most natural” meaning for epistemological terms. It also avoid the assumption that there is a “shared concept” behind our epistemological terms. Instead, I assume that epistemological terms are largely technical terms, flexible to the ends of the technicians that are using them. I thus avoid the general argumentative strategy: “since term t means/implies/expresses p , then...” Instead, I employ the alternative strategy: “if we take term t to mean/imply/express p , then...” At times I try to make explicit the value of using a term for a certain purpose; at times I settle for evidence that many others in fact seem to use it in that way. In either case, my aim is to make clear how exactly I am using my epistemological terminology.

From the standpoint of Reason Consequentialism, my dissertation examines what it would look like for epistemology to take itself as making claims about normative requirements—and what it would look like otherwise. As I read them, the so-called fathers of contemporary epistemology—Descartes and Locke—

took their claims as more than parochial evaluations from an idiosyncratic point-of-view. They rather took themselves as identifying what was required of our ordinary practices of belief. In a representative passage, for example, John Locke in fact prescribes a close connection between evidence and belief that is not at all out of step with contemporary epistemology:

We should keep a perfect indifference for all opinions, not wish any of them true, or try to make them appear so; but being indifferent, receive and embrace them according as evidence, and that alone, gives the attestation of truth. (*Conduct*, §34)

Here and elsewhere, we can say, Descartes and Locke were after an “ethics of belief.” My dissertation examines different aspects of epistemology, understood in this traditional way, from the standpoint of Reason Consequentialism.

One obvious consequence of my method is that one can simply reject my explicit assumptions at every turn. One can say that epistemology, or a certain term, is not about normative requirements; one can say that normative requirements are not about reasons; and so on. Most of my arguments can be blocked that way. I think that's fair. In fact, I think this is true of much of philosophy: we hypothesize certain fundamental explanatory elements, take on certain methodological commitments, and draw their consequences. How one frames a certain problem in fact reveals as much about one's philosophy as how one solves it. I do my best throughout to be crystal clear about the salient background assumptions and the arguments that I launch from them.

1.4 Brief Summary of Each Chapter

In chapter 1, I discuss *the uniformity view*: the claim that the English ‘ought’ is always a raising verb expressing a propositional operator. I here examine what this view would have to look like if it were to allow for the expression of the ‘ought’ of most reason. After arguing that two recently influential versions of the uniformity view—due to John Horty (2001) and Stephen Finlay & Justin Snedegar (2014)—cannot succeed, I articulate a version that can. My diagnosis is that too little attention has been paid, so far, to the contextual standards of evaluation that figure into ought-claims. I conclude by arguing that allowing for the expression of the ‘ought’ of most reason forces the uniformity view into a dilemma.

In chapter 2, I discuss *the deontic puzzle*: the tension between Doxastic Involuntarism (the claim that we do not have voluntary control over our beliefs), Epistemic Deontology (the claim that ‘S ought to believe that p’ is sometimes true), and the Ought-Implies-Can Principle. More exactly, I here criticize what I call the *evasive strategy* against this puzzle: taking doxastic ought-claims as evaluative instead of prescriptive employments of the English ‘ought’. Philosophers as different as Hilary Kornblith (2001), Richard Feldman (2008), and Matthew Chrisman (2012) think that evasion is the least controversial and least theoretically expensive strategy against the deontic puzzle. In reply, I argue that it comes at the costly price of normative reasons for belief.

In chapter 3, I discuss *deontological evidentialism*: the claim that we ought to form and maintain our beliefs in accordance with our evidence. More exactly, I here criticize three arguments in its defense. I begin by discussing Berit

Brogaard's (2014) use of the distinction between narrow-scope and wide-scope requirements against W.K. Clifford's well-known moral defense of deontological evidentialism. I then use this very distinction against a defense of deontological evidentialism inspired by Stephen Grimm's (2009) more recent claims about the moral source of epistemic normativity. I use this distinction once again to argue that Hilary Kornblith's (2001) criticism of Richard Feldman's (2000) defense of deontological evidentialism is incomplete. Finally, I argue that Feldman's defense is insensitive to the relation between normative requirements and what I will call *privileged values*: values that have normative authority over us.

In chapter 4, I continue to discuss deontological evidentialism. I here discuss a promising argument for it, whose central premise says the following: consideration *c* is a normative reason for S to form or maintain a belief that *p* only if *c* is evidence that *p* is true. More exactly, I here discuss the relation between a recent argument for this key premise—offered separately by Nishi Shah (2006) and Ward E. Jones (2009)—and the principle that *ought* implies *can*. I argue that anyone who antecedently accepts or rejects this principle already has a reason to resist either this argument's premises or its intended epistemological implications.

In chapter 5, I discuss *deontological internalism*: the family of views where justification is a positive deontological appraisal of someone's epistemic agency. According to these views, S is justified in believing that *p*, that is, when S is blameless, praiseworthy, or responsible in believing that *p*. Brian Weatherson (2008) discusses very briefly how a plausible principle of ampliative transmission reveals a worry for versions of deontological internalism formulated in terms of

epistemic blame. Weatherson denies, however, that similar principles reveal similar worries for other versions. I disagree. In this chapter, I argue that plausible principles of ampliative transmission reveal a worry for deontological internalism in general.

In chapter 6, I discuss Paul Silva's (2015) recent claim that doxastic justification does not have a basing requirement. An important part of his argument depends on the assumption that doxastic and moral permissibility must have a parallel structure. I here reply to Silva's argument by challenging this assumption. I claim that moral permissibility is an *agential* notion, while doxastic permissibility is not. I then briefly explore the nature of these notions and briefly consider their implications for praise and blame.

CHAPTER 2

EXPRESSING THE ‘OUGHT’ OF MOST REASON

2.1 Introduction

Declarative sentences employing the English ‘ought’—*ought-claims*—display a remarkable grammatical variety. We sometimes say of someone in particular that she ought to volunteer at the local shelter; we sometimes say abstractly that it ought to be that someone volunteers; we sometimes say of the state of things in general that there ought to be volunteers; and so on. Some of these various ought-claims, moreover, can be ambiguous in familiar and uncontroversial ways. An utterance of ‘it ought to rain tomorrow’ sometimes expresses a claim about what would be best to happen (as in ‘it ought to rain tomorrow, if we are to survive this drought’), and it sometimes expresses a claim about what is likely to occur (as in ‘it ought to rain tomorrow, if the readings of the Doppler radar are correct’). According to a widely accepted view, nonetheless, these two familiar features of ought-claims do not reflect an underlying semantic or syntactic difference. At the semantic level, the English ‘ought’ always represents a propositional operator; at the syntactic level, the

English ‘ought’ always functions as a raising verb. Call this the *uniformity view*.¹

According to the orthodox variety of the uniformity view, ought-claims always express that some proposition p is true in all the best contextually salient possible worlds, given a contextually salient standard of evaluation.² Let ‘ \mathbb{W} ’ stand for the set of contextually salient worlds, and let ‘ \mathbb{S} ’ stand for the contextually salient standard of evaluation. We can then say that, according to the orthodox variety of the uniformity view, the semantic structure of (simple) ought-claims is always: $O(p)^{\mathbb{W}, \mathbb{S}}$. Call ought-claims with such underlying semantic structure *evaluative* ought-claims.³

The uniformity view is capable of explaining both the grammatical variety and the common ambiguities of ought-claims by adding a raising verb syntax to this semantic picture. Since English grammar does not allow for sentences without subjects, sentences mirroring the semantic form of evaluative ought-claims—such as ‘ought Sarah volunteers at the local shelter’ or ‘ought to be volunteers at the local shelter’—are ungrammatical. As a result, the grammatical expression of evaluative ought-claims *raises* the subject of the prejacent proposition to the head of the whole sentence, when a subject is

¹For some proponents, see Horty (2001), Chrisman (2012), and Finlay & Snedegar (2014). For proponents of similar claims about the semantics of the English ‘ought’, see Chisholm (1964), Kratzer (1977), Williams (1981a), Broome (1999), and Wedgwood (2007).

²On Kratzer’s (1977; 1981) terminology, the contextually salient worlds form a *modal base* (a set of possible worlds which the ‘ought’ operator quantifies over), and the contextually salient standard consists of an *ordering source* (a conversational background inducing a partial ranking of the worlds in the modal base).

³I will hereafter mean this orthodox variety when I say ‘uniformity view’.

available: ‘Sarah ought to volunteer at the local shelter’. In cases where the prejacent has no subject, the grammatical expression of evaluative ought-claims appends to the whole sentence a non-referring ‘there’ or ‘it’: ‘there ought to be volunteers’. Such raising verb syntax thus explains how claims sharing the same underlying semantic structure ($O(p)^{W, S}$) can look so grammatically various as ought-claims do. Just as easily, the uniformity view explains common ambiguities by noticing that different contexts can make salient different kinds of standards (epistemic, deontic, prudential, etc.).

Some, however, have claimed that the uniformity view faces the following challenge: certain claims about what an agent ought to do cannot be expressed with ought-claims if ‘ought’ is *always* a raising verb expressing a propositional operator.⁴ This is far from a settled debate and my aim here is more modest than taking a stand on it. Instead, my aim is twofold: (a) clarifying how at least one of these challenges to the uniformity view gives way to a general challenge to contextualist accounts of the English ‘ought’ (of which the uniformity view is but one example); and (b) exploring what it would take for such contextualist accounts to meet this nearby challenge.

Here is how I proceed. In section 1, I outline Mark Schroeder’s (2011) argument against the uniformity view. In section 2, I clarify how his argument against the uniformity view gives way to what I call *the normative challenge* against contextualist views in general: providing a principled distinction be-

⁴Versions of this challenge have been suggested with some frequency: see Geach (1982), Harman (1986, appendix B), Feldman (1986, ch. 8), Ross (2010), and Schroeder (2011). Replies have been proposed as well: see Horty (2001, ch. 4), Chrisman (2012), and Finlay & Snedegar (2014).

tween ought-claims that express the ‘ought’ of most reason and those that don’t. In section 3, I argue that two strategies available to the uniformity view against Schroeder’s original challenge (from Horty 2001 and Finlay & Snedegar 2014) cannot succeed against the modified challenge under consideration. My diagnosis of these shortcomings is that not enough attention has been paid to the *standards of evaluation* that figure into ought-claims. In section 4, I articulate a strategy that can, in fact, succeed against the normative challenge set up in section 2. In section 5, however, I consider the way in which succeeding against the normative challenge forces contextualist accounts of the English ‘ought’ into a new dilemma about the contextual availability of normative standards. I do not think that this dilemma is fatal to these accounts, but I take it to reveal some of its hidden costs and to point the way for further exploration.

2.2 Schroeder Against the Uniformity View

Some ought-claims have an agent as their grammatical subject. This allows us to distinguish ‘Sarah ought to volunteer at the local shelter’ from ‘there ought to be volunteers at the local shelter’: the former is an *agential* ought-claim while the latter is a *non-agential* ought-claim. There is a sense in which this is merely a superficial distinction, however, given that any agential ought-claim can be turned into a non-agential ought-claim without thereby becoming ungrammatical. There is nothing ungrammatical about ‘it ought to be that Sarah volunteers at the local shelter’, for example. We can capture this sense of superficiality by saying that every agential ought-claim has a

non-agential grammatical equivalence. In fact, according to the uniformity view, this seemingly superficial distinction is really just that: superficial. That is because the following is a consequence of the uniformity view:

(equivalence) The meaning of an agential ought-claim is preserved by its non-agential grammatical equivalence.

According to the uniformity view, that is, ‘Sarah ought to volunteer at the local shelter’ and ‘it ought to be that Sarah volunteers at the local shelter’ are not just grammatically equivalent but also, all else being equal, equivalent in meaning.

According to Mark Schroeder, however, (equivalence) is false and the uniformity view therefore incorrect. Some agential ought-claims have a meaning that is not preserved by their non-agential grammatical equivalence. Schroeder (2011, 9-11) attempts to focus our attention on this distinct meaning of some agential ought-claims by noting that these are uniquely appropriate in conversational contexts of advice. Only some agential ought-claims (i) matter directly for advice, (ii) are the right kind of thing to close deliberation, (iii) correctly indicate accountability, (iv) appropriately track what the agent is able to perform, and (v) are closely connected to the notion of obligation. We can follow Schroeder in referring to agential ought-claims expressing a claim that is appropriate (in these five ways) in conversational contexts of advice as *deliberative* agential ought-claims. According to Schroeder, then, deliberative agential ought-claims have a meaning that is not preserved by their non-agential grammatical equivalence.

Consider the following example. Suppose a good friend is facing financial hardships. Suppose the following two agential ought-claims are then true:

(a) Larry ought to win the lottery.

(b) Larry ought to file for a loan.

Perhaps (a) is true since Larry is in every respect the most deserving ticket holder. Perhaps (b) is true since filing for a loan is the safest and most effective way for Larry to turn around his awful luck. No doubt (a) and (b) express very different propositions. But at the level of syntactic and semantic structure, according to the uniformity view, they are both the same: they both employ ‘ought’ as a raising verb and have the semantic form $O(p)^{\mathbb{W}, \mathbb{S}}$. Their meaning is thus preserved in their non-agential grammatical equivalences, which share the same syntactic and semantic structure as well:

(a*) It ought to be that Larry wins the lottery.

(b*) It ought to be that Larry files for a loan.

But suppose Larry comes to you for advice. Here, according to Schroeder, there seems to be nothing that we can mean by (a) such that it would matter directly for advice, be the right kind of thing to close Larry’s deliberation, correctly indicate accountability, appropriately respect Larry’s abilities, and relate to Larry’s obligations. On the other hand, there seems to be an available meaning of (b) that would be perfectly appropriate on all four counts. Moreover, when (b) is taken as expressing this distinct meaning—the meaning that is appropriate (in the five ways indicated) in conversational contexts of advice—

Schroeder’s claim is that it is then taken as expressing something that cannot be expressed by (b*): (b), in this distinct sense, is a claim about the requirements on Larry’s agency, while (b*) is merely an evaluation of a possible state of affairs. Schroeder’s (2011, 13-19) argument against the uniformity view, in short, turns on identifying a class of counterexamples to (equivalence), and on offering linguistic evidence for an alternative view of their semantics and syntax.⁵

2.3 The Normative Challenge

Notice that the debate between those who are sympathetic to the uniformity view and those who are not is more than a merely linguistic dispute. This is because both sides acknowledge that the distinct meaning of deliberative agential ought-claims is precisely the meaning that matters most to normative philosophy. It is here, in fact, that we can begin to extract a different challenge from Schroeder’s claims. Though not much of substance is said explicitly about this subtle meaning that competent English speakers are supposed to grasp—on either side of the debate—I will here take it that deliberative agential ought-claims express a claim about what an agent has most reason to do. Call this the *normative assumption*:

⁵Schroeder suggests that deliberative agential ought-claims employ the English ‘ought’ with a *control verb* syntax. A control verb—like ‘wants’—is a verb which (i) can take a further verb as an argument, and (ii) when it does, it has as *its* subject the same subject as its verb argument. This is supposed to explain why their meaning is not preserved in their non-agential grammatical equivalences. Both Chrisman (2012) and Finlay & Snedegar (2014) challenge the linguistic evidence offered by Schroeder.

The Normative Assumption: Deliberative agential ought-claims express a claim about what an agent has most reason to do.

Here, recall, an agential ought-claim is any declarative sentence employing the English ‘ought’ with an agent as its grammatical subject, and a deliberative agential ought-claim is one that is appropriate in conversational contexts of advice in the characteristic ways suggested by Schroeder. This assumption, of course, requires some defense.

There are at least two reasons for accepting the normative assumption. First, this is the sense of ‘ought’ that those within and those outside of this linguistic dispute have acknowledged as the sense that matters most to normative philosophy.⁶ The ‘ought’ of most reason, so to speak, is the sense of ‘ought’ where it expresses an authoritative relation of normative requirement or obligation, where it picks out a “specifically practical relation holding between agents and actions that is relevant to ethics” (Chrisman 2012, 436). As some have put it in different contexts, this notion of normative requirement is the notion of something “having a grip” on us (cf. Korsgaard 1996, 44-6), or being “demanded” of us (cf. Street 2012, 44). Since it is this notion—and the ‘ought’ of most reason that serves as its expression—that matters most

⁶See Harman (1986, 132-3), Broome (1999, 399-400), Zimmerman (2006, 329), Finlay (2009, 316), Ross (2010, 309), Schroeder (2011, 36), and Finlay & Snedegar (2014, 104; 122). Here and throughout this paper I will mean *normative reasons* when I speak of reasons of any kind, as opposed to explanatory or motivating reasons (cf. Schroeder 2007, 10-15). It is notoriously hard to say something informative about normative reasons that is neither vague or metaphorical. A grip on the pre-theoretical and paradigmatic sense in which a innocent person’s unnecessary suffering is a normative reason against torturing her for fun—the sense in which that *fact* is an objectively authoritative consideration against that *action*—will be enough for present purposes.

to normative philosophy, and since it is what matters most to normative philosophy that the uniformity view is challenged by deliberative agential ought-claims to accommodate, we have here one good reason to accept the normative assumption.

The second reason is that this is the only sense of ‘ought’ that could make ought-claims *appropriate*, on a normative reading of ‘appropriate’, in conversational contexts of advice in the characteristic ways suggested by Schroeder. Given idiosyncratic contexts and psychologies, nearly anything can be *effective* as advice—nearly anything can be *taken* as indicating what to do. But notice how deliberative agential ought-claims, according to Schroeder, are supposed to be more than merely effective in this way: they “matter directly”, they are “the right kind of thing”, they “correctly indicate”, they are closely connected to “obligation,” and so on. These are plausibly expressions of normative success. Yet it is only when an ought-claim expresses a claim about what an agent has most reason to do that it is appropriate, or normatively successful, in these ways.⁷ This is another good reason to accept the normative assumption.

If the normative assumption is correct, however, then the distinct meaning of (b) that according to Schroeder cannot be captured by (b*) is: filing for a loan is what Larry has most reason to do. The challenge I am here offering to the uniformity view, then, is explaining how some but not all evaluative ought-claims can express a claim about what one has most reason to do by providing a principled way to distinguish between those that do and those that

⁷See Zimmerman (2008, 6-8), for example, for discussion of the close connection between the ought of most reason and the kind of advice sought by the “conscientious person”.

don't. This is different but related to Schroeder's denial of (equivalence). Call it *the normative challenge*. Importantly, this challenge targets more than just the uniformity view. Any contextualist account of the semantics of the English 'ought' owes a similar explanation.

Let me note one more way in which the normative challenge is distinct from Schroeder's denial of (equivalence). As a reply to Schroeder, Chrisman (2012) suggests that non-agential ought-claims (such as 'rump roast ought to cook slowly') can sometimes carry the kind of meaning that makes it appropriate in conversational contexts of advice. In this way, he intends to cast suspicion on the linguistic intuitions at the foundation of Schroeder's claims. But it is harder to see this maneuver as successful once we have it clear in mind that the distinct meaning under consideration is a claim about what an agent has most reason to do. Consider Chrisman's (2012, 443) example:

If, for example, I've called my mother for advice about how to cook rump roast, and she asserts ['rump roast ought to cook slowly'], knowing it to be true, then surely this settles the question of what is advisable for me to do regarding the rump roast.

Chrisman is here running together two distinct matters: whether his mother's utterance settles what is advisable for him, and whether her utterance settles what is advisable for him simply by virtue of what it expresses. Certainly the former can be true without the latter being the case. Perhaps, unbeknownst to my mother, I haven't yet decided if I wish to cook the rump roast as it ought to be cooked or if I want to ruin the roast as a fun prank on a friend (after arranging for better food to be delivered). This shows that telling me truly and knowingly that 'rump roast ought to cook slowly' does not *simply by virtue of*

what it expresses settle what is advisable for me to do: its truth will sometimes determine what is advisable—a further claim about what I have most reason to do—but sometimes will not. Contrary to Chrisman’s suggestion, then, ‘rump roast ought to cook slowly’ is not an example of a non-agential ought-claim that has the kind of meaning the uniformity view (and contextualist views in general) is here challenged to accommodate. The challenge is not to explain how ought-claims can somehow play roles in settling what is advisable for one to do; the challenge is to explain how ought-claims can sometimes do this simply by virtue of what they express.⁸

Since evaluative ought-claims have only three structural elements that can vary with the context of utterance, there are only three general strategies available to the uniformity view against the normative challenge. The first strategy is identifying some possible feature of prejacent propositions that separates deliberative from non-deliberative agential ought-claims. The second strategy is identifying some similar feature of contextually salient set of worlds, and the third is identifying some similar feature of contextually salient standards of evaluation. (These are the only *pure* general strategies.) In the next section, I argue that the first two of these strategies, though recently influential, cannot succeed. My focus throughout will be on the uniformity view, but it will be clear that similar arguments afflict alternative contextualist accounts.

⁸Perhaps this serves as the kind of non question-begging account of mattering directly and indirectly for advice that Chrisman (2012, 444) asks of Schroeder.

2.4 Privileged Standards of Evaluation

Consider one natural way of developing the first strategy against the normative challenge:

(prejacent) Agential ought-claims are deliberative when the prejacent proposition makes explicit reference to the agent's agency.

The suggestion here is that an agential ought-claim such as 'Sarah ought to volunteer at the local shelter' sometimes takes 'Sarah volunteers at the local shelter' as its prejacent proposition, and sometimes takes 'Sarah *sees to it* that she volunteers at the local shelter' instead (cf. Horty 2001, 53). There is an explicit reference to Sarah's agency in the latter but not in the former, even though the surface grammar is the same in either case. As I will put it, (prejacent) claims that all agential ought-claims whose prejacent proposition makes explicit reference to the agent's agency—call these *evaluative_P* ought-claims (where 'P' stands for 'prejacent')—are deliberative agential ought-claims.

Now consider a natural way of developing the second strategy against the normative challenge:

(world) Agential ought-claims are deliberative when the contextually salient set of possible worlds only contains worlds the grammatical agent has the power to bring about.

The suggestion here is that an agential ought-claim such as 'Larry ought to file for a loan' sometimes takes a set of possible worlds all of which are within the agent's power to bring about, and sometimes not (cf. Finlay & Snedegar

2014, 113). A bit differently, an agential ought-claim such as ‘Larry ought to win the lottery’ never takes a set of possible worlds all of which are within the agent’s power to bring about, since the prejacent proposition itself refers to a world which is beyond the agent’s powers. As I will put it, (world) claims that all agential ought-claims whose contextually salient set of possible worlds only contains worlds the grammatical agent has the power to bring about—call these *evaluative_w* ought-claims (where ‘w’ stands for ‘world’)—are deliberative agential ought-claims.⁹

My argument against both these strategies aims to show that there are evaluative_p and evaluative_w ought-claims that do not express a claim about what an agent has most reason to do. Consider a case that I will call *Tavern*. Suppose Hans and Franz are having some drinks at a German tavern during the time of the second world war. Suppose they notice that a mob is forming: some local has been discovered to be Jewish and drunken men now plan to kill him. If the uniformity view is correct, then it is possible that Hans says something true when he says the following in Tavern:

(join) Franz ought to join the mob.

⁹I am ignoring some differences between Finlay & Snedegar’s (2014) suggestion and (world). First, while the orthodox variety takes ought-claims to be about what’s contextually best, Finlay & Snedegar are *contrastivists* who take ought-claims to be about what’s contextually better (more exactly, they are about what is best *in the contrast set*). They leave open, however, whether the contrast class is realized syntactically or semantically (cf. fn. 19). Second, Finlay & Snedegar take themselves to be offering merely a necessary condition, whereas (world) offers a sufficient condition. This is because they are attempting to meet Schroeder’s original challenge, and not the normative challenge presently under consideration. What follows is *not*, therefore, a argument against their reply to Schroeder, but rather an argument against the use of this reply as a resource against the normative challenge.

If the uniformity view is correct, recall, what (join) expresses is merely that ‘Franz joins the mob’ is true in all of the best ranked worlds given a contextually salient set of worlds and a contextually salient standard of evaluation. But notice that there is a possible standard of evaluation \mathbb{S}_1 that takes anti-Semitism as its guiding value. Given any set of possible worlds, \mathbb{S}_1 ranks them according to how much suffering is experienced by Jewish people in those worlds. If the conversational context provides \mathbb{S}_1 as the relevant standard of evaluation, and provides a set of worlds \mathbb{W}_1 such that ‘Franz joins the mob’ is true in all the worlds that are ranked highest by \mathbb{S}_1 , then (join) is true in Tavern. Call the conversational context that would provide \mathbb{S}_1 and \mathbb{W}_1 to an utterance of (join) in Tavern the *Nazi context*. There is nothing impossible about Hans and Franz being in the Nazi context, so nothing impossible about (join) being true in Tavern.

Similarly, if the uniformity view is correct, then it is possible that Hans says something true when he says the following in Tavern as well:

(no-join) Franz ought not to join the mob.

What (no-join) expresses, after all, is that ‘Franz does not join the mob’ is true in all of the best ranked worlds given a contextually salient set of worlds and a contextually salient standard of evaluation. But notice that there is a possible standard of evaluation \mathbb{S}_2 that takes humanitarianism as its guiding value. Given any set of possible worlds, \mathbb{S}_2 ranks them according to how little harm is done to others in those worlds. If the conversational context provides \mathbb{S}_2 as the relevant standard of evaluation, and provides a set of worlds \mathbb{W}_2 such

that ‘Franz does not join the mob’ is true in all the worlds that are ranked highest by \mathbb{S}_2 , then (no-join) is true in Tavern. Call the conversational context that would provide \mathbb{S}_2 and \mathbb{W}_2 in Tavern the *Humanitarian context*. There is nothing impossible about Hans and Franz being in the Humanitarian context either, so nothing impossible about (no-join) being true in Tavern.¹⁰

Though it seems perfectly possible for Hans and Franz to be in either the Nazi or the Humanitarian contexts, it seems nonetheless false that what Franz has most reason to do in Tavern is different in each case. In both the Nazi and the Humanitarian contexts, that is, Franz has most reason to not join the mob. This is an intuitively plausible judgment, I think, but also a judgment that is buttressed by most accounts of reasons. Whatever we say about the semantics of ought-claims, reasons—the building blocks of that “specifically practical relation holding between agents and actions that is relevant to ethics” (Chrisman 2012, 436)—are simply not context-sensitive in this way. So we can draw the following lesson from reflection on Tavern:

Lesson₁: Not every standard of evaluation has the power to generate reasons.

¹⁰Ross (2010, 308-9) accepts and makes good use of a principle claiming that “there is no proposition, p , such that it ought to be the case that (p and not- p).” This is not in tension with my claims here about (join) and (no-join) both being possibly true in Tavern. My claim is rather that, if the uniformity view is correct, then there can be propositions p such that in a certain conversational context it ought to be the case that p and in another conversational context it ought to be the case that not- p . It does not follow from this that there can be a conversational context where it ought to be the case that p and not- p .

The fact that a certain action of mine scores highly with respect to a certain standard \mathbb{S} , or the fact that a certain world is ranked highest given \mathbb{S} , is not sufficient to make it the case that I have reasons to perform that action, or to bring about the actions that I perform in that highest ranked world.¹¹ Not every standard, we can say, is *privileged* in this way. In Tavern, it seems clear that \mathbb{S}_1 is not a privileged standard: the fact that joining the mob scores highly with respect to \mathbb{S}_1 in no way gives Franz reasons to act in that way, even if Hans and Franz are in a conversational context that provides \mathbb{S}_1 as the relevant standard of evaluation and \mathbb{W}_1 as the relevant set of worlds.¹²

The first lesson from Tavern is a general lesson about standards and their powers to generate reasons. But we can also draw a lesson from Tavern that is more specifically about evaluative_P and evaluative_W ought-claims:

Lesson₂: Not all evaluative_P ought-claims and not all evaluative_W ought-claims express a claim about what an agent has most reason to do.

¹¹Human-independent accounts most clearly think of reasons non-contextually in this way (cf. Skorupski 2010, Enoch 2011, Dworkin 2011, Parfit 2011, and Scanlon 2014). Human-dependent accounts of almost all sorts think of them in this way as well (cf. McDowell 1995, Scanlon 1998, and Gert 2004). Henning (2014), however, argues that normative reasons are, in fact, context-sensitive. But what he has in mind is the view that which reasons there are varies with “the *information* available from a contextually relevant point of view” (601, my italics). Since there is no variation in the information that Hans and Franz have available in the Nazi and the Humanitarian contexts, there is in fact no shift in the kind of context that Henning thinks can affect the existence of reasons. (The same point addresses an extension of the assessor-relativism we find in MacFarlane (2009) to deontic modals, where truth is relative to the information that is available in the context of assessment.) The exception, perhaps, are so-called desire-based accounts (cf. Williams 1981b and Schroeder 2007). I will address these exceptions very briefly at the end of this section.

¹²Strictly speaking, what I am calling standards of evaluation rank *possible worlds* with respect to a certain value, and not *actions*. When I speak of an action of mine being ranked highly by a certain standard, I mean this derivatively: I perform that action in the worlds ranked highest by that standard.

Suppose (join), in the Nazi context, takes ‘Franz *sees to it* that he joins the mob’ as its prejacent proposition. Here’s what results: given the uniformity view, (join) is true; given the normative facts, Franz does not have most reason to join the mob. So (join) is here a true evaluative_P ought-claim that does not express a claim about what Franz has most reason to do. Now suppose (join), in the Nazi context, takes a set of possible worlds all of which are within Franz’s power to bring about. Here’s what results: given the uniformity view, (join) is once again true; given the normative facts, Franz still does not have most reason to join the mob. So (join) is here a true evaluative_W ought-claim that does not express a claim about what Franz has most reason to do.

The second lesson from Tavern, however, is problematic for (prejacent) and (world). If evaluative_P and evaluative_W ought-claims sometimes fail to express a claim about what an agent has most reason to do, then sometimes they are not appropriate in conversational contexts of advice in the characteristic ways suggested by Schroeder. More importantly, they are then not the kinds of claims that pick out that practical relation that matters so much to normative philosophy. We can state this argument more precisely:

The Privileged Standards Argument

1. Not every standard of evaluation is privileged. (Lesson 1)
2. If not every standard of evaluation is privileged, then not all evaluative_P ought-claims and not all evaluative_W ought-claims express a claim about what an agent has most reason to do. (Lesson 2)

3. If not all evaluative_P ought-claims and not all evaluative_W ought-claims express a claim about what an agent has most reason to do, then not all evaluative_P ought-claims and not all evaluative_W ought-claims are deliberative ought-claims. (by The Normative Assumption)

C. So not all evaluative_P ought-claims and not all evaluative_W ought-claims are deliberative ought-claims.

If the Privileged Standards Argument is sound, then it is clear that further refinements of (prejacent) and (world)—or their combination, or further deployments of the first and second strategies in general—will fail as well. The problem with these strategies is not that they identify the wrong feature of the prejacent proposition or of the relevant set of worlds, but rather that features of those parameters are not enough to determine whether an evaluative ought-claim is deliberative or not.¹³

Let me put aside one possible maneuver against the Privileged Standards Argument. My defense of premise 1 depends on the following being a normative fact:

(**fact**) Franz does not have most reason to join the mob in the Nazi context of Tavern.

¹³Another strategy—though one that departs slightly from the orthodox variety of the uniformity view—fails for the same reason: simply adding a further contextual parameter that explicitly identifies an *agent* for the evaluative ought-claim (cf. Broome 1999 and Wedgwood 2007) will similarly fail to distinguish between those with privileged and unprivileged standards, and so fail to distinguish between those that express a claim about what that agent has most reason to do and those that don't.

Perhaps some will think that (fact) is not a fact. I have indeed assumed that the ‘has most reason to’ relation is an *external* relation in the following sense: it holds between a certain person and a certain action irrespective of that person’s actual beliefs and desires. I’m not alone in this assumption. Perhaps this external relation holds because the person would be motivated to perform that action if she had an ideal moral upbringing (cf. McDowell 1995, 74); perhaps it so holds because the person would be motivated to perform that action if her “dispositions of evaluation” were not defective (cf. Scanlon 1998, 371); perhaps it so holds for reasons completely independent of that person’s psychology, even in its idealized forms (cf. Parfit 1997, 129; Scanlon 2014, 44). On any of these externalist accounts, (fact) is a fact. But according to desire-based accounts of reasons, the ‘has most reason to’ relation is an *internal* relation in the following sense: it holds between a certain person and a certain action by virtue of the actual beliefs and desires of that person (cf. Williams 1981b and Schroeder 2007). If these accounts are correct, then perhaps (fact) is not a fact. If joining the mob is what most advances Franz’ terrible desires in the Nazi context, then Franz, on these accounts, has most reason to join the mob.¹⁴

I find this kind of internalism hard to accept, but I do not wish to argue against it here. I take myself to be addressing my fellow externalists instead. What the Privileged Standards Argument shows, then, is that (prejacent) and

¹⁴Notice the similarities between Tavern and Bernard Williams’ (1995, 39) discussion of a man who is cruel to his wife. According to Williams, we can make very many negative evaluations of this man: “that he is ungrateful, inconsiderate, hard, sexist, nasty, selfish, brutal...” But, according to Williams, we cannot say that he has a *reason* not to be cruel.

(world) should be rejected by anyone who pairs the Normative Assumption with externalism about reasons.¹⁵

2.5 Super Standards of Evaluation

Recall the normative challenge: identifying the feature of evaluative ought-claims that distinguishes those that are deliberative from those that are not—those that express a claim about what an agent has most reason to do and those that don't. In the previous section, we've seen that, given plausible assumptions about the nature of reasons, variations in the prejacent proposition or in the contextually salient set of worlds cannot by themselves account for this distinction. Meeting the normative challenge, it turns out, requires paying close attention to the (so far ignored) standards of evaluation that figure into ought-claims.

Consider one natural suggestion:

(standard) Evaluative ought-claims are deliberative when the contextually salient standard of evaluation is privileged.

¹⁵Internalists about reasons, however, should not dismiss the Privileged Standards Argument too quickly. Recall that the uniformity view is supposed to capture all uses of the English 'ought'. One such use, notice, indicates all and only our *epistemic* reasons for believing a certain proposition. When we say 'S ought to believe that p', we are typically not claiming that S has most reason of all kinds to believing that p, but rather that S has most reason of a certain specific kind (the epistemic ones) to believe that p. But when we restrict ourselves to epistemic ought-claims in this way, externalism about reasons becomes even harder to reject. Few are willing to accept that which epistemic reasons we have for believing that p depends on our actual beliefs and desires. So while pairing the uniformity view with internalism about *practical* reasons may help avoid the Privileged Standards Argument (by way of denying that (fact) is a fact), it seems much less plausible to suppose that such internalism would help the uniformity view against a version of the same argument that was constructed with claims about epistemic reasons.

According to (standard), (no-join) is deliberative since \mathbb{S}_2 is privileged: it has the power to generate reasons. Similarly, (join) is not deliberative since \mathbb{S}_1 is not privileged: it does not have the power to generate reasons. As I will put it, (standard) claims that all agential ought-claims whose contextually salient standard of evaluation is privileged—call these *evaluative_s* ought-claims (where ‘s’ stands for ‘standard’)—are deliberative agential ought-claims.

This natural suggestion, however, will not do. To claim that a certain action is what Franz has most reason to do, after all, is to make a comparative claim about the balance of reasons, as opposed to a claim about what is issued by just one source of reasons. We can here welcome a bit of familiar terminology: to claim that not joining the mob is what Franz has most reason to do is to claim that not joining the mob is what Franz has *all-things-considered* reason to do, and not just to claim that it is what Franz has a *pro tanto* reason to do. The fact that some action ϕ scores highly with respect to a privileged standard always gives you a pro tanto reason to perform it, but it does not always give you an all-things-considered reason to perform it as well. (standard) pays no attention to this important distinction. As a consequence, some *evaluative_s* ought-claims simply do not express a claim about what an agent has most reason to do.

It pays to illustrate this point. Recall the Humanitarian context of Tavern. If the uniformity view is correct, then this is a case where (no-join) is true: ‘Franz does not join the mob’ is true in all the best possible worlds given \mathbb{W}_2 and \mathbb{S}_2 . If (standard) is correct, moreover, then this is also a case where (no-join) is a claim about what Franz has most reason to do: \mathbb{S}_2 is a privileged standard,

and all agential ought-claims whose contextually salient standard of evaluation is privileged are deliberative. Now consider three further assumptions. First:

A₁: There are privileged standards beside S_2 .

We know that there are standards of evaluation besides S_2 —terrible ones like S_1 , for example. But perhaps there are *privileged* standards besides S_2 as well. Perhaps there is a privileged standard of evaluation S_3 that takes fidelity (the keeping of promises) as its guiding value. In that case, we have reasons to perform an action that scores highly with respect to S_3 . Perhaps there is also a privileged standard of evaluation S_4 that takes beneficence (the maximization of what’s good) as its guiding value. In that case, we have reasons to perform an action that scores highly with respect to S_4 . Perhaps there are even more privileged standards.¹⁶ Second:

A₂: ‘Franz does not join the mob’ is false in all the best contextually salient possible worlds given S_3 and S_4 .

Perhaps Franz had promised the men now forming a mob that he was going to participate in whatever affair they decided to conduct after a few drinks. (This is an odd promise, alright, but hardly one that has never been made.) Perhaps Franz’s refusal to join the mob would cause the belligerent Nazis to direct their anger at him, at Hans, and at their families, instead of having it directed at just one person. In that case, not joining the mob does not at all

¹⁶It should be clear that I am here drawing on Ross’s (1930, esp. ch. 2) framework of prima facie duties.

advance fidelity and beneficence values. So given (A₁) and (A₂), Franz has at least pro tanto reasons to join the mob. Finally:

A₃: Ranking highly with respect to S₃ and S₄ outweighs ranking highly with respect to just S₂.

Perhaps the fact that an action ϕ advances promise-keeping and the maximization of the good at the cost of humanitarianism gives one reasons to ϕ that are *stronger* than those given by the fact that not ϕ -ing advances humanitarianism at the cost of promise-keeping and the maximization of the good. As I see it, there is nothing impossible about these three assumptions being true in the Humanitarian context of Tavern.¹⁷ But if they are true, then the Humanitarian context of Tavern is now a case where ‘Franz does not join the mob’ is true in all the best possible worlds given *one* privileged standard—the one that happens to be contextually salient in Hans’ utterance of (no-join)—and yet (no-join), though true, is not a claim about what Franz has most reason to do. So (standard) is false: not every evaluative_S ought-claim is deliberative.¹⁸

There is, however, nothing impossible about standards of evaluation that always create all-things-considered reasons. Call such standards *super standards*. Those who think that there is just one source of reasons—just one privileged standard—will already think that such standard is a super standard. Perhaps

¹⁷At any rate, there is nothing impossible about some version of these assumptions being true. Ross (1930, 19-22) himself would disagree with how I have stipulated the weights of S₂, S₃, and S₄, but the argument can be modified to suit his or anyone’s views on which standards are privileged and which standards outweigh which.

¹⁸Ross (2010, 317) seems to think that these considerations create a serious worry for the uniformity view. As we will see presently, this verdict is premature.

consequentialists are correct, for example, in taking \mathbb{S}_4 as the uniquely privileged standard. If scoring highly with respect to a standard taking beneficence as its guiding value is the only thing that can give us a reason to perform a certain action ϕ , then that reason will always be an all-things-considered reason to ϕ . But those who think that there are a variety of privileged standards should welcome the existence of super standards just as well. While some standards rank possible worlds with respect to their advancement of just one value—beneficence, for example—some standards rank possible worlds with respect to their advancement of a variety of values—humanitarianism, beneficence, fidelity, and so on. While privileged standards measuring how well actions advance each of these values indicate our separate pro tanto reasons for action, privileged standards measuring how alternative actions score with respect to all of them, and weighing those scores correctly against each other, indicate our all-things-considered reasons for action. However one thinks about the metaphysics of value—barring axiological nihilism—the existence of super standards of evaluation seems unproblematic.¹⁹

This suggests a second natural way of developing the third strategy:

(super standard) Evaluative ought-claims are deliberative when the contextually salient standard of evaluation is a super standard.

¹⁹I am staying silent on the axiological question: what explains the aggregative and non-aggregative normative relations of relative importance between the relevant values that super standards of evaluation use in their ranking of salient possible worlds? See Chang (2004, 18) for a defense of the interesting claim that these normative relations are determined by the unity of a (sometimes nameless) comprehensive value.

According to (super standard), (no-join) is deliberative since it takes not S_2 but rather SS_2 : a standard of evaluation that always generates all-things-considered reasons (either by being the only privileged standard, or by correctly ranking worlds with respect to how well they score with respect to all privileged standards). As I will put it, (super-standard) claims that all agential ought-claims whose contextually salient standard of evaluation is a super standard—call these *evaluative_{SS}* ought-claims (where ‘ss’ stands for ‘super standard’)—are deliberative agential ought-claims.

If the uniformity view is correct in taking all ought-claims as evaluative ought-claims, and if some of these ought-claims are sometimes appropriate in conversational contexts of advice in the characteristic ways suggested by Schroeder, and if those ought-claims are precisely the ones that normative philosophy is interested in—and so ones that express a claim about what an agent has most reason to do—then I suggest that (super standard) must be true.²⁰ If you come to me for advice, what is appropriate from me is a claim about what you have most reason to do, a claim about what you have all-things-considered reason to do. Only evaluative ought-claims contextually resolved with super standards can truly or falsely express those kinds of claims. Once again, this is a result that any contextualist account of the English ‘ought’ must accommodate.

²⁰Some may complain that I am ignoring the grain of truth in (prejacent) and (world) by not even taking them as stating necessary conditions. But it is not clear to me that they do. Either way, my critical point is that succeeding against the normative challenge requires close attention to the standards of evaluation that figure into ought-claims. (super standard) brings this point home in bold terms. The same point, however, can be used for the development of a position that incorporates either (prejacent), or (world), or both.

Notice how (super standard) mirrors the view of the structure of overall moral obligation that Zimmerman (2008, 5) has called the “ought-value-stuff framework”. According to this view, “any substantive theory of obligation can be cast as one according to which what one ought to do is ranked higher than any alternative.” This may be right, provided we assume that super standards are responsible for these rankings.²¹ Of course, the semantic structure of a claim about what an agent has most reason to do—about the act which is obligatory for an agent, on Zimmerman’s terms—may well differ from the metaphysical structure of what *makes* an act ϕ the act which that agent has most reason to do. But there is no such difference if (super-standard) is true: what makes ϕ -ing the act I have most reason to do is the fact that ϕ -ing is well-ranked by a super standard, and the semantic structure of my claim expressing that fact similarly indicates that ϕ -ing is well-ranked by a super standard. No doubt there is something attractive about this resulting parallel structure.

2.6 The Contextual Standards Dilemma

As presented in section 1, the normative challenge is met: SUPER-STANDARD gives us a principled account of how some but not all evaluative ought-claims can express a claim about what an agent has most reason to do. This is progress, since previous discussions of the uniformity view did not produce the resources required to succeed against that challenge. Nonetheless, accepting

²¹See also Feldman (1986) for the view that what we ought to do is what we do in the best possible accessible worlds.

SUPER-STANDARD forces contextualist views in general into a new dilemma. This is not a fatal dilemma, but it nonetheless reveals some of their hidden costs. In what follows, the uniformity view is once again my focus.

If the uniformity view is correct, then we have seen that the following is true in the Nazi context in Tavern:

(join) Franz ought to join the mob.

This is because the Nazi context provides \mathbb{S}_1 as the relevant standard of evaluation, and provides a set of worlds \mathbb{W}_1 such that ‘Franz joins the mob’ is true in all the worlds that are ranked highest by \mathbb{S}_1 . An utterance of (join) merely expresses that this is the case. But now consider a variation on the original Tavern case. Suppose that Hans and Franz are hopelessly misguided conscientious Nazis: they are sincerely attempting to figure out what Franz has most reason to do, but unfortunately endorse the terrible values embodied in \mathbb{S}_1 . Suppose that, after some deliberation, Hans forms the explicit belief that joining the mob is what Franz has most reason to do and proceeds to utter (join). Call this the *conscientious variation* of Tavern. What should the uniformity view say about this conscientious variation? I want to consider two alternative accounts. One of these accounts, I think, is very problematic. The other account requires a substantive commitment that shouldn’t go unnoticed.

Here is the first account. Since Hans and Franz explicitly endorse the terrible values embodied in \mathbb{S}_1 , the conscientious variation of Tavern is still a case where Hans and Franz are in the Nazi context. By explicitly considering which action will further what *they believe* is valuable, Hans and Franz have thereby

made S_1 the contextually salient standard of evaluation in their conversational background. For many who accept the uniformity view, this is in fact the most natural thing to say. According to Finlay (2009, 328), for example, “every normative ‘ought’ is simply the ordinary modal ‘ought’ under an ‘in order that...’ modifier.” More exactly, Finlay’s view is that the truth of every ought-claim, including those used categorically, is partly determined by the contextually salient ends that the speaker intends to promote.²² In the conscientious variation of Tavern, it is hard to see whose ends would be more relevant than Hans’ and Franz’, and hard to see how those terrible ends would not determine the standard of evaluation that is relevant for Hans’ ought-claim.

But accepting this first account forces the uniformity view to say that (join) is true in the conscientious variation of Tavern. After all, (join) is true in the Nazi Context. And since joining the mob is not what Franz has most reason to do—this is (fact)—the uniformity view cannot as well say that (join) is a claim about what Franz has most reason to do. So, in accepting this first account of the conscientious variation of Tavern, the uniformity view is forced to say that Hans has here tried but failed to express a claim about what Franz has most reason to do, expressing instead a true claim that does not express what he intended. Even though Hans considered which values were relevant to his situation, reflected on their weights and bearings, and conscientiously

²²There are no truly categorical ought-claims on Finlay’s view, merely categorical *uses* of normal end-relational ought-claims. Categorical use is characterized as use where one “presupposes” the relevant ends (Finlay 2009, 331). When the ends that are presupposed are not shared by the audience, moreover, a speaker’s categorical use of an ought-claim “expresses the demand that his audience share his concern for the relevant end, and consequently for the behavior at issue” (Finlay 2009, 333).

formed the belief that joining the mob is what Franz had most reason to do, his attempt to express that belief with an ought-claim was frustrated by the nature of those values, his context, and his choice of words.

This, I think, is a pretty implausible account of what happens in the conscientious variation of Tavern. What happened, rather, is that Hans said something that was *false*: he claimed, with his ought-claim, that joining the mob is what Franz had most reason to do, and he was wrong. This is so despite the fact that joining the mob scores highly with respect to all the contextually salient standards of evaluation which Hans misguidedly endorses (the ones relevant in the Nazi context). But the uniformity view cannot buttress this judgment if Hans is truly in the Nazi context. So if this first account of the conscientious variation of Tavern is correct, then the uniformity view remains problematic, despite having already escaped the normative challenge.

Here is the second account. What seems to generate a problem in the conscientious variation of Tavern is the assumption that the contextually salient standard of evaluation is S_1 —which is admittedly not a *privileged* standard. This is a plausible assumption, of course, since Hans himself—being a hopelessly misguided Nazi—values such terrible things as the suffering of Jewish people. But this assumption is nonetheless mistaken. By ostensibly considering the balance of relevant values, and by explicitly considering what Franz *has most reason to do*, Hans has thereby removed himself from the Nazi context and placed himself into a context providing a super standard of evaluation instead of S_1 . Despite Hans' terrible values, and despite his focus on those very values, the contextually salient standard of evaluation playing a

role in his utterance of ‘Franz ought to join the mob’ is a standard that does not rank as highest the worlds in which Franz joins the mob.

This account of the conscientious variation avoids the problematic feature just discussed. It turns out that the uniformity view can, after all, allow for Hans’ false claim about what Franz has most reason to do. According to SUPER-STANDARD, recall, any ought-claim contextually resolved with a super standard is a claim about what an agent has most reason to do. According to the second account of the conscientious variation of Tavern, conscientiously attempting to make a claim about what someone has most reason to do thereby places one in a conversational context providing a super standard for one’s ought-claims. So Hans’ hopelessly misguided conscientious claim that Franz ought to join the mob is a false claim about what Franz has most reason to do. Hans correctly noticed that Franz joins the mob in the worlds ranked highest by the values he endorses; nonetheless, worlds in which Franz joins the mob rank quite low given the standard of evaluation contextually provided for his ought-claim.

This, I think, is not an implausible account of what happens in the conscientious variation of Tavern. But it requires the uniformity view to take up certain commitments about the pragmatic mechanisms that determine the contextually relevant standard of evaluation. Consider Robert Stalnaker (2014, 24) on contextual relevance in general:

If communication is to be successful the contextual information on which the content of a speech act depends must be information that is *available* to the addressee.

Call this the *Stalnaker Constraint*. If the content of an ought-claim depends on a standard of evaluation that is provided by context, as the uniformity view suggests, and if the content of Hans' particular ought-claim is successfully communicated in the conscientious variation of Tavern, as it is plausible to assume, then the Stalnaker Constraint tells us that the standard of evaluation employed by Hans' ought-claim is one that is in some sense available to Hans and Franz. But how should we understand this notion of availability?

One natural position would be that, after explicit consideration and endorsement of a certain set of terrible values, the standard of evaluation that is available to Hans and Franz is S_1 and not some super standard taking values they neither endorse nor have considered. But this natural position is incompatible with the plausible assumption that Hans makes a false claim about what Franz has most reason to do in the conscientious variation of Tavern, and so incompatible with the second account under consideration. This natural notion of availability, more exactly, would make it the case that Hans and Franz are once again in the Nazi Context. This reveals that accepting the second account of the conscientious variation of Tavern—where Hans and Franz are not in the Nazi Context—requires accepting a weaker and perhaps less natural notion of contextual availability. This is a non-trivial commitment, but not obviously a commitment that leads to serious trouble.

We can generalize my claims in this section and state them in the form of a dilemma. Call it *the contextual standards dilemma*:

The Contextual Standards Dilemma: If contextual accounts of the English 'ought' are correct, then misguided conscientious people either

(a) never make claims about what one has most reason to do, or (b) are placed in a conversational context providing a super standard of evaluation by virtue of their conscientiousness.

(Horn #1) If (a), then contextualist accounts seem exposed to counterexamples.

(Horn #2) If (b), then contextualist accounts depend on a yet to be articulated notion of *weak contextual availability*.

The context dilemma follows from the fact that only evaluatives_{SS} ought-claims can express a claim about what an agent has most reason to do. That is, it follows from my argument in sections 2 and 3 that the uniformity view must accept SUPER-STANDARD in order to escape the normative challenge. I am not suggesting that either horn is fatal, though I think embracing (Horn #1) is more problematic than embracing (Horn #2). Exactly how problematic it is to embrace (Horn #2), of course, depends on further examination of a carefully articulated notion of weak contextual availability.²³

²³Notice how this dilemma differs from the kind of contextual challenge that Dowell (2013, 150-2) has tried to meet. That challenge has to do with the relation between contextual salience and *information that is not available*. (Think here of how the challenge originates in Frank Jackson's (1991) consequentialist discussion of Dr. Jill.) The challenge I am discussing here, however, is instead about the relation between contextual salience and *normative standards*. This is a challenge that arises, importantly, even for cases where there is no bit of relevant information that is lacking. (Think here of an omniscient but devious Dr. Jill.) Relatedly, notice that Finlay's (2014, 231-236) discussion of "pragmatic disagreements from different ends" is a discussion about *interpersonal disagreement*. His discussion, therefore, is not directly relevant to a dilemma about the pragmatic mechanisms required for placing speakers in contexts that match our intuitive evaluation of the truth-value of their claims.

2.7 Conclusion

Let me summarize my argument. Normative philosophy gives pride of place to the ‘has most reason to’ relation. In its discussion of this central relation, normative philosophy typically employs the English ‘ought’. Yet it is remarkably difficult to explain how ‘ought’ could at times express this relation and at times not. I have here called this explanatory difficulty the normative challenge and I have claimed that it is a challenge to any contextualist account of the semantics of the English ‘ought’. As I see it, moreover, succeeding against this challenge recalls going beyond recent deployments of one widely accepted contextualist account—the uniformity view. My diagnosis is that not enough attention has been paid to the contextually relevant standards of evaluation that figure into ought-claims. Accordingly, I have argued that we can succeed against the normative challenge by examining the kinds of standards that can issue in a claim about what someone has most reason to do. But succeeding against the normative challenge forces the uniformity view into what I have called the contextual standards dilemma. I do not think that this is fatal to the uniformity view, but I take it to reveal some of its hidden costs and to point the way for further exploration.

CHAPTER 3

ON EVADING THE DEONTIC PUZZLE

3.1 Introduction

Consider the following three claims:

Epistemic Involuntarism (I): It is not the case that believing that p is within S 's voluntary control.

Epistemic Deontology (D): Some claims of the form 'S ought to believe that p ' are true.

The Ought-Implies-Can Principle (OIC): If ϕ -ing is not within S 's voluntary control, then it is not the case that the claim 'S ought to ϕ ' is true.

Each is independently plausible. But if (I) is true, then the antecedent of (OIC) is always true when ϕ is believing that p ; and if (D) is true, then the consequent of (OIC) sometimes is false when ϕ is believing that p . Yet the antecedent of (OIC) cannot always be true and its consequent sometimes false for the same ϕ . So while (OIC), (D), and (I) all seem true, at least one them must be false. Call this *The Deontic Puzzle*.

Some have attempted to *dispel* the deontic puzzle by denying either (I), (D), or (OIC).¹ Yet none of these options is uncontroversial or theoretically inexpensive: each denial requires silencing strong intuitions, as well as modifying central aspects of one's epistemology. Others have attempted to *evade* the deontic puzzle by suggesting that, properly understood, (I), (D), and (OIC) are not in tension in the first place. This alternative has been recently influential, in fact, with advocates as diverse as Richard Feldman (2000, 2008), Hilary Kornblith (2001), and Matthew Chrisman (2008, 2012). Evasion has seemed to many to be the least controversial and the least costly of the available maneuvers against the puzzle: it does not require silencing strong intuitions (or not as many, or not as strong), and it does not require modifying central aspects of one's epistemology (or not as many, or not as central). Nonetheless, I will here argue that evading the deontic puzzle is much more costly than advertised. More exactly, I will argue that evading the deontic puzzle comes at the cost of normative reasons for belief. The evasive strategy has seemed attractive to many in part because the severity of its consequences has not been fully appreciated.

Here is how I proceed. In the first section, I distinguish two aspects of our normative experience that are too often obscured by generic talk of normativity: evaluations and prescriptions. In the second section, I make use of this distinction to shed light on how the evasive strategy disambiguates (D) and

¹See Steup (2000, 2008, 2012) and Ryan (2003) for denials of (I), Alston (1988, 2005) and Papineau (2014) for denials of (D), and Saka (2000) and Graham (2011) for denials of (OIC).

(OIC): the former as a claim about evaluative normativity and the latter as a claim about prescriptive normativity. In the third section, I discuss and criticize Rik Peels' (2014) attempt to show that the evasive strategy comes at the cost of epistemic responsibility. Here I claim that Peels fails to show that this is truly a consequence of the evasive strategy. In the fourth section, I identify what I take to be the real cost of the evasive strategy: normative reasons for belief. My argument here turns on the claim that normative reasons are the building-blocks of prescriptive normativity. In the fifth and final section, I resist several attempts to brush aside this consequence as not too uncontroversial and not too costly after all.

3.2 Evaluations and Prescriptions

Talk of normativity is often obscure. Sometimes such talk is about certain *speech acts*; sometimes it is about certain *judgments* expressed by certain speech acts; sometimes it is about certain *concepts* deployed in certain judgments; sometimes it is about certain *properties* captured by certain concepts. Even worse, it is often unclear what philosophers are indicating when they claim that a speech act, a judgment, a concept, or a property is *normative* in the first place. I want to avoid this kind of obscurity, so I begin by discussing what I mean by talk of normativity in this paper.

At bottom, talk of normativity is talk of two familiar kinds of experiences. We experience what I will call *evaluative-normativity* when we judge a state of affairs by some standard that we endorse (tacitly or otherwise). When I say sincerely of a guest that she is extremely polite, for example, I am evaluating

her behavior with respect to some standard of etiquette that I find appropriate. When I say sincerely of an actress that she does not deserve to win the award, for another example, I am evaluating her acting with respect to some artistic standard that I find appropriate as well. In one case, the evaluation is positive; in the other, the evaluation is negative. In both cases, I take it, the evaluation is quite familiar.

Ernest Sosa (2009, 70) has in mind precisely these experiences of evaluative-normativity when he says:

We humans are zestfully judgmental across the gamut of our experience: in art, literature, science, politics, sports, food, wine, and even coffee; and so on, across many other domains. We love to evaluate even when no practical interest is in play. We judge performances, whether artistic or athletic; grade products of craft or ingenuity; evaluate attitudes, emotions, institutions, and much more.

So when we say that some speech act, judgment, concept, or property is normative, sometimes we simply mean to indicate their place within these familiar experiences of evaluative-normativity.

A bit differently, we experience what I will call *prescriptive-normativity* whenever we judge that some action or attitude is required of someone or to be done. When I say sincerely to a friend that he ought to be faithful to his partner, for example, I typically don't take myself as merely saying that cheating scores poorly with respect to a standard that I happen to endorse. Some such claim seems too parochial to capture what I am trying to convey. What is characteristic of our experiences of prescriptive-normativity—and what distinguishes them from our experiences of evaluative-normativity—is a perception of certain actions as “having a grip” on us (cf. Korsgaard 1996,

44-6) or being “demanded” of us (cf. Street 2012, 44); a perception of “being called on” to behave in a certain way (cf. Thomson 2008, 207).²

Stephen Grimm (2009, 253-4) is alluding to these experiences of prescriptive-normativity when he says:

When we judge a belief to be unjustified or irrational, we seem to be doing more than just evaluating (in this case, in a negative way) the skill or virtuosity of the believer’s performance... To judge someone’s belief to be unjustified or irrational is thus to judge that the person’s attitude towards the content of the belief should be reconsidered, in some apparently binding sense of ‘should.’

So when we say that some speech act, or judgment, or concept, or property is normative, sometimes we mean to indicate something more than merely the fact that they belong to our familiar experiences of evaluative-normativity. Sometimes we mean to indicate their place within our equally familiar experiences of prescriptive-normativity instead.

So talk of normativity is at times talk of our experiences of evaluative-normativity and at times talk of our experiences of prescriptive-normativity. Sometimes we simply mean to indicate our personal endorsement of some standard (tacit or otherwise); sometimes we also mean to indicate the perception of something a bit more binding. A theory of normativity, of course, aims at *explaining* what makes it appropriate for some speech act, or judgment, or concept, or property to play a role in these familiar experiences. But simply

²It is this feature of morality—the fact that we experience its claims as prescriptive-normative—that Kant (1788, 86) eulogized (“Duty! Thou sublime and mighty name that dost embrace nothing charming or insinuating but requirest submission”) and that Mackie (1977, 38) took as “utterly different from anything else in the universe,” and as ultimately reflecting some of our mistaken beliefs about a purely natural world.

distinguishing between these two kinds of normativity will be sufficient for now.³

Passages such as the above by Grimm, however, may seem to suggest that the English ‘ought’ is only appropriate in the context of prescriptive-normativity. The passage may suggest, that is, that ‘ought’ always serves to express that something is required or to be done. But this is not the case.

No doubt there is a use of ‘ought’ closely connected to our experiences of prescriptive-normativity. ‘Ought’, in what we can call *prescriptive-normative ought-claims*, expresses that there is an authoritative relation of normative requirement or obligation holding between a certain agent S and a certain action ϕ . This is likely what we mean when we tell a tempted friend that he ought to be faithful to his partner.⁴ But there are uses of the English ‘ought’ closely connected to our experiences of evaluative-normativity as well. ‘Ought’, in what we can call *evaluative-normative ought-claims*, merely expresses the ranking of a state of affairs with respect to a standard that is being endorsed (tacitly or otherwise). This is likely what we mean when we tell a visiting friend that she ought to take the second exit on the left. The English ‘ought’ is thus appropriate in contexts of prescriptive-normativity as well as in contexts of

³See Nolfi (2014, 98) for the claim that our experience of prescriptive-normativity is connected to the giving of advice, guidance, or direction. See Eklund (2013) for discussion of alternative explanations of what makes speech acts, concepts, or properties appropriate in normative contexts.

⁴The source of this normative authority is a matter of debate. Since the term ‘obligation’ is so often and so naturally associated with *moral* obligation, I will here give preference to the more neutral term ‘normative requirement’. I will elide the ‘authoritative’ qualifier throughout.

evaluative-normativity, depending on what it is used, at that time, to express: a relation of normative requirement between agents and actions, or merely a ranking-relation between standards and states of affairs.⁵

This feature of the English ‘ought’ is quite important. Since evaluative-normative ought-claims and prescriptive-normative ought-claims express different kinds of relations, they can be expected to have different truth conditions. In particular, it is plausible that only evaluative-normative ought-claims can be true even when the grammatical subject of the sentence is not capable of bringing about the relevant state of affairs (cf. Chrisman 2008, 56-7). Consider the claims:

(C₁) The clock ought to strike on the quarter hour.

(C₂) Wealthy Americans ought to donate to Oxfam.

Since clocks have no agency, and since striking on the quarter hour is not an action which clocks can choose to perform or not, (C₁) would be false or incoherent if it claimed that there is a relation of normative requirement holding between a certain agent (clocks) and a certain action (striking on the quarter hour). Yet (C₂) seems coherent and even true when taken in that way. Wealthy Americans are agents, after all, and we typically think of them as perfectly capable of choosing from alternative actions. But the surface-grammar similarity between (C₁) and (C₂) does not reflect a normative similarity. While (C₁) is simply an evaluative-normative claim, (C₂) is a prescriptive-normative

⁵See Schroeder (2011) for the claim that this distinction reflects a distinction in the underlying semantic structure of ought-claims; see Chrisman (2012a) for resistance.

claim; while (C₁) is simply a positive evaluation of a certain state of affairs, (C₂) is a claim about a certain relation of normative requirement that holds between certain agents and a certain action. As a consequence, claims like (C₁) can be true in kinds of situations where claims like (C₂) cannot.

The English ‘ought’ is thus an instructive example of the importance of being clear about our talk of normativity. I have suggested that, at bottom, such talk is about our familiar experiences of what I have called evaluative- and prescriptive-normativity. What the English ‘ought’ illustrates, however, is that what makes a speech act, a judgment, a concept, or a property appropriate in the context of one of these experiences may well differ from what makes it appropriate in the context of the other. This general distinction and this particular illustration will both be instrumental in the discussion of the evasive strategy that follows below.⁶

3.3 The Evasive Strategy

Recall the deontic puzzle: if (I) is true—if it is not the case that forming the belief that p is within S’s voluntary control—then the antecedent of (OIC) is always true when ϕ is believing that p. And if (D) is true—if some claims of the form ‘S ought to believe that p’ are true—then the consequent of (OIC) is sometimes false when ϕ is believing that p. Since the antecedent of (OIC)

⁶Passages such as the above by Grimm also suggest that speech acts, judgments, concepts, or properties *other than* those related to the English ‘ought’ can be part of our experience of prescriptive-normativity. This suggestion will be important for the discussion in section 5 below.

cannot always be true and its consequent sometimes false for the same ϕ , at least one of the independently plausible (OIC), (D), and (I) must be false.

We can turn the deontic puzzle into an argument against (D):

The Main Argument Against Epistemic Deontology:

1. If ϕ -ing is not within S's voluntary control, then it is not the case that the claim 'S ought to ϕ ' is true.
2. It is not the case that believing that p is within S's voluntary control.
3. So it is not the case that the claim 'S ought to believe that p' is true.

The Main Argument seems to show that a commitment to (OIC) and (I) forces a rejection of (D). But this appraisal is premature. We have already noted the surface-grammar ambiguity between evaluative-normative ought-claims and prescriptive-normative ought-claims. Now we must clarify which of these two kinds of ought-claims are being deployed in the premises of the Main Argument.

The first stage of the evasive strategy is precisely the suggestion that true claims of the form 'S ought to believe that p' are always evaluative-normative ought-claims and never prescriptive-normative ought-claims. When we say truly that S ought to believe that p, that is, we are always expressing something of a kind with what we express by saying that the clock ought to strike on the quarter hour; we are always merely claiming that S's believing that p is well-ranked by a certain standard we endorse, and we are never claiming that there is a relation of normative requirement holding between S and believing that p. Let 'ought_E' stand for the evaluative-normative sense of the English

‘ought’ and let ‘ought_P’ stand for the prescriptive-normative sense. We can then say that, according to the evasive strategy, the proper understanding of (D) is this:

Epistemic *Evaluative-normative* Deontology (D_E): Some claims of the form ‘S ought_E to believe that p’ are true.

And not this:

Epistemic *prescriptive-normative* Deontology (D_P): Some claims of the form ‘S ought_P to believe that p’ are true.

If this interpretation is correct, then the truth of claims of the form ‘S ought to believe that p’ does not require that S be capable of bringing about the state of affairs of believing that p.⁷

The second stage of the evasive strategy is the suggestion that (OIC) cannot be true if its consequent refers to evaluative-normative ought-claims. Having voluntary control, after all, is not required for their truth. According to the evasive strategy, that is, the proper understanding of (OIC) is this:

(OIC_P): If ϕ -ing is not within S’s voluntary control, then it is not the case that the claim ‘S ought_P to ϕ ’ is true.

And not this:

⁷Kornblith (2001, 238) takes true doxastic ought-claims as expressions of our *epistemic ideals*, Chrisman (2008, 364; 2012b, 603) takes them as *rules of criticism* or *state-norms* (respectively), and Feldman (2008, 351) takes them as *role oughts*. These are three different accounts of why doxastic oughts are not the kind of ought that requires voluntary control.

(OIC_E): If ϕ -ing is not within S's voluntary control, then it is not the case that the claim 'S ought_E to ϕ ' is true.

Properly understood, (OIC) is the claim that a certain relation of normative requirement holds between an agent and an action only if that action is within that agent's voluntary control. Properly understood, (OIC) places no constraints on the truth of any positive or negative mere evaluation.

The third and final stage of the evasive strategy simply replaces the original premises of the Main Argument with their appropriate disambiguations:

- 1*.** If ϕ -ing is not within S's voluntary control, then it is not the case that the claim 'S ought_P to ϕ ' is true.
- 2.** It is not the case that believing that p is within S's voluntary control.
- 3*.** So it is not the case that the claim 'S ought_P to believe that p' is true.

(3*), however, is no denial of (D_E). While (I) and (OIC_P) show that no prescriptive-normative ought-claims are true about an individual and her beliefs, they nonetheless spell no trouble for (D_E). Whatever we say about S's agency, sometimes it is true that S believing that p is well-ranked according to the standards that we endorse. The evasive strategy thus amounts to accepting the conjunction of two claims:

The Evasive Strategy:

(ES₁) (I), (OIC_P), and (D_E).

(ES₂) \neg (OIC_E) and \neg (D_P)

This is why evasion has seemed to many to be the least controversial and the least costly of the available maneuvers against the deontic puzzle: it does not require a rejection of (I), it does not require a rejection of (OIC), and it does not require abandoning normative talk of ‘ought to believe’.

3.4 Epistemic Responsibility

My aim is to argue that the evasive strategy is, nonetheless, quite costly and controversial. This means, first, identifying some controversial consequence of it and, second, making the case that this consequence is in fact costly. In this section, I will consider and criticize Rik Peels’ recent attempt to do just that.

According to Peels (2014, 682-3), the evasive strategy comes at the cost of *epistemic responsibility*: of it being appropriate to praise or blame S for, specifically, having the belief that p. You may think, for example, that some racists are blameworthy for their racist beliefs. That is, you may think that it is appropriate to blame them for their beliefs and not just for whichever actions (if any) result from those beliefs. In such a case, you will thereby think that some racists are epistemically responsible for their racist beliefs. According to

Peels, the evasive strategy comes at the cost of holding people responsible in this way.⁸

The heart of Peels' (2014, 691) argument is the claim that responsibility, thus understood, requires voluntary control:

(control) S is responsible for ϕ -ing only if ϕ -ing is under S's voluntary control.⁹

There is nothing implausible about (control). But part of the attraction of the evasive strategy, part of what makes it less costly than alternative strategies against the deontic puzzle, is that it accepts (I); it accepts that believing that p is not within our voluntary control. So if (control) is true, then the evasive strategy comes at the cost of the possibility of being responsible for believing that p. If (control) is true, that is, then the evasive strategy comes at the cost of the possibility of ever being blameworthy or praiseworthy for our beliefs. If both (I) and (control) are true, then, whatever we say about the racist's actions, it is simply inappropriate to blame him for what he believes.

Perhaps this is, in fact, a costly consequence of the evasive strategy. Perhaps we (or some or many of us) have strong intuitions that there really is such a thing as epistemic responsibility; perhaps sacrificing epistemic responsibility amounts to, or leads to, modifications in central aspects of our epistemology.

⁸See Van Woudenberg (2009) for the claim that we do, ordinarily, hold people responsible for their beliefs in precisely this way.

⁹According to Peels (2014, 691), moreover, "someone has voluntary control over ϕ -ing if and only if one can ϕ as the result of an intention to ϕ and one can $\neg\phi$ as the result of an intention to $\neg\phi$." My criticisms will not turn on whether one accepts this or alternative accounts of voluntary control.

I do not wish to consider these possibilities at length here. Instead, my aim is showing that Peels' argument has a serious shortcoming. There are widely accepted reasons for thinking that (control) is false, and Peels has not succeeded in his attempt to resist them.

According to a widely accepted account, what responsibility requires is not voluntary control, but rather *reason-responsiveness*. When S ϕ 's as a result of responding to what S perceives as sufficient reasons to ϕ , according to this account, then S is the proper subject of praise and blame for ϕ -ing. This is because assessments of responsibility are assessments of the quality of someone's will, and because the quality of someone's will is revealed by how one responds to what one perceives as one's reasons.¹⁰

A reason-responsiveness account of responsibility, in fact, not only motivates a denial of (control), but it also explains why it is appropriate to hold people responsible, specifically, for their beliefs. That's because we typically form beliefs as a result of appreciating our reasons. As Ryan (2003, 65) puts it, "my doxastic decisions are guided by what seems to me to be good evidence." This hardly seems controversial. But in believing that p as a result of appreciating our reasons, we thereby reveal something about "our take on the world, on what is or is not true or important or worthwhile in it" (cf. Hieronymi 2008, 370). So just as we can be blameworthy for how we respond to our perceived reasons for punching someone in the face, we can be blameworthy for how we respond to our perceived reasons for believing that p. In believing that p,

¹⁰There is a variety of such accounts on offer. See, e. g., Fischer & Ravizza (1998, 62-91), Ryan (2003, 70-74), Arpaly (2006, 16-7; 19), and Hieronymi (2008, 359-363).

just as much, we typically reveal something about the quality of our will. Yet in both the punching and believing cases, according to reason-responsiveness accounts, the question of whether the relevant ϕ -ing was under our voluntary control is simply irrelevant for assessments of praise and blame.

Since Peels' argument against the evasive strategy depends on (control), his argument falls short if such a reason-responsiveness account of responsibility is correct. Peels (2014, 693) is sensitive to this. He offers the following dilemma as a reply:

Peels' Dilemma: S either has or does not have voluntary control over the higher-order beliefs and reason-responsive processes responsible for S's belief that p.

(Horn #1) If S does, then S's belief that p is under S's voluntary control after all.

(Horn #2) If S does not, then S is simply not responsible for S's belief that p.

(Horn #1) allows for what we can call *indirect* voluntary control. The idea is that our beliefs are under our indirect voluntary control in the same way that a certain arrangement of furniture in our living room is under our indirect voluntary control: in both cases, we have control by virtue of several other actions over which we have *direct* voluntary control. (Horn #2) claims that there simply cannot be any responsibility in the absence of either direct or indirect control.

I think this reply by Peels is seriously misguided. Peels' Dilemma fails to appreciate the full force of the reason-responsiveness account. This account, after all, is formulated in the backdrop of causal determinism; it is intended to deliver responsibility, that is, in the absence of *any kind* of voluntary control. These accounts, that is, are *compatibilist*. Consider Ryan (2003, 71):

We cannot, except in unusual cases, control the fact that our beliefs will be determined by our awareness of the evidence *in much the same way* that we cannot control the fact that our actions will be determined by the laws of nature, and events in the past, if determinism is true. (My emphasis)

The reason-responsiveness account grants from the very beginning that responsibility is compatible with complete lack of control. So (Horn #2) simply begs the question against their very central claim, namely, that responding to reasons—*qua* response to reasons, and not *qua* voluntary action—reveals something about the quality of our will, and is thus an appropriate object of praise and blame.¹¹ My claim here, notice, is not that these considerations settle whether (control) is false. My claim is merely that Peels' Dilemma places no pressure on those already committed to a denial of (control). Peels' Dilemma does not motivate a rejection of the reason-responsiveness account, that is, it merely assumes it.

So Peels fails to show that the evasive strategy comes at the cost of epistemic responsibility. His argument depends on the plausible but widely rejected

¹¹Peels is no doubt entitled to his recalcitrant intuitions. (See, e.g., Peels' (2014, 693-4) thought experiment.) But Ryan (2003, 70) is quite explicit about her intentions when she says: "I do not expect anyone who is unmoved by compatibilism to find plausible the view I am about to articulate and defend." See also Steup (2012, 153) for the claim that what matters for control is not *whether* our actions are determined but *how* they are determined.

(control), and his argument against this widespread rejection falls short. Importantly, this is not to say that the evasive strategy *does not* come at the cost of epistemic responsibility. In fact, I will briefly suggest a different argument for this conclusion at the end of the next section. But appealing to (control), for the reasons just mentioned, is a dead end.

Before moving on, let me note one more reason for thinking that Peels' argument against the evasive strategy is less than overwhelming. Suppose that sacrificing epistemic responsibility is a real consequence of the evasive strategy. Even then, it is not clear that this is a *costly* consequence. Both Kornblith and Feldman, in fact, are quite willing to concede that there is no such thing as epistemic responsibility: as being praiseworthy or blameworthy for one's *beliefs*.¹² Whatever the plausibility of their dismissal of the notion of epistemic responsibility, there is something at least dialectically misguided about claiming that a costly consequence of the evasive strategy is something that two of the three most prominent evaders are antecedently quite willing to give up. They may well reply that this is not a bug but rather a feature of their views. So even if Peels had succeeded in showing that sacrificing epistemic responsibility was a consequence of the evasive strategy—which I have argued he did not—his argument would still fall short of showing that, by requiring extensive modifications to central aspects of one's epistemology, the evasive

¹²Kornblith (1983, 38) accepts that being justified for believing that *p* is (at least partially) a matter of being responsible, but he takes the relevant kind of responsibility here as notion that applies to epistemically-relevant actions other than belief formation and maintenance. Feldman (2008, 353) is quite explicit about his doubts regarding ordinary talk of epistemic responsibility, as well as about what would be lost if such talk turned out to be inappropriate after all.

strategy has a consequence that is costly. Much better, that is, if we can identify a consequence that will seem costly even to those attracted to the evasive strategy in the first place.

3.5 Undefeated Normative Reasons

I think the real price of the evasive strategy lies elsewhere: it sacrifices the possibility of normative reasons for belief.¹³ In this section, I will argue that this is really a consequence of the evasive strategy. In this next and final section, I will argue that it is a consequence that is truly costly. My argument here turns on understanding how normative reasons belong within our experiences of prescriptive-normativity. While Peels identifies (I) as the commitment that is problematic for the evasive strategy, I identify $\neg(D_P)$ instead. We will arrive at my argument in three steps.

The first step is accepting the familiar picture of normative reasons as the building blocks of relations of normative requirement: it is precisely *because* S has undefeated authoritative considerations in favor of ϕ -ing that an authoritative relation of normative requirement holds between S and ϕ -ing. While our reasons are considerations of various strengths in favor of various actions at a certain time, the balance of these considerations determines what we are

¹³See Schroeder (2007, 10-15; 2015, sec. 3.1) for the contrast between normative, explanatory, and motivating reasons. It is notoriously hard to say something informative about normative reasons that is neither vague or metaphorical. A grip on the pre-theoretical and paradigmatic sense in which an innocent person's unnecessary suffering is a normative reason against torturing her for fun—the sense in which that *fact* is an authoritative consideration against that *action*—will be enough for present purposes.

required to do. As Broome (1999, 400) puts it, “If you have a reason to q and no reason not to q , then you ought to q .” As Skorupski (2010, 48) puts it, “‘ x should ψ ’, ‘ x ought to ψ ’ and ‘it is right for x to ψ ’ can all be used to say that there is a sufficient reason for x to ψ .” Using familiar terminology, the idea we find here in both Broome and Skorupski is that the prescriptive-normative English ‘ought’—ought_P, the ‘ought’ that expresses the existence of a relation of normative requirement between an agent and an action—expresses what we have *all-things-considered* reasons to do. I will call this *normative rationalism*:

(NR) There is a normative requirement for S to ϕ if and only if (and because) there is most normative reason for S to ϕ .¹⁴

Accepting (NR), notice, does not require taking a stand on the controversial question of whether considerations other than evidence can count as a reason to believe that p .¹⁵ It also does not require taking a stand on whether reasons for belief are reasons to promote some value or to respect some norm.¹⁶ (NR) is simply an account of the source of normative requirements, of the kind of normativity that we experience as binding. As such, (NR) is compatible with whatever views one has about what counts as reasons to believe, as well as whatever views one has about the sources of these reasons. It is important not

¹⁴I take this to be one way of expressing the view made famous by W.D. Ross (1930). For more recent developments, see Broome (1999), Raz (2002), Skorupski (2010), Parfit (2011), and Smith (2013). See Schroeder (2007) for the contrast between normative, explanatory, and motivating reasons.

¹⁵See Raz (2011, 36) for a negative answer; see Schroeder (2012, 471) for a positive answer.

¹⁶See Talbot (2014, 602-3) for defense of the promotion account; see LittleJohn (2013, 359) for defense of the respect account.

to confuse one's reasons for rejecting one of these other views with a reason to reject (NR).

The second step towards my argument is noticing an implication of (NR): if there is no relation of normative requirement holding between S and ϕ -ing, then either (a) there is no reason for S to ϕ or (b) these reasons are defeated by reasons in favor of doing something else. At nearly all times, there are available actions such that it is not the case that S ought_P to perform them at that time. But this can be so for different reasons. Sometimes this is the case because S has no reason whatsoever to perform them. Perhaps right now, for example, I have absolutely no reason to raise my index finger straight up into the air; consequently, it is not the case that I ought_P to raise my finger in that way: raising my finger, right now, is not something that I am required to do. But sometimes it is not the case that S ought_P to perform a certain action despite in fact having reasons for performing it. Perhaps right now, for example, I have a reason to have a cup of coffee (say, because I have a certain desire for it), but have even stronger reasons not to have it (say, because I've already had too many); consequently, once again, it is not the case that I ought_P to have a cup of coffee: having a cup of coffee is not something that I am required to do.

It is worth stressing that this implication of (NR) does not show that (NR) forces a choice on the two controversial questions mentioned just above. Consider Littlejohn (2013, 359):

To think of epistemic reasons as reasons to promote the attaining of an end, one must think that there is 'room' to explain an epistemic ought in terms of a kind of weighing explanation that's common from the practical case. There is no room for that sort of reasoning

in the theoretical sphere because unlike action, belief serves a single master.

Littlejohn seems to be making two suggestions here. First, that since only evidence can be a reason for belief (that is the “single master” he is alluding to in this passage), a relation of normative requirement holding between S and a certain belief cannot be explained in terms of the weighing of reasons. Second, that since a relation of normative requirement holding between S and a certain belief cannot be explained in terms of the weighing of reasons, reasons for belief cannot be reasons to promote. If Littlejohn is correct in his first suggestion, however, then the implication of (NR) under consideration forces a rejection of the view that only evidence can be a reason for belief. But Littlejohn is not correct in his first suggestion. Suppose only evidence can be a reason for belief; suppose some of my evidence points to the butler being guilty and some of it points to the driver. I see no conceptual reason to think that this could not be case where my evidence for the butler being guilty is *stronger* than my evidence for the driver being guilty, and that therefore my reasons for believing that the butler did it *outweigh* my reasons for believing that he did not. So even if only evidence can be a reason for belief, a requirement-relation holding between S and a certain belief can nonetheless be explained in terms of the weighing of reasons.

The third and final step towards my argument is noticing an implication of the conjunction of (NR) and the evasive strategy’s denial that it is ever true that S ought_p to believe that p. Given the implication of (NR) just mentioned, this denial means accepting that it is *always* true that either:

(a₁) There is no reason for S to believe that p.

Or:

(b₁) S's reasons for believing that p are defeated by S's reasons in favor of doing something else.

Since reasons are the building blocks of relations of normative requirement, denying that a certain relation of normative requirement holds between S and ϕ means denying that there are *undefeated reasons* for S to ϕ . Denying that there is *ever* a relation of normative requirement holding between S and ϕ , similarly, means denying that there are *ever* undefeated reasons for S to ϕ . It is here that we see the true cost of the evasive strategy: while it succeeds at escaping the deontic puzzle by taking doxastic oughts as evaluative-normative ought-claims that do not require agency, it sacrifices the existence of undefeated normative reasons for believing that p. It thus retains one kind of epistemic normativity, alright, but it abandons the kind of epistemic normativity that we experience as binding.

We can state my argument more precisely in the following way:

The Undefeated Reasons Argument

1. If the evasive strategy is correct, then claims of the form 'S ought_P to believe that p' are always false.
2. If claims of the form 'S ought_P to believe that p' are always false, then there are never undefeated normative reasons for believing that p.

3. If the evasive strategy is correct, then there are never undefeated normative reasons for believing that p.

Premise 1 follows from the fact that the evasive strategy accepts both (I) and (OIC_P). This is an important element of the evasive strategy, since it thereby avoids some of the costs of alternative strategies against the deontic puzzle. Premise 2 follows from the familiar picture of normative reasons as the building blocks of requirement-relations: (NR). The conclusion is not a rejection of the evasive strategy, of course, but it may well be taken as good grounds for it.¹⁷

Before moving on, let me indicate rather briefly how (NR) can be used in two further arguments against the evasive strategy. The first argument is similar to the Undefeated Reasons argument, but stronger. Its strength, however, comes from accepting two further claims. First, that only evidence can be a reason for believing that p. Second, that one's evidence always supports believing, disbelieving, or suspending judgment with respect to p—alternatively, that one's evidence always supports having a specific degree of confidence regarding p. There is nothing implausible about either of these claims, but they of course require some defense.¹⁸ At any rate, if (NR) and both these claims are true,

¹⁷Notice that there is nothing odd or unusual about bits of evaluative-normativity that are disconnected from normative reasons. Certain thick-concepts—slurs, for example—are clear examples: they are often sincere evaluations of ϕ -ing, but evaluations that nonetheless fail to indicate normative reasons for ϕ -ing or not. So my argument is not that it is in principle implausible to claim, as the evasive strategy does, that certain evaluations can be normative (in one sense) without indicating the existence of normative reasons. See Papineau (2014), in fact, for an explicit endorsement of the claim that epistemic evaluations do not indicate normative reasons for belief.

¹⁸Schroeder (2015, sec. 4.2), in fact, suggests that these two claims cannot be held together since evidence can only be a reason for a belief or its negation. He does not, however, consider the possibility of *higher-order evidence* being a reason for suspending judgment.

then the cost of the evasive strategy is not only the possibility of undefeated normative reasons for belief, but the possibility of normative reasons for belief altogether. I will not, however, pursue this stronger argument here.

The second argument has the same conclusion as Peels' argument. The argument, however, does not depend on Peels' (control), but rather on the claim that S is epistemically responsible for believing that p only if there are reasons for S to believe that p . There is nothing implausible about this claim either. As Skorupski (2010, 295) puts it: "there can be no reason to blame someone who had no reason not to do what he did." This claim, in fact, is not only compatible with the reason-responsiveness account of responsibility—the account that allows one to resist Peels' argument—it may even be essential to it. (This will depend on whether one takes responding to reasons as *one of* or *the only* way in which we reveal the quality of our will.) But if (NR) and this claim are both true, then the evasive strategy sacrifices not only the existence of normative reasons for belief, but also the possibility of epistemic responsibility. I will also not pursue this argument here.

3.6 Letting Reasons Go

I have suggested that it is a consequence of the evasive strategy that we never have undefeated normative reasons to believe that p . I now wish to consider two defensive maneuvers. These are not, however, attempts to reject the Undefeated Reasons argument. Instead, these are attempts to show that denying that we ever have reasons to believe that p is not so costly after all. The first maneuver is to claim that while we can never have reasons to believe that p , we can

nonetheless have reasons to perform other epistemically relevant actions. The second maneuver is to claim that while we can never have reasons to believe that p , in some sense, we can nonetheless have constitutive reasons to believe that p . I will suggest that neither of these maneuvers succeeds.

Chrisman (2008, 369-70) suggests that evaluative-normative doxastic ought-claims materially imply prescriptive-normative non-doxastic ought-claims. Consider:

When we say, ‘You ought to disbelieve that the earth is flat’, it is plausible to suppose that this... implies both intra-personal rules of action such as ‘You ought to have read your science books and listened to your parents and teachers’ and inter-personal rules of action such as ‘Your parents and teachers ought to have taught you that the earth is not flat’.

Here the first use of the English ‘ought’ is an evaluative-normative ought-claim, while the second and third uses are prescriptive-normative ought-claims.¹⁹ If Chrisman is correct, then true claims of the form ‘S ought_E to believe that p ’ imply true claims of the form ‘S ought_P to ϕ ’ (where ϕ is not believing that p). While we may never have reasons to believe that p , that is, we may nonetheless have reasons to perform a host of epistemically relevant actions such as re-considering the evidence as to whether p , gathering more evidence as to whether p , disseminating reliable information as to whether p to those around us, and so on. Perhaps the existence of reasons for epistemically relevant actions

¹⁹See also Chrisman (2012b, 609): “The validity of state-norms would, I think, come into question if there were no discernable connection between things that people can choose to do and conformity to state-norms.”

makes the consequence of the evasive strategy brought out by the Undefeated Reasons argument seem less significant.

Whatever we say about the implications of evaluative-normative ought-claims, I think they hardly suffice as substitutes for the typical role that talk of reasons has played in epistemology. I have in mind the way that many take the fact that ‘there is evidence for p’ as either identical or constitutive of the fact that ‘there is a reason for believing that p’. Call this the *evidence* principle:

(evidence) If *e* is evidence for S that p, then *e* is a reason for S to believe that p.

There are competing accounts of what can constitute S’s evidence for p, of course, but (evidence) is compatible with all of them.²⁰ What the Undefeated Reasons argument shows, however, is that the evasive strategy is committed to a denial of (evidence). Whatever we say about our reasons to gather evidence and think carefully about it, I suggest that it is a costly consequence of evasion that we cannot as well say that the evidence gathered and considered is a reason to believe that p.

There is, however, another reason why appealing to the implications of evaluative-normative ought-claims is not enough as a defensive maneuver

²⁰See Turri (2009, 503-504) and Conee & Feldman (2011, 321-23) for the claim that S’s evidence consists of some of S’s mental states; See Dougherty (2011, 230-31) and Littlejohn (2012, 99-109) for the claim that S’s evidence consists of certain propositions. All of these author’s subscribe to (evidence). (evidence) is also neutral on the further question, mentioned above, of whether *only* evidence provides reasons for belief.

against the Undefeated Reasons argument. Recall the passage from Grimm (2009, 254) quoted in section 1 above:

To judge someone's belief to be unjustified or irrational is thus to judge that the person's attitude towards the content of the belief should be reconsidered, in some apparently binding sense of 'should.'

Grimm is here giving expression to a widely accepted picture of epistemic evaluations as belonging to our experiences of prescriptive-normativity. To say that someone's belief is unjustified, according to Grimm, is to indicate that there is a relation of normative requirement between S and that belief. Given (NR), this picture of epistemic evaluations thus entails that we sometimes have undefeated normative reasons for believing (or not believing) that p. Cuneo (2007, 58) goes so far as calling this feature of epistemic evaluations a "platitude."²¹ So another consequence of the evasive strategy brought out by the Undefeated Reasons argument is the denial of this picture of epistemic evaluations as binding. I suggest that this is also a costly consequence of evasion.

I have suggested that *even if* Chrisman is correct in his claim that evaluative-normative ought-claims imply prescriptive-normative ought-claims, this is not enough to mitigate the cost of the evasive strategy brought out by the Undefeated Reasons argument. This first defensive maneuver, that is, provides

²¹Cuneo (2007, 58-9) says "epistemic facts are authoritative in the following twofold sense. They are authoritative, in the first place, insofar as they are prescriptive. That is to say, epistemic facts are, imply, or indicate reasons for properly situated agents to behave in certain ways... The second sense in which epistemic facts are authoritative is that some such facts inescapably govern our conduct. The fundamental idea in this case is that some epistemic facts are, imply, or indicate categorical reasons for agents to behave in certain ways."

no defense at all. There is, however, a second defensive maneuver available to the evasive strategy. Consider Chrisman's (2008, 349-50) distinction between *extrinsic reasons* for believing that p and *constitutive reasons* for believing that p:

Constitutive reasons for the belief that p are reasons that bear on the question whether p... By contrast, extrinsic reasons for the belief that p are reasons that count in favor of believing p independently of whether p... Evidence provides constitutive reasons for belief, while other considerations provide (if anything) extrinsic reasons for belief.

Perhaps the Undefeated Reasons argument shows only that the evasive strategy comes at the cost of extrinsic reasons to believe that p; perhaps there is then room for the claim that we can nonetheless have constitutive reasons for believing that p. Perhaps the existence of constitutive reasons for belief, even in the absence of extrinsic reasons for belief, makes the consequence of the evasive strategy brought out by the Undefeated Reasons argument seem less significant.

Once again, I do not think that this maneuver—even if we accept it—is enough to mitigate the force of the Undefeated Reasons argument. Notice that claiming that the Undefeated Reasons argument is consistent with there nonetheless being constitutive reasons for believing that p is tantamount to suggesting that evidence for p constitutes a kind of reason for believing that p that is not a building-block of relations of normative requirement. One immediate consequence of this maneuver, then, is once again the denial of the familiar picture of epistemic evaluations as prescriptive-normative. Most of us don't think of evidence as providing us with a *lesser*, non-binding kind

of reason; yet constitutive reasons would have to be some such lesser kind if they are taken as compatible with the Undefeated Reasons argument. So this maneuver does not eliminate at least one the costly consequences of evasion already mentioned.²²

I suggest that these considerations—if not individually, then together—show that *if* the Undefeated Reasons argument is sound, then the evasive strategy is quite costly indeed. In the previous section, of course, I have argued that the Undefeated Reasons argument is, in fact, sound. So I conclude that the real costly consequence of evasion is sacrificing the possibility of undefeated normative reasons for belief.

3.7 Conclusion

These considerations, notice, do not suggest that (D_E) is false.

(D_E) : Some claims of the form ‘S ought_E to believe that p’ are true.

Perhaps some claims about what we ought to believe really are evaluative-normative ought-claims. But contrary to what some have suggested, it matters to normative epistemology that (D_P) is true *as well*.

(D_P) : Some claims of the form ‘S ought_P to believe that p’ are true.

²²Also, there are reasons—good reasons, I think—to reject the claim that the Undefeated Reasons argument is consistent with there being constitutive reasons for believing that p. That’s because this would require epistemic reasons to be a kind of reason that is radically different from practical reasons. This suggestion, however, is implausible. See, e.g. Kornblith (2002, ch. 5), Turri (2011, 384), Littlejohn (2012, 105-6), Booth (2012, 511-512), and Schroeder (2015, sec. 3.2).

It is important that some claims about what we ought to believe really are prescriptive-normative. It is only if there are relations of normative requirement binding S as an agent to some of her beliefs—that is, only if there are undefeated normative reasons for believing that p—that we can make sense of common claims, such as Grimm’s above, that epistemic evaluations are somehow binding. But when (D) is understood as (D_P) , the deontic puzzle is still very much alive: the Main Argument entails that (D_P) is false, and resisting it seems to require a denial of the independently plausible (I) or (OIC_P) . So while the evasive strategy allows us to retain most of the surface-features of ordinary normative discourse, it requires silencing powerful intuitions (that we sometimes have normative reasons to believe that p) as well as modifying other aspects of one’s epistemology (the deontic aspects connected to having reasons, such as the bindingness of epistemic evaluations). Whatever the merits of the evasive strategy against the deontic puzzle, the strategy seems to me far from uncontroversial and theoretically inexpensive.

CHAPTER 4

DEONTOLOGICAL EVIDENTIALISM, WIDE-SCOPE REQUIREMENTS, AND PRIVILEGED VALUES

4.1 Introduction

John Locke is famous for prescribing a close connection between evidence and belief. Here is a representative passage:

We should keep a perfect indifference for all opinions, not wish any of them true, or try to make them appear so; but being indifferent, receive and embrace them according as evidence, and that alone, gives the attestation of truth. (*Conduct*, §34)

On one natural reading of this and other passages, Locke seems to endorse the claim that I will call *deontological evidentialism*:

(DE) S ought to form and maintain S's beliefs in accordance with S's evidence.

The English 'ought' here expresses an authoritative relation of normative requirement or obligation. Fully stated, (DE) is thus the claim that there is an authoritative relation of normative requirement that holds between each individual and the complex of actions and attitudes that constitute forming

and maintaining one's beliefs in accordance with one's evidence.¹ As some have put it in different contexts, (DE) sees this normative requirement as "having a grip" on us (cf. Korsgaard 1996, 44-6), as something that is "demanded" of us (cf. Street 2012, 44), and as something "utterly different from anything else in the universe" (cf. Mackie 1977, 38).

Locke's argument for (DE) has two descriptive premises and one normative premise. The first descriptive premise is the claim that God wants us to do our best to avoid believing that p if p is false (cf. *Essay*, IV, XVII, 24). The normative premise is the claim that if God wants us to ϕ then we have a normative requirement to ϕ (cf. *Essay*, XX, XXVIII, 7-8). From these it follows that we have a normative requirement to do our best to avoid believing that p if p is false. The second descriptive premise is the claim that doing our best to avoid believing that p if p is false is constituted by forming and maintaining our beliefs in accordance with our evidence (cf. *Essay*, IV, XX, 3). From these it follows that (DE) is true. Locke's defense of (DE) is thus both theistic and moral. We have a normative requirement to form and maintain our beliefs in accordance with our evidence because of a God-created moral requirement to do our best to avoid believing that p if p is false.²

¹The source of this normative authority is a matter of debate. Since the term 'obligation' is so often and so naturally associated with *moral* obligation, I will here give preference to the more neutral term 'normative requirement'. I will elide the 'authoritative' qualifier throughout.

²Some argue that Locke took his evidentialism to be restricted to those propositions that are of most importance to us (e.g. those about religion and morality). See Wolterstorff (1996, 63-66) for a defense of this reading. Since Locke's views are not my main focus, I will put this exegetical detail to the side.

To my knowledge, Locke is the first to argue explicitly for (DE). Yet those who do not believe in God will naturally find this argument unconvincing. Even some who *do* believe in God will likely disagree with Locke's description of what God wants of us and our beliefs. Perhaps, for example, God wants us to believe in his existence, in his salvific actions, and so forth, not on the basis of evidence but rather on the basis of trust and love and faith.³ There is a very small audience, that is, for whom Locke's argument has any pull. (This is no refutation of his argument, of course; I will not offer one here.) Nonetheless, (DE) remains alive and well. My interest in this paper is in examining two alternative arguments in its defense.

Here is how I proceed. In the first section, I discuss Berit Brogaard's (2014) reply to W.K. Clifford's well-known defense of (DE). My aim here is clarifying her use of the distinction between narrow-scope and wide-scope requirements against Clifford. In the second section, I discuss how we can turn Stephen Grimm's (2009) recent claims about the moral source of epistemic normativity into a novel argument for (DE). I argue, however, that the distinction between narrow-scope and wide-scope requirements—used by Brogaard against Clifford—is effective when used against Grimm as well. In the third section, I take time to fill-in the details of Richard Feldman's (2000, 2001, 2008) defense of (DE) and I clarify the unstated commitments that it involves. In the fourth section, I discuss Hilary Kornblith's (2001) reply to Feldman's defense. I argue that Kornblith's use of the distinction between narrow-scope and wide-scope

³See Plantinga and Wolterstorff (1983).

requirements only allows for a reply that is incomplete: it correctly identifies *what* is wrong with Feldman’s defense, but it provides the incorrect explanation for *why* this is so. In the fifth and final section, I provide an alternative reply. I argue that Feldman fails to appreciate the distinction between ought-claims that are true relative to some arbitrary value and ought-claims that are true relative to *privileged values*: values that have normative authority over us.⁴

4.2 The Consequentialist Moral Argument

According to Clifford, false beliefs always have negative moral consequences. His prime example is of a shipowner whose false belief that his ship is seaworthy costs the lives of several innocent families. Because of such inescapable negative moral consequences, Clifford takes it that we have a normative requirement to do our best to avoid believing that p if p false. Since doing our best to avoid believing that p if p is false is constituted by forming and maintaining our beliefs in accordance with our evidence, it once again follows that we have a

⁴Cowie (2014, 4003) argues that there is a presumptive case in favor of what he calls *instrumentalism* about epistemic normativity: the claim that “there is reason to believe in accordance with one’s evidence because this is an excellent means of fulfilling the goals that one has, or should have.” His argument, however, is entirely negative. It is a product of his (2014, 4004) criticism of what he takes as the only non-error-theoretic prominent alternative, which he calls *intrinsicism* about epistemic normativity: the claim that “there is reason to believe in accordance with one’s evidence in virtue of a brutally epistemic normative truth relating belief to evidence, or to some other epistemic property such as truth, or epistemic rationality”. The Lockean views of epistemic normativity discussed here, however, do not fit within either of these prominent categories. They explain epistemic normativity—unlike intrinsicism—but do not appeal in any way to our goals—unlike instrumentalism. This is another reason—besides their historical influence and surface plausibility—why they deserve the separate and careful treatment they receive here.

normative requirement to form and maintain our beliefs in accordance with our evidence. Call this the *consequentialist moral argument* for (DE):

The Consequentialist Moral Argument

1. False beliefs always have negative moral consequences.
 2. If false beliefs always have negative moral consequences, then we have a normative requirement to do our best to avoid believing that p if p is false.
 3. So we have a normative requirement to do our best to avoid believing that p if p is false.
 4. Doing our best to avoid believing that p if p is false is constituted by forming and maintaining our beliefs in accordance with our evidence.
- C. So we have a normative requirement to form and maintain our beliefs in accordance with our evidence.

Notice that the difference between the Locke's theistic moral argument and Clifford's consequentialist moral argument is minimal. The latter appeals to negative moral consequences while the former appeals to God's binding desires. Otherwise the arguments are identical. Yet notice the irony as well: Locke relies on his religious beliefs in his defense of (DE) and Clifford offers a nearly

identical argument, for the very same position, in his famous polemic against religion. (DE), we can say, is an equal opportunity employer.⁵

Clifford's moral consequentialist argument is often rejected on the basis of simple counter-examples. Sometimes, for example, believing against the evidence is life-saving and nothing else seems to hang in the balance. Perhaps a patient in critical care will increase her chances of survival by a non-trivial degree if she believes, against the evidence, that she is very likely to recover (cf. Feldman 2006, 30). Even if this belief is false, it does not seem to have negative moral consequences. So this seems to be a counter-example to premise (1). Sometimes, for another example, believing against the evidence is simply trivial and isolated. Perhaps a certain shopkeeper quite uncritically believes, against the evidence, that the apples she just picked are the best apples in the supermarket (cf. Haack 2001, 24). Even if this belief is false, it once again does not seem to have negative moral consequences. So this seems to be another counter-example to premise (1).

Recently, however, Berit Brogaard (2014) has offered an alternative reply to Clifford's consequentialist moral argument. Brogaard notes that beliefs only give way to actions when coupled with dominant desires. If I believe that kicking a dog will not cause it pain and yet have no desire to kick a dog, then my belief will not by itself produce an act of kicking the dog. Similarly, if I have a desire to kick a dog but do not believe that the dog in front of me is a dog, then the desire will not by itself produce an act of kicking the dog. Brogaard thus

⁵I read Clifford differently from Brogaard (2014, 2). She takes him as committed to the premise that "false beliefs *could have* morally harmful consequences" (my emphasis).

infers that Clifford was mistaken in thinking that negative moral consequences gave rise to a normative requirement against forming and maintaining certain *beliefs*. Since negative moral consequences are the consequences of actions, they at best give rise to normative requirements against forming and maintaining those things that can properly give way to actions: *dominant belief-desire pairs*.

While Clifford claims that we have a normative requirement to avoid false beliefs, Brogaard claims that we have a normative requirement to avoid harmful belief-desire pairs. But the latter requirement can be satisfied in two different ways: we can either refrain from forming and maintaining the relevant belief or we can refrain from forming and maintaining the relevant desire. As Brogaard notices, this suggests that the difference between Clifford's claim and hers is in fact a difference in the *scope* of premise (2). While Clifford sees the connection between belief and action as giving rise to a narrow-scope normative requirement, Brogaard sees the same connection as giving rise to a wide-scope requirement instead:

(narrow₁) If false beliefs always have negative moral consequences, then we have a normative requirement to do our best to avoid believing that p if p is false.

(wide₁) If false beliefs always have negative moral consequences, then we have a normative requirement to be such that, *if* we have a false belief that p, then we refrain from forming the dominant desire d which, when coupled with p, would give way to an action that has negative moral consequences.

If (narrow₁) is true, then the Clifford’s consequentialist moral argument goes through. But if (narrow₁) is false and (wide₁) is true instead, then the argument is unsound. In fact, if Brogaard is right, then there is nothing particularly worrisome about *false* beliefs in themselves. Any belief can become part of a harmful belief-desire pair, so any belief can be such that one way of satisfying our normative requirements is to refrain from believing it.⁶

Notice, however, that Brogaard’s reply to Clifford does not show that we *do not* have a normative requirement to do our best to avoid believing that p if p is false. It does not show, that is, that premise (3) is false. What it shows instead is that Clifford’s defense of this claim—by way of premises (1) and (2) of his consequentialist moral argument—is unsound. This is because premise (2) is false: even if false beliefs always have negative moral consequences, it does not follow from this that we have a normative requirement to do our best to refrain from believing them. Put a bit differently, what Brogaard succeeds in showing is that correctly understanding the relation between beliefs and actions supports (wide₁) instead of (narrow₁). Yet there may be different arguments for the claim that we have a normative requirement to do our best to avoid believing that p if p is false—arguments for Clifford’s premise (3) that do not

⁶Talk of wide and narrow *scope* is here talk about the place of the deontic operator in the underlying logical structure of premise (2). Let ‘ \square ’ represent a normative requirement, let ‘F’ stand for ‘false beliefs always have negative moral consequences’, and let ‘E’ stand for ‘avoids believing that p if p is false’. Then (narrow₁) has the form $\forall x (Fx \rightarrow \square Ex)$. Now let ‘B’ stand for ‘has a false belief that p’ and let ‘R’ stand for ‘refrains from forming the dominant desire d which, when coupled with p, would give way to an action that has negative moral consequences’. Then (wide₁) has instead the form $\forall x (Fx \rightarrow \square (Bx \rightarrow Rx))$. I will leave the details about underlying logical structure merely implicit throughout. For a broader discussion of wide and narrow scope requirements in connection to *rationality*, see Kolodny (2005), Broome (2007), Schroeder (2009), and Way (2011).

rely on Clifford's premise (2). In fact, we have already seen one such argument. For Locke, recall, God simply does not want us to have false beliefs about His world, whatever the actions such beliefs may or may not produce. This shows that Brogaard's success against Clifford is compatible with (DE) being true. The question I ask in the next section is thus the following: is there a sound defense of the claim that we have a normative requirement to do our best to avoid believing that p if p is false that does not appeal to either God's binding desires or to (narrows_1)? I will consider and reject one such attempt.

4.3 The Grimm-Inspired Moral Argument

Stephen Grimm (2009, 253-4) claims that epistemic evaluations have a characteristic kind of *normative force*:

To judge someone's belief to be unjustified or irrational is thus to judge that the person's attitude towards the content of the belief should be reconsidered, in some apparently binding sense of 'should'... If I accept that a certain belief of mine is 'inapt' I seem now to have a reason to do something about my attitude toward the content of the belief.

On one natural reading, Grimm's talk of epistemic evaluations as intimately connected to a "binding sense of should" and to "reasons" reveals that he takes epistemic evaluations to convey that certain normative requirements hold between individuals and their beliefs. That is, though Grimm's claims are about epistemic evaluations such as "S is justified in believing that p ," he seems to either take them as the same kinds of claim as (DE) or as intimately connected to such claims—by entailment, implication, pragmatics, and so on. Grimm (2009, 258-9), moreover, claims that these relations of normative requirement,

conveyed by epistemic evaluations and holding between individuals and their beliefs, have a moral source. So although Grimm does not defend (DE) himself, it seems appropriate to examine the prospects for a Grimm-inspired attempt to rescue the moral arguments we find in Locke and Clifford.⁷

As I've mentioned, Grimm's own interest is in explaining the normative force of epistemic evaluations—the fact that they are intimately connected to a “binding sense of should” and to “reasons”. On one hand, it is implausible to take this force as a product of the *intrinsic value* of all true beliefs. Intuitively, true beliefs about the number of blades of grass on my lawn are not intrinsically valuable, even if they may sometimes serve esoteric purposes or take part in promoting idiosyncratic desires. On the other, restricting one's views of which true beliefs have intrinsic values while at the same time holding that the normative force of epistemic evaluations is relative to these values seems to commit us to the view that epistemic evaluations do not apply to certain beliefs. If true beliefs about the number of blades of grass on my lawn are not intrinsically valuable, then I could not be justified in having them even after careful and diligent field research. Grimm's claims about the nature of the normative force of epistemic evaluations thus aim at explaining how those evaluations apply to any and all beliefs without being committed to the implausible claim that all of them have intrinsic value.

Here is how Grimm (2009, 258-9) puts his suggested explanation:

Even though we might not care less about some belief (or better, some topic), it is nonetheless the case that other people might care

⁷Grimm (2009, 259 fn. 32) is in fact aware that his claims are, as he puts it, “Cliffordian.”

about the topic a great deal... Given that someone... might depend on us as potential sources of information about this topic, it seems that we have an obligation not to be cavalier when we form beliefs about the question... As a potential source of information for others, we have an obligation to treat any topic or any question with due respect.

There are at least two controversial claims worth highlighting from this passage. First, the claim that, for any belief or topic, it is a relevant possibility that someone might care about that topic a great deal. Second, the claim that, for any belief or topic, we might find ourselves in a situation where someone who cares a great deal about that belief or topic depends on us for information about that belief or topic. These two claims are not obviously true, yet I will not examine them here in any detail. Now notice what Grimm concludes from these two claims: we have a normative requirement not to be *cavalier* when we form our beliefs, and we have a normative requirement to treat any belief with *due respect*. These are, of course, rather vague conclusions and it is unclear to me how Grimm would prefer to make them more precise. I therefore do not intend what follows as an interpretation of his own views. What is clear, however, is that one way of making these conclusions more precise produces a new moral argument for (DE). Call this the *Grimm-inspired moral argument* for (DE):

The Grimm-Inspired Moral Argument

1. It is possible that there is someone who cares a great deal about whether p and who will at some time depend on us for information about whether p.

2. If it is possible that there is someone who cares a great deal about whether p and who will at some time depend on us for information about whether p, then we have a normative requirement to do our best to avoid believing that p if p is false.

3. So we have a normative requirement to do our best to avoid believing that p if p is false.

4. Doing our best to avoid believing that p if p is false is constituted by forming and maintaining our beliefs in accordance with our evidence.

C. So we have a normative requirement to form and maintain our beliefs in accordance with our evidence.

Here we take “doing one’s best to avoid believing that p if p is false” as an interpretation of what it means to refrain from being cavalier about our belief formation and of what it means to treat beliefs with due respect. Besides premises (1) and (2), notice, the Grimm-inspired moral argument is identical to the theistic and the consequentialist moral arguments that we find in Locke and Clifford. Premises (1) and (2), that is, attempt to offer a defense of the claim that we have a normative requirement to do our best to avoid believing that p if p is false, yet one that does not appeal to either God’s binding desires or to the negative moral consequences of certain actions.⁸

⁸There is a consequentialist reading of premise (2) that makes this last claim false. According to this reading, what makes us required to be good sources of information are the negative moral consequences of the action of providing bad information. This is neither the appropriate reading of Grimm’s intentions—to my mind—or the sense of premise (2) that I intend here. I take it instead as the claim that there is something *disrespectful*

This defense, however, fails. Premise (2) is once again false. It reflects a mistaken picture of the normative requirements that we have in virtue of the interests and needs of others. The picture is the following: we are each required to put ourselves in a position where we can best assist anyone who may need our assistance in promoting their (non-immoral or otherwise conflicting) interests. But this seems excessive. It places, in fact, an intolerable moral burden on each of us. There is an incredible variety of possible ways in which the many interests and many needs of others may come to somehow depend on us. In fact, there is simply no consistent combination of actions and attitudes that amounts to putting ourselves in a position where we can best assist all of these potential dependencies. Perhaps there is *some reason* in favor of putting ourselves in a position where we can best assist each of these potential dependencies. But the claims that we are *required* to put ourselves in a position where we can best assist anyone who may need our assistance in promoting their (non-immoral or otherwise conflicting) interests entails a widespread and inescapable proliferation of normative dilemmas. This gives us good reason to reject the picture of the normative requirements that we have in virtue of the interests and needs of others that is reflected in premise (2). This also gives us good reason to reject premise (2) itself.

There is, nonetheless, something that rings true about Grimm's claims. There seems to be some important connection between belief, testimony, and

about providing someone with bad information, whether or not there are negative moral consequences to it. In this sense, premise (2) is akin to a Kantian appeal to the *inherent dignity* of individuals.

the respect we owe to individuals. But we can capture what is true and important in the vicinity without accepting premise (2). While Grimm is mistaken when he says that “as a potential source of information for others, we have an obligation to treat any topic or any question with due respect,” it seems true that as an *actual* source of information for someone, we have an obligation to treat *that someone* with due respect. This is not an intolerable moral burden. Though there is an incredible variety of possible ways in which the many interests and many needs of others may come to somehow depend on us, there is only a small amount of actual such dependencies at any given time. In fact, the structure of the requirement that this reformulation of Grimm’s suggestion yields shows that here there is a consistent combination of actions and attitudes that consists in treating every person with an actual dependence on us with due respect.

Here I take a page from Brogaard’s reply to Clifford. While Grimm claims that we have a normative requirement to be good sources of information, I claim that we have a normative requirement not to be bad sources of information. But the latter requirement can be satisfied in two different ways: we can either be good sources of information or we can refrain from being a source of information in the first place. This suggests that the difference between Grimm’s claim and mine is also a difference in the scope of premise (2). While Grimm sees the respect we owe to individuals as giving rise to a narrow-scope normative requirement, I see the same bond as giving rise to a wide-scope requirement instead:

(narrow₂) If it is possible that there is someone who cares a great deal about p and who will at some time depend on us for information about whether p, then we have a normative requirement to do our best to avoid believing that p if p is false.

(wide₂) If it is possible that there is someone who cares a great deal about p and who will at some time depend on us for information about whether p, then we have a normative requirement to be such that, if we inform someone about p, then we are a good source of information about p.

If (narrow₂) is true, then the Grimm-inspired moral argument goes through. But if (narrow₂) is false and (wide₂) is true instead, then the argument is unsound. It does not follow from (wide₂) that we have a normative requirement to do what puts us in a position to be good sources of information: doing our best to avoid believing that p if p is false. What follows from (wide₂) is merely that we have a normative requirement to *either* do what puts us in a position to be good sources of information with respect to whether p *or* refrain from being a source of information about whether p at all.⁹

⁹One may worry that in trying to avoid the demandingness of Grimm's view we have swung too far towards the opposite extreme. On the suggested alternative, our epistemic normative requirements may seem too easy to fulfill. In particular, the alternative seems to lose an apparent virtue of Grimm's view, namely, the fact that the interests of *future* people give rise to normative requirements that bind us *now*. This worry, as I see it, misconstrues the proposed alternative. The requirement to be such that, if we inform someone about p, then we are a good source of information about p, does bind us now and, moreover, it arises due to the possibility that someone who cares about p might come to depend on us in the future. On this matter, the Grimmean view and the alternative are in agreement; they differ merely in the nature of the requirement that they identify. I am grateful to an anonymous reviewer for bringing this worry to my attention.

Just as Brogaard's reply to Clifford, of course, my claims here do not show that we *do not* have a normative requirement to do our best to avoid believing that p if p is false. What they show—if successful—is rather that a defense of this claim by way of premises (1) and (2) of the Grimm-inspired moral argument is unsound. Since it is implausible that the respect we owe to individuals gives rise to an intolerable moral burden, (wide_2) is false. And premise (3) does not follow from the more plausible (wide_2). Is there a defense of premise (3) that does not appeal to God's binding desires, (narrow_1), or (narrow_2)? I will consider and reject one such suggestion next.

4.4 The Epistemic Point-Of-View Argument

Richard Feldman (2000, 679) accepts that we ought to form and maintain our beliefs in accordance with our evidence. Moreover, he sees himself as showing that Locke's evidentialism can be separated from Locke's own apparent commitment to doxastic voluntarism (cf. Feldman 2001, 89-90), and as showing that Clifford's evidentialism can be separated from Clifford's own moral defense of it (cf. Feldman 2006, 20). It seems we have good reason to take Feldman as attempting to provide new and better grounds for (DE). In this section, I want to clarify the argument that he offers. The argument appeals to the epistemic point-of-view, to what's valuable from that point-of-view, and to a

certain sense of English ‘ought’. I will first clarify each of these three elements before stating his argument with some precision.¹⁰

I have an uncle who is a teacher. I also have an uncle who is a businessman. With some abstraction, we can say that I have an uncle who at times *plays the role* of the teacher, and that I have an uncle who at times *plays the role* of the businessman. With a bit more abstraction, we can say that there is a way to see things from the *teacher point-of-view*, and that there is a way to see things from the *businessman point-of-view*. We can make sense, for example, of a school administrator who says:

“I used to be a teacher, you know? So, from the teacher point-of-view, I see that we need smaller classrooms. But now I am a businesswoman. And, from the businesswoman point of view, I see that we need our classrooms to get even bigger.”

There is a tacit appeal here to the different roles that one can play in life and an explicit appeal to the point-of-views that we take from within these roles. At any rate, I take it that we have an intuitive grasp of what this administrator means.

In the sense just mentioned, it seems we each play a variety of roles in life. Some of these roles, as the two examples just above suggest, correspond to our professions. But not every role that we play are jobs. Some of us play the role of the husband, for example, and some of us play the role of the father. All of us, however, play the role of *the believer*. That is, we are all engaged in the activity of forming, maintaining, and revising our beliefs. Accordingly, just as

¹⁰My discussion here puts together as a unified picture the claims we find in Feldman (2000, 676), Feldman (2001, 87-9), and Feldman (2008, 349-52).

with every role, there is a way to see things from that point-of-view. This is the *believer's point-of-view*. If we have an intuitive grasp of what it means to make claims about the teacher's point of view, perhaps we have an intuitive grasp of what it means to make claims about the believer's point-of-view as well.

We can move from an understanding of the believer's point-of-view to an argument for (DE) once we accept three substantive principles. The first is the claim that there is a correct way to perform each role. We can call this the *correctness* principle:

(correctness) For each role, there is a correct way to perform that role.

Take the role of the teacher. It seems there is a correct way to perform that role, such that one can perform it badly and even try but fail to perform it in the first place. A teacher who grades her student's math work on the basis of how many times her favorite number is mentioned, for example, is not performing well in the role of the teacher. Take the role of the businessman. It seems there is a correct way to perform that role as well. A businessman who routinely sells his product for less than what it costs to produce it, for example, is not performing well in the role of the businessman. The same is true, of course, of the role of the believer. It seems there is a correct way to perform that role as well.

The second principle answers a very natural question: what determines the correct way to perform a certain role? The answer is that the correct way to perform a role is determined by what is valuable from the point-of-view of that

role. For each role R, that is, there is an associated notion of R-value. We can determine the correct way to perform a certain role R by examining which actions are most conducive to the things or states that are R-valuable. We can call this the *value* principle:

(value) The correct way to perform some role R is determined by what is R-valuable.

Take the role of the teacher once again. Some things and states are valuable from the point-of-view of the teacher—they are teacher-valuable. Perhaps these include explaining things clearly, being patient, being unbiased, being a fair grader, and so on. The correct way to perform the role of the teacher, then, is determined by reference to these teacher-valuable actions and states. This is why grading a student’s math work on the basis of how many times your favorite number is mentioned counts as performing badly in the role of the teacher. The same is true of the role of the believer. There are actions and states that are valuable from the point-of-view of the believer—actions and states that are believer-valuable—and these are the actions and states that determine the correct way to perform the role of the believer.

We are almost in a position to see how we can move from an understanding of the believer’s point-of-view to an argument for (DE). What is missing is a third substantive principle, this time about a certain sense of the English ‘ought’. It seems there is a sense of ‘ought’ that is used to indicate correct performance in a certain role. These are what we can call *role oughts*. Ought-claims that employ role oughts—claims of the form “S ought_R to ϕ ,” where

‘ought_R’ indicates the employment of a role ought—are made true by what is valuable from the point-of-view of the relevant role. We can call this the *ought* principle:

(ought) The claim ‘S ought_R to ϕ ’ is true iff ϕ -ing maximizes what is R-valuable.

Take the role of the teacher for one last time. Consider the claim that a teacher ought_R to give good lectures. If this is true, it is true because giving good lectures maximizes what is valuable from the teacher’s point of view. Similarly, consider the claim that a businessman ought_R to make profitable deals. If this is true, it is true because making profitable deals maximizes what is valuable from the businessman’s point of view. In precisely this sense, the claim that a believer ought_R to ϕ is true depending on whether ϕ -ing maximizes what is valuable from the believer point-of-view.

Consider how Feldman (2008, 351) puts all of this together:

There is a sense of ‘ought’ that depends upon the proper carrying out of a particular role... What I ought to do in my role as a teacher is give interesting lectures and grade in an unbiased way... Similarly, I think, as a believer I ought to follow my evidence. That is the right way to carry out that role.

I will follow Feldman in referring to the point-of-view of the believer as the *epistemic* point-of-view. I will refer to relevant sense of ‘ought’ as the *epistemic* role ought: ought_{ER}. The heart of Feldman’s defense of (DE) can thus be stated as the following claim: we ought_{ER} to form and maintain our beliefs in accordance with our evidence since forming and maintaining beliefs according to our evidence maximizes what is valuable from the epistemic point-of-view.

Since Feldman takes himself as rescuing Locke's and Clifford's evidentialism from doxastic voluntarism and from appeals to morality, respectively, we have reason to believe that Feldman takes epistemic role oughts to express our *sui generis* epistemic normative requirements.¹¹ Call this the *point-of-view argument* for (DE):

The Point-of-View Argument

1. We are all performing the role of the believer.
 2. If S is performing the role of the believer, then S ought_{ER} to form and maintain one's beliefs according to one's evidence.
 3. So we ought_{ER} to form and maintain our beliefs according to our evidence.
 4. If we ought_{ER} to form and maintain our beliefs according to our evidence, then we have a normative requirement to form and maintain our beliefs in accordance with our evidence.
- C. So we have a normative requirement to form and maintain our beliefs in accordance with our evidence.

With appropriate restrictions on the quantifier, perhaps premise (1) is true. Premise (2), however, requires some careful defense. Feldman's defense of it

¹¹In this way, Feldman rejects his earlier claim (Feldman 1988, 240-3) that epistemic ought-claims should be understood in the same way as we understand the 'ought' in 'we ought to pay our mortgage'. As he sees it, the latter is a case of a *contractual* ought, yet there is no sense in which we are bound by any contract (explicit or implicit) to believe in a certain way.

turns on the claim that *rationality* is what is fundamentally valuable from the epistemic point-of view, and on his claim that rationality is a matter of forming and maintaining one's beliefs according to one's evidence.¹² This is a controversial and substantive claim, yet one that I will not examine here in any detail. My interest is rather on premise (4). It attempts to capture the connection between role oughts and normative requirements. In the next section, I will discuss very briefly Kornblith's argument that premise (4) is false. As we will see, considerations of scope are once again recruited against (DE)—this time, however, without success.¹³

4.5 Kornblith's Incomplete Reply to Premise (4)

Kornblith (2001, 238) claims that role oughts lack *normative force*. I take this to mean that role oughts do not “have a grip” on us, do not state something that is “demanded” of us, do not involve the “binding sense of should,” and do not indicate anything about our “reasons.” The relations expressed by role oughts, that is, do not have the kind of *authority* over us that is characteristic of relations of normative requirement. If this is right, however, then premise (4) is false. Nothing follows about our normative requirements from the mere fact that the kind of relation expressed by a role ought obtains. I think this is exactly right. But Kornblith's explanation of why this is the case is mistaken.

¹²See Feldman (2008, 346-7).

¹³See Jon Altschul (2014, 252-254) for the claim that, contra Feldman, role-oughts are not counterexamples to the principle that ‘ought’ implies ‘can’. I will not discuss this worry here.

Though Kornblith correctly identifies *what* is wrong with Feldman's defense of (DE), he nonetheless provides an incorrect explanation for *why* this is so. His reply to Feldman, as I will put it, is incomplete.

According to Kornblith (2001, 237), role oughts do not carry normative force because they do not *detach*:

If one wants to be a good tyrant, perhaps one ought to be particularly brutal. Nevertheless, we would not say, even of someone who did in fact want to be a good tyrant, that he ought to be particularly brutal... This contrasts with the epistemic case, where we not only want to say that if someone wants to be a good believer, he or she should believe in certain ways; we also wish to endorse the claim that individuals ought, without qualification, to believe in those ways which, as a matter of fact, flow from good performance of the role of being a believer. Since being a tyrant or a con artist or a thief is just as much of a role as being a believer, what is the relevant difference here that allows us to detach the ought judgment in the case of believers, but prevents us from detaching the ought judgment in the case of tyrants, con artists, and thieves?

Kornblith here is not denying the existence of role oughts. He is not denying that, if one is a tyrant, then one *ought_R* to be particularly brutal. Kornblith is instead claiming that role oughts such as this, even if true, do not indicate that those who satisfy the antecedent thereby have a normative requirement to what is stated in the consequent. This, however, is just to say that role oughts give rise to wide-scope, instead of narrow-scope, normative requirements. Let '□' represent a normative requirement, let 'T' stand for 'playing the role of the tyrant', and let 'C' stand for 'being particularly cruel'. Now consider for comparison:

(tyrant_{Narrow}) If we are playing the role of the tyrant, then we have a normative requirement to be particularly cruel. $\forall x (Tx \rightarrow \Box Cx)$

(tyrant_{Wide}) We have a normative requirement to be such that, if we are playing the role of the tyrant, then we are particularly cruel. $\forall x \Box(Tx \rightarrow Cx)$

If the normative requirements produced by role oughts have wide-scope structure, then Kornblith is correct in suggesting that they do not detach. We cannot infer that S has a normative requirement to be particularly cruel from (tyrant_{Wide}) and the claim that S is playing the role of the tyrant. That is, though the following inference is valid:

1. $\forall x (Tx \rightarrow \Box Cx)$
2. Tx
- C. $\Box Cx$

The following inference *is not*:

1. $\forall x \Box(Tx \rightarrow Cx)$
2. Tx
- C. $\Box Cx \#$

But if this is true of role oughts in general, as Kornblith suggests, then premise (4) of the point-of-view argument is false and the inference from (3) to (C)

invalid. Instead of (4), the connection between role oughts and normative requirements is best captured by the following:

4*. We have a normative requirement to be such that, if we are playing the role of the believer, then we form and maintain our beliefs in accordance with our evidence.

And it does not follow from (4*) that we have a normative requirement to form and maintain our beliefs according to our evidence, even if we cannot help but satisfy the antecedent. In this *non-detachable* sense, Kornblith seems correct in claiming that role oughts in general, and epistemic role oughts in particular, seem to lack normative force.¹⁴

But this is not quite right. Wide-scope ought-claims can have normative force despite their non-detachability. The difference between wide and narrow scope is merely that the former gives the relevant agent two ways of satisfying their normative requirements: either explain things clearly or cease being a teacher; either be particularly cruel or cease to be a tyrant; either believe according to your evidence or cease being a believer. Recall the earlier example:

(tyrant_{wide}) We have a normative requirement to be such that, if we are playing the role of the tyrant, then we are particularly cruel. $\forall x$
 $\Box(Tx \rightarrow Cx)$

Though, as mentioned, the following inference is invalid:

¹⁴See Brogaard (2014) for the appeal to wide-scope requirements in a maneuver against Clifford's defense of (DE).

1. $\forall x \Box(Tx \rightarrow Cx)$

2. Tx

C. $\Box Cx \#$

The following inference *is not*:

1. $\forall x \Box(Tx \rightarrow Cx)$

C. $\forall x \Box \neg(Tx \wedge \neg Cx)$

Even if we grant Kornblith’s criticism—that role oughts only give rise to wide-scope normative requirements—we can still say, for example, that we have a normative requirement not to be a tyrant who is not particularly cruel. Though role ought claims do not detach, that is, this fact alone does not show that they fail to make authoritative demands on us. It can be binding that we satisfy our requirements in one of the two ways; it can be binding that we do not *both* play a role and fail to act as we ought_R. Feldman is still allowed the (weaker) claim that we have a normative requirement not to be a believer who does not form and maintains his beliefs in accordance with his evidence. So though Kornblith’s reply to Feldman may be sufficient for showing that his defense of (DE) fails—since the non-detachability of role oughts is sufficient for showing that premise (4) is false—there is a gap between that reply and the explanation for it that is on offer—the claim that role oughts lack normative force. In the next section, I offer an alternative explanation for why premise (4) is false that leaves no similar gap. It shows that premise (4) is false precisely because role

oughts lack normative force, whatever the logical structure of the normative relation they express.¹⁵

4.6 Normative Requirements and Privileged Values

No doubt the English ‘ought’ is ambiguous in various ways. Sometimes an ought-claim indicates something about what is likely to occur, as in “your ankle ought to heal in two weeks time.” Sometimes an ought-claim indicates what would be ideal, as in “someone ought to volunteer at the shelter.” Sometimes an ought-claim indicates the best way to achieve a certain end of ours, as in “you ought steal that car in order to escape from the police officers who are chasing you.” And so on.¹⁶ So perhaps Feldman is right that there is also a sense of ‘ought’ where it indicates the correct way to perform a certain role. But not all ought-claims express something about the relation of normative requirement that can hold between an individual and a certain action or attitude. This is perhaps trivial if we consider ought-claims that are not about individuals and ought-claims that are not about anyone in particular. But these are not the only cases. Some ought-claims of the form ‘S ought to ϕ ’ can be true of

¹⁵In later work, in fact, Kornblith (2002, ch.5) may well be relying on the normative force of wide-scope normative requirements. This is because he argues that epistemic normativity is regular instrumental normativity where the antecedent is always satisfied. Since it is an open and lively question whether instrumental normativity has a narrow-scope or a wide-scope structure (see the aforementioned Kolodny (2005) and Broome (2007), for example), Kornblith’s own account of epistemic normativity may well be committed to wide-scope structures being capable of normative force.

¹⁶This is not to suggest these various senses of the English ‘ought’ betray a difference in semantic structure and/or syntactic behavior. See, e.g. Finlay and Snedegar (2014).

someone S and a certain action or attitude ϕ without thereby expressing that ϕ -ing is required of S.

It will be useful to distinguish, in general, between ought-claims that do and ought-claims that do not express a claim about a relation of normative requirement. I will call the former *prescriptive* ought-claims and the latter *evaluative* ought-claims. We can say that evaluative ought-claims express a mere evaluation of a certain state of affairs, given a certain guiding value. More exactly, evaluative ought-claims express that a certain state of affairs obtains in those possible worlds ranked highest by a certain value. Similarly, we can say that prescriptive ought-claims express evaluations in this way as well. But these are not *mere* evaluations. Prescriptive ought-claims, instead, express that a certain state of affairs obtains in those possible worlds ranked highest by a certain *privileged* value: a value that has normative authority over us. What distinguishes ought-claims of the form ‘S ought to ϕ ’ that express a normative requirement, then, is the kind of value that guides the evaluation. Only those that are guided by privileged values “have a grip” on us, state something that is “demanded” of us, involve the “binding sense of should,” and indicate something about our “reasons.” None of this, notice, makes any reference to scope.¹⁷

¹⁷I take this to be a Moorean point. In the second preface to the *Principia*, Moore (1903, 3) says: “it cannot be too emphatically insisted that the predicate which... I call ‘good’, and which I declare to be indefinable, is only one of the predicates for which the word ‘good’ is commonly used to stand... [T]he predicate I am concerned with is that sense of the word ‘good’ which has to the conceptions of ‘right’ and ‘wrong’ a relation which makes it the sense which is of the most fundamental importance for Ethics.” I am here suggesting something similar regarding ‘ought’, values, and normative requirements. Thanks to Miles Tucker for pointing this out to me.

Admittedly, it is not easy to give an account of what makes a certain value privileged, and thereby of what gives an ought-claim normative authority over us. Without some such account, perhaps the distinction between evaluative and prescriptive ought-claims that I am suggesting is not entirely clear. (I will have more to say about privileged values below.) Nonetheless, it is quite easy to see that ought-claims employing role oughts are paradigmatic examples of evaluative ought-claims. Consider an illustration. Suppose a father tells his son “you ought to be a Patriots fan.” Suppose the son has no interest in sports in general and so demurs. There seems to be a clear sense in which it is not at all true that the son ought to be a Patriots fan: there is no normative requirement for him doing so. But now suppose the father explains himself in the following way:

“Listen, you were born in Boston. You have no choice but to play the role of the New Englander, though you can do it poorly or do it well. Performing well in this role, however, requires being a Patriots fan.”

The father is here explaining what he meant by his original claim. He intended to make a claim about a role ought. He has identified what maximizes *New England value*—being a Patriots fan—and has inferred that one therefore *New-England-role-ought* (ought_{NER}) to be a Patriots fan. If we grant that the father is correct about what is valuable from the point-of-view of the New Englander (and the details about this are irrelevant for the present point), and grant that the ‘ought’ in his utterance was the New-England-role-ought, then we must also grant that his claim was true: his son ought_{NER} to be a Patriots fan. Nonetheless, our initial assessment remains unaltered: there is no normative

requirement for the son to be a Patriots fan. So his father's claim, though true, must not be a claim about normative requirements. New England value, that is, is not a privileged value.

One may worry that the role of the New Englander is a gerrymandered sort of role in a way that the role of the teacher or the believer is not. (I myself find it hard to see a principled distinction here.) But similar illustrations can be multiplied. Suppose that father and son are robbing a bank. Suppose the father tells the son "you ought to threaten the life of the security guard." Suppose the son has no interest in threatening anyone's life and so demurs. Once again, there seems to be a clear sense in which it is not at all true that the son ought to threaten the life of the security guard: there is no normative requirement for him doing so. But now suppose the father explains himself in the following way:

"Listen, you are currently robbing a bank. You have no choice but to play the role of the bank robber, though you can do it poorly or do it well. Performing well in this role, moreover, requires threatening the life of the security guard."

The father is here explaining what he meant by his original claim. He intended to make a claim about a role 'ought'. He has identified what maximizes *bank-robbing value*—threatening the life of the security guard—and has inferred that one therefore *bank-robbing-role-ought* ($\text{ought}_{\text{BRR}}$) to threaten the life of the security guard. If we grant that the father is correct about what is valuable from the point-of-view of the bank robber (and again the details about this are irrelevant for the present point), and grant that the 'ought' in his utterance was the bank-robber-role-ought, then we must also grant that his claim was

true: his son ought_{BRR} to threaten the life of the security guard. Nonetheless, our initial assessment once again remains unaltered: there is no normative requirement for the son to threaten the life of the security guard. So his father's claim, though true, must not be a claim about normative requirements. bank-robbing value, that is, is also not a privileged value.

If we recognize the role of the teacher and teacher-role-oughts, then I think we must also recognize the role of the bank robber and bank-robber-role-oughts. Either way, the point can be made quite generally. If we hold fixed some value x , whatever x is, then there is an easily definable sense of 'ought'—*ought_x*—according to which it is true that S ought _{x} to do what maximizes x . But not all of these evaluations—and not all of the useful ones, and not all of the ones salient in several different contexts, and so on—express normative requirements. As I'm suggesting, we have a normative requirement to perform a certain action, or to take on a certain attitude, only when doing so promotes a privileged value, a value that has normative authority over us. And ought-claims express these normative requirements only when they express evaluations that are guided by these privileged values. What the two illustrations just above are intended to show, then, is that role-values, *qua* role-values, are not privileged values: they do not have normative authority over us. Role oughts, therefore, express mere evaluations instead of normative requirements. If this is right, however, then premise (4) is false. Nothing follows about our normative requirements from the mere fact that the kind of relation expressed by a role ought obtains. This time, moreover, premise (4) is false precisely because role oughts lack normative force, whatever their scope and structure.

Perhaps Feldman disagrees. When speaking of the role of the believer, in particular, Feldman (2000, 676) says that “anyone engaged in this activity ought to do it right.” On one reading, Feldman is here claiming that, for any activity A, anyone engaged in A ought to do it right. If we take him here as making a prescriptive ought-claim—that is, as claiming that we have a normative requirement to do it right—then we can take him as disagreeing with me on the normative authority that role oughts have over us. This seems just false to me. New-Englander-role-oughts and bank-robber-role-oughts have no normative authority over anyone. I’m not sure what else to say to convince someone who thinks otherwise. On another reading, however, Feldman is here saying that there is something special about the *epistemic* role ought that distinguishes it from the kinds of role oughts that I’ve been considering. Here we take him as saying something specific about the role of the believer, namely, that we are required to perform it right. I find this suggestion much more plausible. As I see it, this is just to claim that epistemic value is a privileged value. Perhaps that is true. What my criticism aims to show, however, is that epistemic value, if truly authoritative over us, is not so by virtue of being a role-value. Defending the claim that epistemic values are privileged values, that is, requires a very different kind account.¹⁸

¹⁸Let me make five clarifications about my criticism of Feldman’s defense of (DE). First, I am not claiming that there are no epistemic role oughts. Second, I am not assuming that only the *moral* sense of the English ‘ought’ expresses a claim about normative requirements. Third, I am not assuming that there is a sense of ‘ought’ that expresses the notion of an *all-things-considered* ought, a claim about what is best once we take into consideration all true ought-claims. Fourth, I am not suggesting that *sui generis* epistemic value cannot give rise to *sui generis* epistemic normative requirements. Fifth, I am not ignoring Feldman’s (2000, 676) injunction that “it is our plight to be believers.” This seems to suggest a feature

None of this, however, suggests that Feldman (2008, 355) is mistaken when he concludes that he has identified a sense of ‘ought’ that can make certain deontological claims in epistemology true. For all I’ve said, it may well be true that we ought_{ER} to form and maintain our beliefs according to our evidence. But to the extent that one finds plausible that the claim that we ought to form and maintain our beliefs according to our evidence expresses a relation of normative requirement, then to that extent one should find implausible that Feldman’s point-of-view argument for (DE) goes through. The fact that we ought_{ER} to form and maintain our beliefs according to our evidence goes no distance towards showing that there is an authoritative relation of normative requirement that holds between each individual and the complex of actions and attitudes that constitute forming and maintaining one’s beliefs in accordance with one’s evidence. Feldman’s claims about role oughts, that is, go no distance towards his professed goal of placing the views of Locke and Clifford on newer and better grounds.

of epistemic role oughts that distinguishes it from more ordinary role oughts, such as the teacher role ought. Moreover, one might think that this distinguishing feature is enough to justify the claim that epistemic role oughts express normative requirements after all. This is not the case. On this point, Kornblith (2001, 237-8) said it best: “Many people are forced into horrible roles; they are put in positions over which they have no choice. Some are forced into slavery; others into prostitution. Much as they may have no choice about playing certain roles, we don’t want to say in these cases that, whatever the role, they ought to perform them well.”

4.7 Privileged Values and the Nature of Normativity

What distinguishes ought-claims expressing a normative requirement is the kind of value that guides the evaluation. Only those that are guided by privileged values “have a grip” on us, state something that is “demanded” of us, involve the “binding sense of should,” and indicate something about our “reasons.” My argument in the previous section has been that role-values, as a *kind* of value, are not privileged values. This means that indicating that some value X is a role value is not an adequate way of indicating that X is a privileged value—no more than indicating that someone’s name is Peter is an adequate way of indicating that he is American. My argument, however, leaves open whether any particular role-value is, in fact, a privileged value, even if not in virtue of the very fact of being a role value as well. More importantly, with respect to epistemic normativity, my argument leaves open whether Feldman’s basic axiology, though not his argument for it, is correct. Perhaps it is, in fact, the value of *rationality* that grounds our normative requirement to believe according to our evidence. Even if Feldman’s defense of this claim in terms of the epistemic point-of-view fails, perhaps rationality is a privileged value for some other reason nonetheless. Two natural questions suggest themselves: which values are, then, privileged? And why are they so? I don’t have the space here to answer these questions in the detail that they deserve, but I want to conclude by clarifying what it would take to answer them and by indicating the outlines of my own view.¹⁹

¹⁹Notice that my criticism of Feldman in the previous section is independent of any positive account of privileged values. So long as we have it clear that role values, as a *kind*

What needs explaining is why some values have authority over us while other do not. One option is to claim that the privileged status of some values as authoritative is simply a primitive. It is a brute fact that the value of beneficence gives rise to reasons that make demand on us, while the value of maleficence does not; it is a brute fact that the value of justice gives rise to reasons that make demand on us, while the value of injustice does not. It is unclear to me, however, whether this option offers any explanation at all. Another option, which I endorse, takes a *constructivist* route by claiming that there is an explanatory connection between the privileged status of some values and those things that we most deeply care about: privileged values are those that we would endorse under conditions of ideal reflection. On this account, these values give rise to reasons that make binding demands on us because they reflect who we really are.²⁰

Similarly, an account of epistemic normativity must explain the connection between epistemic requirements and privileged values. One option is identifying certain epistemic values as privileged—truth or rationality, for example—and justifying that identification. Yet what reason do we have for thinking that truth or rationality, like beneficence, make binding demands on us? I don't know of a good answer to this question.²¹ Another option is to connect epistemic requirements to non-epistemic privileged values. We have seen

of value, fail to have authority over us, claiming that X is a role-value will be an inadequate explanation of its authority.

²⁰For the constructivist approach to normativity, see Railton (1986), Lewis (1989), Smith (1994), Koorsgard (1996), and Street (2008).

²¹See Ahlstrom-Vij (2013) for discussion.

this option deployed already with Clifford and Grimm, where the privileged values giving rise to epistemic normative requirements were the non-epistemic values of welfare and dignity. I endorse a similar *instrumentalist* account of epistemic normativity, together with the constructivist account of privileged values outlined just above. Whatever the normative requirements we have with respect to the forming and maintaining of our beliefs, on my view, we have them as instrumental requirements for the fulfillment and advancement of those things that we most deeply care about.²²

4.8 Conclusion

All three arguments for (DE) that I've considered seem defective. Clifford's consequentialist moral argument depends on a mistaken view of the structure of the normative requirement that arises from the relation between belief and action. Similarly, the Grimm-inspired moral argument depends on a mistaken view of the structure of the normative requirement that arises from the relation between beliefs, testimony, and the respect we owe to individuals. Lastly, Feldman's point-of-view argument depends on a mistaken view of the authority that role-values have over us. Locke's deontological evidentialism—the claim that there is a relation of normative requirement that holds between each of us and the complex of actions and attitudes consisting in forming and maintaining

²²For the instrumentalist approach to epistemic normativity, see Stich (1990), Kornblith (2002), and Cowie (2014). For a complete articulation and defense of a constructivist epistemic instrumentalism, see [omitted for blind review]. I am grateful to an anonymous reviewer for the encouragement to add the brief positive remarks that compose this section.

our beliefs in accordance with our evidence—seems to remain hostage to his wildly unpopular claims about God’s binding desires. Perhaps it remains itself as a product of belief against the evidence.

CHAPTER 5

DEONTOLOGICAL EVIDENTIALISM AND THE PRINCIPLE THAT *OUGHT* IMPLIES *CAN*

5.1 Introduction

John Locke (1706, §34) is famous for prescribing a close connection between evidence and belief:

We should keep a perfect indifference for all opinions, not wish any of them true, or try to make them appear so; but being indifferent, receive and embrace them according as evidence, and that alone, gives the attestation of truth.

W.K. Clifford (1877, 70) is famous for prescribing quite the same:

It is wrong always, everywhere, and for anyone to believe anything on insufficient evidence.

On one natural reading of these and other passages, Locke and Clifford seem to endorse *deontological evidentialism*:

(DE) S ought to form and maintain S's beliefs in accordance with S's evidence.¹

¹Richard Feldman (2000, 679; 2008, 351) accepts (DE) as well. He argues, moreover, that (DE) can be separated from Locke's apparent commitment to doxastic voluntarism (cf. Feldman 2001, 89-90) and from Clifford's moral argument in its defense (cf. Feldman 2006, 20).

The English ‘ought’ here expresses an authoritative relation of normative requirement or obligation. Fully stated, (DE) is thus the claim that there is an authoritative relation of normative requirement that holds between each individual and the complex of actions and attitudes that constitutes forming and maintaining one’s beliefs in accordance with one’s evidence.² As some have put it in different contexts, (DE) sees this normative requirement as “having a grip” on us (cf. Korsgaard 1996, 44-6), as something that is “demanded” of us (cf. Street 2012, 44), and as something “utterly different from anything else in the universe” (cf. Mackie 1977, 38).

This paper resists a certain argument for (DE). The argument depends on two principles. The first principle is a particular account of the nature of the relation of normative requirement. Call it *normative rationalism*:

(NR) There is a normative requirement for S to ϕ if and only if (and because) there is most normative reason for S to ϕ .³

²The source of this normative authority is a matter of debate. Since the term ‘obligation’ is so often and so naturally associated with *moral* obligation, I will here give preference to the more neutral term ‘normative requirement’. I will elide the ‘authoritative’ qualifier throughout.

³This is one way of expressing the view made famous by W.D. Ross (1930). We find it more recently in Scanlon (1998), Broome (1999), Raz (2002), Skorupski (2010), and Parfit (2011). See Schroeder (2007, 10-15) for the difference between normative, explanatory, and motivating reasons. Notice that (NR) is neutral with respect to the nature of normative reasons themselves. It favors neither consequentialism nor Kantianism, neither objectivism nor subjectivism, neither realism nor anti-realism, and so on. Whatever the nature of normative reasons, (NR) simply claims that they are ontologically prior to, or are the grounds for, the relation of normative requirement that can hold between a certain individual and a certain attitude or action.

The second principle places a constraint on what can count as a normative reason for belief. Call it *only evidence*:

(OE) Consideration *c* is a normative reason for S to form or maintain a belief that *p* only if *c* is evidence that *p* is true.

From (NR) and (OE), it follows that only evidence that *p* is true can ground a normative requirement to believe that *p*. Support for (DE) follows from the addition of modest non-skeptical assumptions.

My focus here is on (OE). Why should we think that it is true? Suppose I am an average meteorologist with an interview scheduled with the local news. Suppose my chances of doing well in this interview would significantly increase if I formed the belief that I am the best meteorologist in town. This consideration—that believing that *p* would increase my chances of getting something that I want—seems to be a normative reason for forming the belief that I am the best meteorologist in town. Yet (OE) says this is not so. Since this result is not obviously correct, the appeal to (OE) in an argument for (DE) requires antecedent support for (OE) itself. In this paper, I discuss the relation between a recent argument for (OE) and the principle that *ought* implies *can*. I argue, more exactly, that anyone who antecedently accepts or rejects the principle that *ought* implies *can* already has a reason to resist either (a) one of the argument's premises or (b) the argument's role in providing support for (DE).

Here is how I proceed. In the first section, I explain how Nishi Shah (2006) and Ward Jones (2009) extract a constraint on what can count as a reason

for *belief* from a constraint on what can count as a reason for *action* that is familiar from the work of Bernard Williams. I then explain how this constraint serves as a premise in what I will call the Transparency Argument for (OE). In the second section, I defend the first direction of my argument: anyone who antecedently accepts that the principle that *ought* implies *can* is false already has good reasons to reject premise (1) of the Transparency Argument. In the third section, I defend the second direction of my argument: anyone who accepts that the principle that *ought* implies *can* is true already has good reasons for either rejecting premise (2) of the Transparency Argument or for denying the argument's role in providing support for (DE). In the concluding section, I discuss the conditions under which one can, in fact, accept both the premises of the Transparency Argument as well as its role in providing support of (DE). Nonetheless, I identify a serious worry about the possibility of these conditions obtaining.

5.2 Reasons, Deliberation, and Transparency

It is well known that Bernard Williams took there to be an essential connection between *reasons for action* and *deliberation*. This is because reasons are essential in the explanation of our intentional actions and because they could not play this explanatory role if they were not also essential elements of our process of deliberation. Williams' example of Owen Wingrave illustrates what he had in mind. As Williams (1981, 106) puts it, Owen's family wants him to join the army, though he himself has no such desire and in fact "hates everything about military life and what it means." This is supposed to be

a case where no amount of careful deliberation about his desires, projects, and commitments could motivate Owen to heed his family's wishes. There is simply nothing about Owen's own self which deliberation could identify as being furthered by that action. For Williams, this is tantamount to simply saying that there is no reason for Owen to join the army.⁴

There is some disagreement, however, about how to state Williams' claims with more precision. On one reading, Williams is here identifying an essential connection between normative reasons for action and *motivation*. If a certain consideration is a normative reason for me to ϕ , that is, then I must be capable of becoming motivated to ϕ after careful deliberation about what to do. Call this the *motivation* constraint on reasons:

(motivation) If consideration c is a normative reason for S to ϕ , then, as a result of careful deliberation, S can *be motivated to* ϕ .⁵

On another reading, Williams is here identifying an essential connection between normative reasons for action and *effective deliberation*. If a certain consideration is a normative reason for me to ϕ , that is, then I must be capable of not only becoming motivated to ϕ but also of ϕ -ing on the basis that consideration. Call this the *basis* constraint on reasons:

⁴Williams focuses on the contrast between claims of the form 'A has reason to ϕ ' and claims of the form 'there is a reason for A to ϕ .' He claims that any attempt to understand the latter expression in a way where it does not collapse into the former expression—in a way, that is, which expresses what he calls an *external reason*—sacrifices the essential connection between reasons and deliberation. I am here using 'there is a reason for A to ϕ ' neutrally throughout.

⁵See Garcia (2004, 233-5) and Smith (2013, 102) for examples of this reading.

(basis) If consideration c is a normative reason for S to ϕ , then, as a result of careful deliberation, S can ϕ on the basis of c .⁶

Notice the contrast between (motivation) and (basis). According to the former, careful deliberation must be capable of producing a certain kind of *attitude*; according to the latter, it must be capable of producing a certain kind of *action*. While both constraints on normative reasons are inspired by Williams' claim that there is an essential connection between reasons for action and deliberation, the former constraint is weaker than the latter.⁷

It will not matter here which of these constraints best captures what Williams himself had in mind. Nor will it matter that Williams' claims, and these principles, are controversial in themselves. What will matter, instead, is that only (basis) can be used as an argument for (OE). (The reasons for leaving (motivation) aside will become evident just below.) For our purposes, moreover, what will matter is a version of (basis) where it states a general constraint on normative reasons *for belief*:

(b-basis) If consideration c is a normative reason for S to form or maintain a belief that p , then, as a result of careful deliberation, S can form or maintain a belief that p on the basis of c .⁸

⁶See Shah (2006, 484) for an example of this reading.

⁷It is difficult to make sense of the modality of the 'can' in these principles. Part of difficulty bears a family resemblance to issues about deviant causal chains in the literature on intentional action and turns on understanding the nature of what Williams (1993, 35) later called a "sound deliberative route." I will leave these difficulties aside for now.

⁸See Shah (2006, 486) and Jones (2009, 149).

We can infer (b-basis) from (basis) if we take it that what is true of normative reasons for performing some action is also true of normative reasons for taking on some attitude, such as belief. I find this a rather plausible suggestion.⁹ Of course, some particular consideration may count as a normative reason for action while not counting as a normative reason for belief as well. Nonetheless, if a general principle is true of normative reasons for action, then whatever the elements that count as normative reasons for belief in particular, that principle is true of them as well. The argument for (OE) that I will consider, at any rate, depends on this assumption.¹⁰

We can see the argument from (b-basis) to (OE) once we turn our focus to empirical facts about our belief formation. No one denies that many of our beliefs are formed on the basis of wishful thinking. Sometimes a desire to win an argument produces a rather strong belief that one's interlocutor is misinformed. Sometimes a deep fear of failure produces the belief that everything will work out fine. Sometimes, in fact, the mere spatial-location of a consumer good—being the rightmost—seems to produce the belief that it

⁹I am not alone. See Kearns & Star (2009, 219-221), Turri (2011, 384), and Booth (2012, 511-512).

¹⁰Notice that (b-basis) allows one to resist the first premise of Cowie's (2014, 4007) argument for *instrumentalism* about the normativity of evidence: "there is reason to believe in accordance with one's evidence (or value in so believing) *because* it is of great practical utility" (my emphasis). His argument, roughly, is that those who find *other* sources for the normativity of evidence must explain the "striking coincidence" of these normative facts. If (b-basis) is true, however, then practical utility alone is not sufficient to ground the existence of a reason for belief. If (b-basis) is true, that is, then there is no coincidence to explain in the first place. In his discussion of attempts to brush away the existence of a coincidence, Cowie (2014, 4012-3) is remiss in not considering this possibility.

is better than otherwise identical goods.¹¹ In these and other cases, the force behind the formation and maintenance of our beliefs is not an indication that *p* is true. These beliefs, that is, are not formed on the basis of the evidence that we have for *p*. Yet it seems as if nothing but an indication of *p*'s truth—nothing but our evidence for *p*—can affect our *conscious* forming and maintaining of beliefs. Call this constraint *transparency*:

(**T**) If, as a result of careful deliberation, *S* can form or maintain a belief that *p* on the basis of *c*, then *c* is evidence that *p* is true.¹²

If I am offered a million dollars in exchange for forming the belief that the Moon is made of cheese, for example, I may become convinced that forming this belief is desirable. But I will not be able to consciously form this belief on the basis of its desirability. To consciously form a certain belief, it seems, I must have an indication that this belief is true. (**T**) seems to be an empirically observed psychological fact.

Together, however, (b-basis) and (**T**) entail (**OE**):

The Transparency Argument for (OE**)**

1. If consideration *c* is a normative reason for *S* to form or maintain a belief that *p*, then, as a result of careful deliberation, *S* can form or maintain a belief that *p* on the basis of *c*. (b-basis)

¹¹See Nisbett & Wilson (1977) for these and other well-known experiments. See Kornblith (2014, 20-26) for a detailed summary of the large body of empirical literature documenting the misfortunes of unreflective belief-formation.

¹²See Moran (2001) and Shah & Velleman (2005) for careful discussions of transparency.

2. If, as a result of careful deliberation, S can form or maintain a belief that p on the basis of c, then c is evidence that p is true. (T)

3. Consideration c is a normative reason for S to form or maintain a belief that p only if c is evidence that p is true. (1,2)¹³

The reason why (motivation) cannot be used in a similar argument should be clear. There is no psychological constraint that prevents a non-evidential consideration from producing a mere motivation for believing that p. As Pascal's (1670, 124) *wager* illustrates so well, becoming convinced that belief in God is prudentially best can motivate me into active attempts to obtain it. Pascal himself recognized this point and recommended disciplined church-going as an effective method for inducing belief in God. But being so convinced cannot, by way of conscious deliberation, itself produce the desired belief. The way in which the Transparency Argument is part of a longer argument for (DE) should be clear as well. If only evidence can give us normative reasons for belief, and if normative requirements are produced by the balance of undefeated normative reasons, then a normative requirement to believe that p can only obtain when there is undefeated evidence for p. In the relevant sense of 'ought', modest non-skeptical assumptions deliver that one ought to form and maintain one's beliefs in accordance with one's evidence.

¹³We find versions of this argument in Shah (2006, 487) and Jones (2009, 150).

5.3 What if OIC is False?

The sense of ‘ought’ that matters presently, recall, is the sense where it express that a relation of normative requirement obtains, that someone is obligated to perform a certain action or to take on a certain attitude. The version of the principle that *ought* implies *can* that is of present interest, then, is a version that places a certain control-constraint on when such facts about requirements can be the case. Roughly, the principle says that S ought to ϕ only if S *has a choice* as to whether or not to ϕ , or only if it is *up to S* whether or not S ϕ 's. Here are two corresponding endorsements (my italics):

By the time honored principle that “ought implies can,” one can be obliged to do A only if one *has an effective choice* as to whether to do A. (Alston 1988, 259)

If the fact is I am sad and I consider this proposition, then whether or not I accept it is simply *not up to me*; but then accepting this proposition cannot be a way in which I can fulfill my obligation to the truth, or, indeed, any obligation to try to bring about some state of affairs. (Plantinga 1993, 38)

These are different ways to express a familiar and widely accepted control-constraint on normative requirements. Putting together these ways of thinking about *ought* and *can*, we can say more carefully that the version of the principle that *ought-implies-can* which matters presently is therefore the following:

(OIC): If there is a normative requirement for S to ϕ , then S has control over whether S ϕ 's.

In this section, I will argue that anyone who accepts some of the most common arguments against (OIC) already has good reasons to reject premise (1) of the Transparency Argument: (b-basis). I will focus on two counter-examples.

One reason to think that (OIC) is false is the existence of *direct* counter-examples. Suppose Karen is a kleptomaniac. Suppose her kleptomania is so compulsive that whenever she has the opportunity to shoplift with perceived impunity she quite literally lacks the power to do otherwise. Suppose Karen even describes these cases by saying that “something takes over and I’m no longer in control.” If all of this is right, then according to (OIC) it is not the case that Karen is required to refrain from stealing, say, a new watch at Macy’s when she notices that she is alone. At that time, she has simply no control over whether she steals it. *We*, of course, would be required to refrain from stealing it in a similar situation. But not Karen. Karen flouts no normative requirement when she goes ahead and steals the watch. To some, this assessment of Karen’s situation seems straightaway false: of course Karen has a requirement to refrain from stealing the watch (cf. Blum 2000, 287). Saying otherwise makes hash of our intuitions and our linguistic and social practices. But one cannot reject this assessment without also rejecting (OIC). So those who think that Karen, in this case, has a requirement to refrain from stealing the watch thereby have a counter-example to (OIC).¹⁴

¹⁴Some direct counterexamples to (OIC), however, are more contrived. Graham (2011a, 345-6), for example, exploits cases where (i) “it is plausible that the moral permissibility of A’s ϕ -ing depends on the moral impermissibility of B’s ψ -ing,” and where (ii) “it is not plausible that rendering B incapable of refraining from ψ -ing would render A’s ϕ -ing morally impermissible.” If there are cases of which (i) and (ii) are true, then there are

Another reason to think that (OIC) is false is the existence of *indirect* counter-examples. Suppose Black, an evil neurosurgeon, can manipulate Jones in such a way that will ensure that Jones kills Smith. Black observes Jones and will only intervene if Jones is about to decide not to kill Smith. But Black never has to show his hand. Jones, for reasons of his own, decides to and does kill Smith. This is Frankfurt's (1969, 836) famous counter-example to the claim that being *blameworthy* for ϕ -ing requires having control over whether one ϕ 's. But this case doubles as an indirect counter-example to (OIC) if one accepts the following bridge principle:

(bridge) S is blameworthy for ϕ -ing only if S ought to refrain from ϕ -ing.¹⁵

This is a plausible principle. Why would someone be blameworthy for doing what one is permitted to do? Assuming a broadly Strawsonian approach, why should we *resent* someone who acts in perfect accordance with their normative requirements? If (bridge) is true, however, and if Frankfurt's famous case indeed establishes what it intended to establish, then (OIC) is false. So those who think that Black is blameworthy and think that (bridge) is true thereby have a counter-example to (OIC) as well.

Perhaps these are two good reasons to be suspicious of (OIC). My aim here is not defending them. Instead, my aim is showing that those who take these

counterexamples to (OIC). Graham goes on to identify one such case, and to defend various attempts to resist the assessment that (i) and (ii) are true of it.

¹⁵Fischer (2006, 25) uses this principle to argue against (OIC). See Graham (2011b, 6) for resistance.

considerations as reasons for rejecting (OIC) already have good reasons for rejecting premise (1) of the Transparency Argument. We can see this in three short steps.

First, notice that combining (NR) with a rejection of (OIC) is incompatible with a control-constraint on normative reasons in general. Given (NR) and a denial of (OIC), we can infer that it is possible for there to be most reason for S to ϕ even if S does not have control over whether S ϕ 's. If the cases of Karen the kleptomaniac and of Jones' Frankfurtian murder of Smith are cases of normative requirements without agential control, that is, then (NR) tells us that these are also cases of undefeated normative reasons without agential control. Yet if there is most reason for S to ϕ , then there is a consideration c which is a normative reason for S to ϕ —there being *most* reason, of course, just being the claim that there is an abundance of a certain thing, namely, normative reasons. So given (NR) and a denial of (OIC) we can rather infer that it is possible for there to be a consideration c which is a normative reason for S to ϕ even if S does not have control over whether S ϕ 's.

Second, notice that combining (NR) with a rejection of (OIC) is also incompatible with with (basis). Consider:

The Argument Against (basis):

1. It is possible that there is a consideration c which is a normative reason for S to ϕ and S does not have control over whether S ϕ 's.
2. If it is possible that there is a consideration c which is a normative reason for S to ϕ and S does not have control over whether S ϕ 's, then it

is possible that there is a consideration c which is a normative reason for S to ϕ and S cannot, as a result of careful deliberation, ϕ on the basis of c .

3. So it is possible that there is a consideration c which is a normative reason for S to ϕ and S cannot, as a result of careful deliberation, ϕ on the basis of c . (1,2)

The reasoning captured just above shows that premise (1), and if (3) is true, of course, then (basis) is false. Moreover, it is hard to resist premise (2).

Recall the cases that have convinced some that (OIC) is false. In Karen's case, she is a kleptomaniac who is required to refrain from stealing a watch even though she has no control over whether she steals it. Karen simply cannot refrain from stealing the watch in the situation that she is in. Once she finds herself with the opportunity to steal with impunity, she cannot help but steal. This very same case, however, is also a case where Karen cannot refrain from stealing the watch, as a result of careful deliberation, on the basis of a consideration in favor of not stealing it. In general, in fact, if Karen simply cannot ϕ at t , then Karen *a fortiori* cannot ϕ at t , as a result of careful deliberation, on the basis of a consideration for ϕ -ing. Perhaps this is because she cannot carefully deliberate about that consideration in her situation at all; perhaps this is because her careful deliberation about that consideration in that situation is ineffective. Either way, it remains true that when Karen's kleptomania takes over, Karen cannot refrain from stealing the watch, as a result of careful deliberation, on the basis of the consideration that there is for

not stealing it. But since Karen is nonetheless required to refrain from stealing the watch, there is, in fact, a normative reason to refrain from stealing it. So there is a normative reason for Karen to refrain from stealing the watch even though she cannot, as a result of careful deliberation, refrain from stealing the watch on the basis of that normative reason. The same is true, *mutatis mutandis*, in Black's case: if Black simply cannot refrain from killing Smith, then Black *a fortiori* cannot refrain, as a result of careful deliberation, from killing Smith on the basis of a consideration that there is for not killing him.¹⁶

Third, notice that the Argument Against (basis) is effective against (b-basis) as well. Recall that (b-basis) derives from (basis) under the assumption that if a general principle is true of normative reasons for action, then, whatever the elements that count as normative reasons for belief in particular, that principle is true of them as well.¹⁷ But the conclusion of the Argument Against (basis) is itself a general principle governing (at least) reasons for action. I see no principled way of claiming that (basis) gives grounds for (b-basis)—as the Transparency Argument requires us to say—while at the same time resisting that (3) similarly gives grounds for the following:

¹⁶Notice that even if one thinks that Karen is not required to refrain from stealing the watch, or that Karen does not lack control over whether she steals it, one can still be moved by this bit of reasoning in defense of premise (2). Premise (2), that is, captures my claim that those who *already* deny (OIC) already have a good reason to reject (basis) as well.

¹⁷This is different from saying that every true *ought-claim* is constrained by the same general principles. There are many counterexamples to the suggestion that every ought-claim requires control. Feldman (1988, 674-6) discusses *contractual* and *role* oughts, for example, and Chrisman (2012, 603) discusses oughts that express *state norms*. These ought-claims, however, are not claims about normative reasons.

3*. It is possible that there is a consideration c which is a normative reason for S to form or maintain a belief that p and S cannot, as a result of careful deliberation, form or maintain that belief on the basis of c .

If (3*) is true, however, then (b-basis) is false. If this is right, then those who antecedently reject (OIC) already have a good reason to reject (b-basis) as well. That is, those who antecedently reject (OIC) already have a good reason to reject premise (1) of the Transparency Argument for (OE).

5.4 What if OIC is True?

Not everyone denies (OIC). For our purposes, it is enough to see why some think that both of the counter-examples suggested above fall short. Wedgwood (2013, 76), for example, thinks that cases of compulsion such as Karen's are not in fact cases where Karen does not have the power to refrain from stealing the watch:

Such compulsions seem broadly similar to powerful addictions; and although such addictions are typically at least partial excuses or mitigating factors... it seems that they typically do not completely remove the agent's ability to choose otherwise.

According to Wedgwood, that is, Karen does have control over whether she steals the watch after all, even after having the opportunity to steal it with perceived impunity and feeling, in some weaker sense, compelled to do it. If Wedgwood is right, then Karen's case, and cases of compulsion and addiction in general, cannot be direct counter-examples to (OIC). This is one way to defend (OIC). Graham (2011b, 6), for another example, thinks that the bridge

principle required to go from an assessment of Black as blameworthy to a denial of (OIC) is false. This is because (bridge) is in tension with cases of *blameworthy permissible-doing*. Here is one of the cases Graham offers as illustration:

Unbeknownst to Caleb, Jack is about to mercilessly torture 15 innocent children to death. The only way to prevent Jack from doing so is to kill him. On a whim, Caleb kills Jack merely for the fun of it.¹⁸

According to Graham, this is a case where Caleb is blameworthy for killing Jack even though it is not the case that Caleb ought to refrain from killing Jack. If Graham is right, then Frankfurt cases such as Black's cannot be indirect counter-examples to (OIC): since one can be blameworthy for ϕ -ing without being required to refrain from ϕ -ing, showing that one can be blameworthy for something that is out of one's control *is not* tantamount to showing that one can be required to perform actions or to take on attitudes that are beyond one's control. This is another way to defend (OIC).¹⁹

Perhaps these are two good reasons to accept (OIC). My aim here, once again, is not defending them. Instead, my aim is showing that those who accept (OIC) already have good reasons for either rejecting premise (2) of

¹⁸Graham (2011b, 6-7) offers three more cases against (bridge). One of them aims to show the possibility of *blameworthy obligation-fulfilling*, and two of them aim to show that these cases do not depend on ignorance of any kind.

¹⁹Littlejohn (2012, 3) argues against Graham (2011a) that he has not identified a counterexample to (OIC). His argument turns on identifying an alternative explanation to the moral permissibility of the relevant action which does not depend on the moral impermissibility of the other relevant action (see fn. 14). His argument depends on accepting a certain principle about the legitimate use of violence on passive threats, and on rejecting Graham's rejection of that principle. This is yet another way to defend (OIC).

the Transparency Argument or for denying the argument's role in providing support for (DE). We can see this in three short steps as well.

First, notice that (OIC) places a control-constraint on normative requirements in general. But given the grounding of normative requirements on normative reasons provided by (NR) and the assumption that the principles governing normative reasons are uniform irrespective of their objects, there is no reason to resist a version of (OIC) where the variable ' ϕ ' picks out the specific attitude of *belief*. Similarly, there is no reason to resist a version of (OIC) where the variable ' ϕ ' picks out the complex of actions and attitudes that constitutes forming and maintaining one's beliefs in accordance with one's evidence:

(OIC*): If there is a normative requirement for S to form and maintain one's beliefs in accordance with one's evidence, then S has control over whether S forms and maintains one's beliefs in accordance with one's evidence.

If one resists the counter-examples that are offered against the control-constraint (OIC) imposes on normative requirements in general, that is, then one *ipso facto* resists similar qualms against the control-constraint that (OIC*) imposes on normative requirements for forming and maintaining beliefs in particular.

In the terminology we find in Alston and Plantinga, (OIC*) is the claim that the relevant normative requirement obtains only if we have a choice as to whether or not we form or maintain our beliefs in accordance with our evidence, only if it is up to us whether or not we form or maintain our beliefs in

accordance with our evidence. Do we have such a choice? Call an affirmative answer to this question *doxastic voluntarism*; call a negative answer to it *doxastic involuntarism*. We can now state the argument from this section in the form of a dilemma:

The Voluntarist Dilemma: Doxastic voluntarism is either true or false:

(Horn #1) If doxastic voluntarism is true, then (T) is false.

(Horn #2) If doxastic voluntarism is false, then (DE) is false.

Enlisting the Transparency Argument as support for (DE), it turns out, is in trouble either way.

Second, suppose someone accepts (OIC*) and accepts doxastic voluntarism. This is someone who accepts two claims: (a) we can have a normative requirement to form and maintain our beliefs in accordance with our evidence only if forming and maintaining our beliefs in accordance with our evidence is within our control, and (b) it *is* within our control whether we form and maintain our beliefs in accordance with our evidence. This seems to be a position Wedgwood (2013, 77-79) is willing to defend:

If a reasoner exercises the capacity to come to have a certain level of confidence in a proposition *p* at a certain time *t*, the reasoner must also have had the power not to come to have that level of confidence at that time... We can also make sense of the suggestion that we have control over our beliefs and intentions because we can exercise our reasoning capacities, and the way in which we exercise those reasoning capacities will determine what we believe and intend.

His aim, in fact, is precisely that of articulating a sense of control that can allow us to accept deontic claims about beliefs without having to abandon (OIC). On its face, of course, there is nothing objectionable about this position. But if one says that whether or not we form and maintain our beliefs in accordance with our evidence *is*, in fact, within our control, then one is thereby denying the alleged psychological fact of transparency, captured above by (T). (T) is precisely the claim that forming and maintaining our beliefs in accordance with our evidence *is not* a choice. When we form and maintain beliefs unconsciously, we have no control; when we form and maintain them consciously, we have no control over whether to follow the evidence. According to (T), that is, we are never in a position where it is up to us whether we believe in accordance with our evidence. So anyone who accepts (OIC*) and doxastic voluntarism already has a good reason to reject premise (2) of the Transparency Argument for (OE).

Third, suppose that someone accepts (OIC*) and accepts doxastic involuntarism instead. This is someone who accepts two claims: (a) we can have a normative requirement to form and maintain our beliefs in accordance with our evidence only if forming and maintaining our beliefs in accordance with our evidence is within our control, and (b) it *is not* within our control whether we form and maintain our beliefs in accordance with our evidence. This is finally a position which allows one to accept both premises of the Transparency Argument: (b-basis) and (T). Moreover, there seems to be nothing objectionable about this position on its face. But, given (OIC*), (DE) is true only if forming and maintaining one's beliefs in accordance with one's evidence is within one's

control. And, given doxastic involuntarism, it is not within our control to form and maintain our beliefs in this way. So anyone who accepts (OIC*) and doxastic involuntarism already has a good reason to reject (DE) itself. Perhaps surprisingly, the only position which allows one to accept both premises of the Transparency Argument for (OE) is a position which precludes the use of (OE) in an argument for (DE).

5.5 Conclusion

I have been examining the prospects of an argument for (DE) that turns on (NR) and (OE). Since (OE) is not particularly obvious, it requires rather substantive support. Yet the most clearly stated argument in its defense precludes the use of (OE) in an argument for (DE). The conjunction of three claims delivers this result: those who antecedently reject (OIC) already have good reason to reject premise (1) of the Transparency Argument; those who antecedently accept (OIC) and accept doxastic voluntarism already have good reason to reject premise (2) of the Transparency Argument; and those who accept (OIC) and accept doxastic involuntarism already have good reason to reject (DE) itself. We can now state my overall argument in this paper in the form of a trilemma:

The OIC Trilemma: (OIC) is either true or false:

(Horn #1) If (OIC) is false, then (b-basis) is false.

(Horn #2) If (OIC) is true and doxastic voluntarism is true, then (T) is false.

(Horn #3) If (OIC) is true and doxastic involuntarism is true, then (DE) is false.

This, I think, is bad news for the Transparency Argument, for (OE), and for (DE).

Yet this is not, perhaps, the final word. Perhaps the OIC Trilemma merely reveals the conditions required for accepting the Transparency Argument and its role in providing support for (DE). First, one must resist the common counter-examples to (OIC). Second, one must accept doxastic involuntarism. This allows one to accept both premises of the Transparency Argument: (b-basis) and (T), respectively. But one cannot accept (OIC). Since accepting premise (1) requires accepting doxastic involuntarism, accepting (OIC) with it would force accepting the denial of (DE). The third condition, then, is providing a novel rejection of (OIC) that is tailor-made for avoiding any tension with (basis). Nothing of what I said above excludes this possibility. Perhaps, that is, there is a way to reject (Horn #1).

This possibility, however, is problematic for the Transparency Argument. At the beginning of section 1, recall, I suggested that Williams' defense of (basis) turned on the connection between normative reasons and the explanation of our intentional actions. But this defense of (basis) has been criticized even by those who agree with William's constraint. The worry is simple: we can explain intentional actions by appealing to our *beliefs* about our reasons just as well as by appealing to the reasons themselves (cf. Anomaly 2008, 475). Normative reasons, that is, are explanatorily superfluous in this case. Despite the shortcomings of this defense of (basis), however, Williams-sympathizers

suggest that we can accept it on different grounds: by appealing to (OIC). Accepting (OIC), that is, allows us to say that if someone simply *could not* perform some action or take on some attitude after careful deliberation, then this person cannot as well have a normative reason for performing this action or for taking on this attitude. Consider Anomaly (2008, 476):

In the example [of Owen Wingrave] there is no discernible route from his current motivational set (or, if we wish, psychological profile) to some utterly different set in virtue of which Owen would be convinced that honor or tradition requires him to join [the army]. In the absence of such a route, we naturally conclude that he lacks a reason to join because he *cannot* be motivated by the considerations advanced by his father without becoming an utterly different person. (His emphasis)

I will not examine here whether Anomaly is correct in suggesting that we find this argument from (OIC) to (basis) in Williams. I will also not examine whether the first argument—based on the explanation of intentional actions—is defective in the ways that he and others have suggested. I mention these suggestions only to highlight the problem that they seem to raise. The third condition for accepting the Transparency Argument and its role as support for (DE) is providing a novel *rejection* of (OIC) that is tailor-made for avoiding any tension with (basis). Naturally, this will be a terribly thorny condition to satisfy if the best argument *for* (basis) depends on (OIC) itself.

CHAPTER 6

AMPLIATIVE TRANSMISSION AND DEONTOLOGICAL INTERNALISM

6.1 Introduction

The literature on the inferential transmission of justification has so far focused almost exclusively on competent deductive inferences: on cases where (i) p follows deductively from q and r , and (ii) S infers p from q and r because of awareness of that relation. Since competent deductions seem so secure, transmission *failure* has taken center stage.¹ But this narrow focus is misguided. Ampliative inferences present philosophically interesting cases of *successful* transmission instead.

Consider the following transmission principle discussed very briefly by Weatherson (2008, 568):

Blamelessness Ampliative Transmission (BAT):

If:

- (a) S believes that believing that p is justified (call this higher-order belief q)

¹See, for example, Wright (2003, 2004) and Davies (2004).

(b) S is epistemically blameless in believing that q

Then:

(c) If S forms the belief that p on the basis of q , then S is epistemically blameless in believing that p .

Condition (c) defines a certain ampliative inference that is claimed to transmit blamelessness whenever conditions (a) and (b) obtain. Though (BAT) seems harmless, Weatherson has argued that it causes a serious problem for versions of deontological internalism formulated in terms of epistemic *blame*: views where being justified in believing that p is being blameless in believing that p . I agree. But Weatherson denies that a similar argument can be pushed against versions formulated in terms of epistemic *praise*. Here I disagree. (BAT), I will argue, creates a serious problem for any version of deontological internalism where justification is a positive appraisal of someone's epistemic agency.

Here is how I proceed. After some clarifications (section 1), I mount a two-stage defense of (BAT): I argue that (BAT) is intuitively plausible (section 2), and I argue that the most common alleged causes of transmission failure do not apply to it (section 3). I then argue that my two-stage defense of (BAT) supports similar ampliative transmission principles focused on the related deontological notions of epistemic *praise* and epistemic *responsibility* (section 4). Finally, I consider Weatherson's argument from (BAT) and argue that his resistance to the soundness of similar versions is implausible (section 5).

6.2 Clarifying (BAT)

I begin by answering three clarificatory questions: What exactly does it mean to transmit? What exactly does the transmitting? What exactly gets transmitted?

Transmission is the phenomena where a belief that p has some epistemic property *in virtue of* a distinct belief that q having that same epistemic property. At a minimum, transmission requires causation: if my belief that q plays no causal role in the formation or maintenance of my belief that p , then my belief that q cannot transmit any epistemic properties to my belief that p . (BAT) is thus a claim about how a belief can have the epistemic property of blamelessness in virtue of a causally relevant belief having that epistemic property as well. This answers the first question.

I will understand the inference codified in “forms the belief that p on the basis of q ” as the mental act of (i) believing that q and (ii) having that belief as a partial cause of one’s belief in p . In the case of (BAT), this amounts to performing the following inference:

1. I am justified in believing that p .
2. Therefore p .

Such inferences are *level-lowering* since the believer infers the lower-level belief that p from the higher-level belief that she is justified in believing that p . Such inferences are *ampliative* since q does not entail p —I will assume throughout that justification does not entail truth.² So whenever someone accepts (2) on

²See Littlejohn (2012) for discussion and disagreement on this point.

the basis of (1)—in the sense just defined—I will say that she has performed a *level-lowering ampliative inference*. And whenever someone accepts (2) on the basis of (1) in the absence of any defeaters for (1), I will say that she has performed a *competent level-lowering ampliative inference*. (BAT) is thus the claim that blamelessness transmits through competent level-lowering ampliative inferences. This answers the second question.

Mirroring the distinction between warrant and doxastic justification in the literature on transmission, I will here distinguish between (i) the *ground* for blamelessness (that which makes it blameless for S to believe that *p* at *t*), and (ii) *blamelessly believing* that *p* (a belief that is appropriately connected to existing grounds for blamelessness). Just as one can have warrant for a certain belief and yet fail to be doxastically justified in believing it, one can similarly have grounds for blamelessly believing a certain belief and yet fail to be blameless in believing it. I will here take (BAT) as a claim about the transmission of blameless belief and not as a claim about the transmission of grounds for blamelessness.³ (BAT) is thus the claim that blameless belief transmits through competent level-lowering ampliative inferences. This answers the third question.

³I am here following Tucker (2010, 502-507). This is important since it allows for inferences that *do not* transmit warrant (or grounds for blamelessness) to sometimes count as transmission success by virtue of transmitting doxastic justification (or blameless belief). This will be relevant in section 3 below.

6.3 The Intuitive Defense of (BAT)

A subject satisfies conditions (a) and (b) of (BAT) when she is blameless in believing that she is justified in believing that p . How should we understand the notion of blamelessness at stake? Here it is helpful to turn to a hypothetical case sometimes used as an argument for internalism about justification. Suppose that John's inferential habits accord with all the true internal epistemic standards for belief formation and maintenance. Perhaps John always gathers as much evidence as it is reasonable to expect of him; perhaps John never believes what his total evidence does not support; perhaps John always believes what he believes on the basis of appropriate epistemic grounds; perhaps John even fully endorses these true internal epistemic standards; and so on. We can say that John is in this way *reflectively ideal*. Now suppose that an evil demon is ensuring that all of John's beliefs about the external world are nonetheless false. The argument for internalism about justification that I have in mind turns on the conviction that while John lacks knowledge (by having false beliefs), and lacks reliability (by only having processes that produce a preponderance of false beliefs), he is nonetheless justified in his false beliefs about the external world. He is justified, that is, by virtue of his ideal reflection.⁴

Here is how this case is relevant for our present purposes: whatever we say about John's epistemic justification for believing that, say, the sun is shining, it seems hard to deny that John is *blameless* in believing it. John, after all, is being deceived *despite his best efforts*: everything that is under John's control

⁴See Cohen (1984) and Pryor (2001), for example, for arguments of this sort. See Weatherson (2008, 564-567) for discussion.

is being executed to perfection; John is not being derelict in any of his epistemic duties; and so on. There are simply no grounds for blaming John for having that false belief. Quite the contrary, in fact: John's ideal reflection is the very ground for his blamelessness. While it is controversial whether or not justification has a substantive connection to truth—and, more exactly, how that connection should be understood—the connection to truth seems entirely beside the point when the evaluation in question is focused on *blame*. Only facts about the *agent* matter for this kind of deontological evaluation, and John seems beyond reproach.

Similarly, it seems hard to deny that John's ideal reflection would be just as appropriate for grounding higher-order beliefs. Suppose that John has been informed by an unreliable source that he is being deceived by an evil demon. Knowing that the source is unreliable, but wanting to be extra careful nonetheless (unreliable sources sometimes are right), John suspends his belief that the sun is shining and proceeds to examine whether he is justified in holding that belief after all. "I can't account for the behavior of evil demons," he thinks to himself, "but I can account for mine." After the kind of ideal reflection defined above, suppose John forms the belief that he is justified in believing that the sun is shining. "As far as I can tell," he concludes, "believing that *p* is entirely appropriate from my epistemic position." Just as ideal reflection seems enough to ground John's blamelessness in believing that the sun is shining, ideal reflection seems here enough to ground John's blamelessness in believing that he is justified in believing that the sun is shining.

A case of ideal reflection like John's (after being challenged and having confirmed the internal credentials of his belief), then, seems to be a case where conditions (a) and (b) of (BAT) are satisfied: John believes that he is justified in believing that the sun is shining and John is blameless in that higher-order belief. Now suppose that John performs the following competent level-lowering ampliative inference:

1. I am justified in believing that the sun is shining.
2. Therefore the sun is shining.

According to (BAT), John is now blamelessly believing that the sun is shining. This seems correct to me. John's ideal reflection, on its own, seemed sufficient to make him blameless in believing that the sun is shining. It seems odd to suggest that the same ideal reflection cannot make him blameless in believing that very same thing once we add a blameless higher-order belief as an intermediary step.

6.4 (BAT) and Transmission Failure

Put generally, an inference will fail to transmit blameless belief just in case the conclusion is not blamelessly believed in virtue of being based on a blamelessly believed premise. In John's case, we have already seen that the relevant premise—'I am justified in believing that the sun is shining' (q)—is, plausibly, blamelessly believed. What could then prevent q from transmitting this epistemic property to John's conclusion that 'the sun is shining' (p)? There

are four main alleged causes of transmission failure that are discussed in the literature. Neither, however, applies to (BAT).⁵

One alleged cause of transmission failure is *simple premise circularity*—an inference of the form “ x therefore z therefore x .” But John’s inference is clearly not premise-circular in this simple way. Another alleged cause of transmission failure is *failure to resolve doubt*. If John’s inference from q to p cannot resolve John’s doubts about p , that is, then perhaps it has failed to transmit blamelessness from q to p (cf. Wright 2003, 63). Whatever we think about the truth of this conditional, its antecedent seems false in John’s case anyway. In our story, John first suspends his belief that the sun is shining, but only to remove his doubts about it by way of his competent level-lowering ampliative inference. John’s inference seems in this way entirely capable of resolving doubt. This is a stipulation, of course, but there is no reason to think that John’s case is conceptually incoherent.

A more involved alleged cause of transmission failure is *having an inefficient structure*. This happens when the following two conditions are satisfied. First, there is some ground G that is sufficient by itself to make belief in premise q blameless and is also sufficient by itself to make belief in conclusion p blameless. Second, S nonetheless bases her belief that p in her belief that q , which is itself supported by G (cf. Wright 2002, 334). In our case, John’s inference

⁵Discussion of transmission failure has so far focused on the epistemic properties of *warrant* and *doxastic justification*. My discussion below adapts these claims to the present context where the relevant epistemic property is *blamelessness*. My references in the text should then be understood as indications of where to find similar claims about either warrant or doxastic justification, as opposed to where to find corroborations of the claims I am making about blame.

does, in fact, have an inefficient structure. John's ideal reflection is sufficient to make him blameless in believing that he is justified in believing that the sun is shining, as well as straightaway blameless in believing that the sun is shining. And yet John arrives at the latter via the former. So if having an inefficient structure is in fact a cause for transmission failure, then (BAT) is false.

Having an inefficient structure is not, however, cause for transmission failure. We can see this first by counterexample, following Tucker (2010, 512):

The perceptual evidence that warrants me in believing... that there are exactly twenty-five people in the room, seems to be an equal warrant for the conclusion... that there are fewer than one hundred people in the room. The deduction [in this case], therefore, has an inefficient structure... It nonetheless seems clear that it transmits (doxastic) justification to its conclusion.

This seems correct and, moreover, applicable to the epistemic property of blamelessness. What it brings out is the importance of distinguishing between transmission of grounds for blamelessness (or warrant) and transmission of blameless belief (or doxastic justification). An inference with ineffective structure cannot provide a ground for its conclusions since the ground for at least one of its premises is already itself ground for the conclusion. John's ideal reflection, for example, is already grounds for both q and p and John's inference is not what is providing p with those grounds. But inferences with inefficient structures can nonetheless provide the conclusion with an *appropriate connection* to the relevant grounds (cf. Tucker 2010, 513-4). John has grounds for blamelessly believing that p , but is only blameless in that belief when his belief is appropriately connected to those grounds for blamelessness. This is the sense in which John's inference transmits blameless belief without

transmitting grounds for blamelessness: it provides an appropriate connection between John's ideal reflection and his belief that the sun is shining. (BAT)'s inefficient structure, therefore, is no reason to think that it is false.

Consider one final alleged cause of transmission failure: *robust premise circularity*. This happens when one is justified in believing some premise in virtue of already being justified in believing the conclusion (cf. Tucker 2010, 214-16). But John's is not a case of this kind either. John's blamelessness for believing that q does not in any way depend on his blamelessness for believing that p . John's beliefs share a common ground for blamelessness—John's ideal reflection—but neither belief is itself a part of that ground.

6.5 Extending (BAT)

I have mounted a two-stage defense of (BAT) in sections 2 and 3. I've argued that (BAT) is intuitively plausible and that the most common properties suggested in the literature as reasons for transmission failure do not apply to it. It seems we have good reason to think that (BAT) is true.

(BAT), however, can be extended to the related deontological notions of epistemic *praise* and epistemic *responsibility*. Consider:

Praiseworthiness Ampliative Transmission (PAT):

If:

- (a) S believes that believing that p is justified (call this higher-order belief q)

(b₁) S is epistemically praiseworthy in believing that q

Then:

(c₁) If S forms the belief that p on the basis of q , then S is epistemically praiseworthy in believing that p .

and

Responsibility Ampliative Transmission (RAT):

If:

(a) S believes that believing that p is justified (call this higher-order belief q)

(b₂) S is epistemically responsible in believing that q

Then:

(c₂) If S forms the belief that p on the basis of q , then S is epistemically responsible in believing that p .

Since the most common properties suggested in the literature as reasons for transmission failure are all *formal* properties as opposed to properties of the *content* of the relevant inferences, it follows that these properties do not apply to (PAT) and (RAT) if they do not apply to (BAT). This means that my defense of (BAT) in section 3 serves as a defense of (PAT) and (RAT) as well. This is one reason for thinking that they are true.

(PAT) and (RAT), moreover, seem equally intuitive. Think of John once again. Everything that is under John's control is being executed to perfection;

John is not being derelict in any of his epistemic duties; and so on. I've already suggested that there are therefore no grounds for blaming John for having the false belief that the sun is shining. What grounds could there be for withholding praise or claiming that John is not being epistemically responsible? I can see no such grounds. John's ideal reflection, instead, seems quite appropriate as ground for his epistemic praiseworthiness and responsibility. In both these cases, as with blamelessness, the connection to truth seems entirely beside the point and the efforts of the agent seem to measure up. This is another reason for thinking that (PAT) and (RAT) are true.

6.6 Trouble for Deontological Internalism

Deontological internalism, as I will have it, is the claim that justification should be defined in agent-focused deontic terms. Many versions of deontological internalism so construed are currently in vogue.⁶ Weatherson, however, has argued that (BAT) creates a serious problem for versions formulated in terms of epistemic *blame* and has denied that a similar problem can be created for versions focused on epistemic *praise*. We are now in a position to see that Weatherson is mistaken in this last point.

Consider the following version of deontological internalism:

⁶See, e. g., BonJour (2010), Booth (2012), Petersen (2013), and Smithies (forthcoming). Notice that my definition of deontological internalism leaves out deontological views where justification is *permissible* belief. Some of these views are externalist (e.g. Goldman 2009), but some are internalist as well (e.g. Wedgwood 2012). So there is a sense of "deontological internalism" that is not my target here.

Blame (B): S is justified in believing that p iff S is blameless in believing that p .

Think of John one more time. (B) suggests that John's belief that the sun is shining is justified precisely because it is blameless. Now consider the following principle of *false reflection*:

False Reflection (F): It is possible for S to have a justified but false belief that her belief that p is justified.

Those who deny that justification and truth can ever come apart will naturally deny (F) as well. Those who do not deny that justification and truth can ever come apart will find it difficult to reject (F) in a way that is not *ad hoc*. Suppose that John considers whether to believe that p . He first reflects on what he takes as his evidential base; he then reflects on which inferential principles lead from his evidence to justified beliefs; he then finally reflects on whether there are any defeaters that are relevant in this situation. If justification and truth can sometimes come apart, then it is possible for John to have a justified but false belief about the contents of his evidential base, or about the relation between that base and p , or even about the presence of defeaters. *Human* ideal reflection, after all, is not *perfectly* ideal.⁷

Together, however, (BAT), (B) and (F) entail a contradiction (cf. Weather-son 2008, 568-9):

⁷Importantly, none of these mistakes need to be attributable to John, as long as the relevant mental states are *non-transparent*. See Williamson (2000, ch. 4) for an influential argument that all interesting mental states are non-transparent.

The Transmission Argument Against Blame:

1. S justifiedly, but falsely, believes that she is justified in believing p .
(Instance of F)
2. On the basis of this belief, S comes to believe that p .
3. S blamelessly believes that she is justified in believing that p . (1, B)
4. S blamelessly believes that p . (2, 3, BAT)
5. S is justified in believing that p . (4, B)
6. S is justified in believing that p and S is not justified in believing that p . (1, 5)

The Transmission Argument suggests that (B) should be rejected.⁸ If I'm correct about (PAT) and (RAT), moreover, then similar transmission arguments can

⁸One might worry that there is no contradiction between (1) and (5) since (1) is a claim about what S is justified in believing *before* performing the relevant level-lowering inference. Taken in this way, that is, (1) says nothing about what S is justified in believing *after* performing that inference. Let me make two comments about this worry. First, I take it that (1) is, in fact, a claim about what S is justified in believing after the relevant level-lowering inference. The appropriate paraphrase of S's justified belief in this case would be: "I'm looking at my reasons for p and I believe that believing that p is justified for me." When one forms this justified higher-order belief, that is, one does not yet have the lower-order belief as well. We can say that (1) is future-directed in this way. Second, notice that the level-lowering inference that S performs subsequently simply *connects* p to the grounds that S mistakenly believes are good enough. The level-lowering inference, that is, is not itself a part of those grounds. (See my discussion of inferences with ineffective structure in section 2.) This means that the justifying power of those grounds with respect to p stay the same throughout. This is important, otherwise (1) would be false. When taken in this way—where (1) is future-directed and where the subsequent inference does not alter the grounds for p —(1) is both a true instance of F and in tension with (5). I'm grateful to an anonymous referee for bringing this worry to my attention and to Chris Tucker for discussion of the issue.

be pushed against versions of deontological internalism formulated in terms of epistemic *praise* and epistemic *responsibility*. What are Weatherson's reasons for resisting these extensions?⁹

(PAT) is a conditional, so it is false only if there can be instances where (a) and (b₁) are true while (c₁) is false. Weatherson (2008, 569) suggests one such instance: 'the inference from *I am justified in believing that p* to *p* is not praiseworthy if the premise is false' (his emphasis). Perhaps it is fair to take him here as suggesting, more exactly, that S's belief that *p* is never praiseworthy if arrived at via a competent level-lowering ampliative inference where *q* is false. But this is very hard to maintain.

The extreme version of this suggestion would be that one can never be praiseworthy in believing something false. This would mean denying that only facts about the agent matter for the kind of epistemic evaluation that is focused on praise. But there are at least three reasons for resisting this denial and one further reason for resisting this suggestion in general.

First, the denial seems intuitively false: John's ideal reflection does seem to ground his praiseworthiness in believing that the sun is shining. As I've mentioned in section 4, the intuitive defense of (BAT) in section 2 seems equally good as a defense of (PAT), and the agent-centeredness of deontic evaluations is central to that defense. Second, the denial seems at odds with our evaluations of *moral* praiseworthiness. Suppose a soldier jumps on top of a

⁹Weatherson only discusses epistemic praise, so I'll focus exclusively on his reasons against (PAT). My claims in its defense, however, apply *mutatis mutandis* as a defense of (RAT).

live grenade intending to sacrifice his life for the safety of the group. Suppose that his action causes the opposing side to send a mortar bomb next, this time killing five-times more people than they would've otherwise. Whatever we say about the rightness or wrongness of this soldier's action—perhaps we think the consequences matter for that assessment—it seems intuitive that he is nonetheless praiseworthy for his sacrifice. Moral praiseworthiness, that is, seems agent-centered. Third, both of these intuitive judgments are buttressed by widely accepted accounts of moral and epistemic praiseworthiness. On these accounts, both praise and blame are evaluations of someone's *reason-responsiveness*: S is praiseworthy or blameworthy for ϕ -ing, that is, when S ϕ 's as a result of responding to what S perceives as sufficient reasons to ϕ . This is because praise and blame are assessments of the quality of someone's will, and because the quality of someone's will is revealed by how one responds to what one perceives as one's reasons.¹⁰ Taken together, these three considerations give us good reason for thinking that only facts about the agent matter for the kind of epistemic evaluation that is focused on praise.

But these reasons, at any rate, can be put aside. This extreme suggestion would make it impossible for the antecedent of (PAT) to be true in the kinds of cases that are supposed to show it false. Challenging (PAT), that is, requires accepting that one can in fact be praiseworthy in believing something false.

¹⁰For versions focused on *moral* praise and blame, see, e.g., Fischer & Ravizza (1998, 62-91) and Arpaly (2006, 16-7; 19); for versions focused on *epistemic* praise and blame, see, e.g., Ryan (2003, 70-74) and Hieronymi (2008, 359-363). I have argued elsewhere that our lack of voluntary control over our beliefs makes talk of epistemic praise and blame simply inappropriate (cf. Oliveira 2015, 393). My argument here is independent of that argument.

The moderate version of this suggestion would be that one can never transmit the praiseworthiness of a false belief. Here one avoids the main problem facing the extreme version just considered. But given the kind of transmission that is at stake in (PAT), this would mean denying that a false belief which is appropriately connected to grounds that make it praiseworthy can ever serve to appropriately connect another belief to grounds that would make *it* praiseworthy. Intuitively, once again, this seems false, specially when we notice that such failure cannot be explained by claiming that the relevant grounds are not good enough to make the further belief praiseworthy. The suggestion here would have to be that the quality of one's reflection would have been enough to make one praiseworthy in believing that p , had one not instead inferred p from some further belief also made praiseworthy by the quality of that reflection. But I've already suggested that this is implausible (section 2). Once we reflect on the kind of transmission that is at stake in these cases, that is, then we see that even the moderate version of Weatherson's suggestion should be rejected.

6.7 Conclusion

Since the two-stage defense of (BAT) mounted on sections 2 and 3 gives us good reason to think that (PAT) and (RAT) are true, versions of the Transmission Argument work equally well against versions of deontological internalism formulated in terms of epistemic praise and epistemic responsibility. As I have just argued, attempts to resist these damaging extensions are implausible. The moral, in fact, is quite general: if one attempts to draw a close connection be-

tween justification and an evaluation of how well the believer has carried herself given her situation—has she been blameless? has she been responsible? has she been praiseworthy?—then very plausible ampliative transmission principles and very plausible transmission arguments will lead one into a contradiction. Justification, that is, should be distinguished from the deontological appraisal of epistemic agency.

CHAPTER 7

NON-AGENTIAL PERMISSIBILITY IN EPISTEMOLOGY

7.1 Introduction

There is a sense in which you are justified in believing a certain proposition p by simply having reasons that support it. Call this *propositional justification*. There is another sense in which you are justified in believing that p only if you base your belief that p on those very reasons you have that support it. Call this *doxastic justification*. So you can be propositionally justified in believing that the defendant is guilty after hearing all of the evidence, for example, while at the same time failing to be doxastically justified: perhaps you simply have not formed that belief in the first place; perhaps you have formed it on the basis of reasons that do not support it after all. Call all of this *the received wisdom*.

Some have recently disagreed. Specifically, some have claimed that the requirements for doxastic justification are not as strong as the received wisdom seems to suggest. I here argue that at least one instance of such disagreement is based on faulty grounds. Close examination of these grounds and faults, however, illuminates the nature of permissibility in epistemology.

Consider Paul Silva's argument against the received wisdom.¹ He begins with a plausible guiding assumption:

A: Doxastic justification and moral justification are, essentially, the notions of doxastic permissibility and moral permissibility.

To judge that some act or doxastic attitude is justified, according to Silva (2014, 4), "is to judge that that act or attitude is permissible." I here accept this guiding assumption (cf. Goldman 2009 and Wedgwood 2012). But talk of justification is in general ambiguous between what we can call a *bare* and a *rich* conception:

Bare Moral Justification (MJ_b): S is morally justified iff S does the right thing.

Rich Moral Justification (MJ_r): S is morally justified iff S does the right thing for the right moral reasons.

Bare Doxastic Justification (DJ_b): S is doxastically justified iff S believes the right thing.

Rich Doxastic Justification (DJ_r): S is doxastically justified iff S believes the right thing for the right doxastic reasons.

¹Silva's argument has a positive and a negative stage. The positive stage consists in an argument against the received wisdom, while the negative stage consists in criticisms of alternative motivations for it. I here focus exclusively on the positive stage of his argument, and I take my criticisms as further motivation for the received wisdom.

Each of these conceptions associates the notion of justification with a certain positive evaluative feature. Given (A), however, deciding between DJ_b and DJ_r requires determining which of the relevant features maps onto the specific notion of doxastic permissibility. Similarly, given (A), deciding between MJ_b and MJ_r requires determining which of the relevant features maps onto the specific notion of moral permissibility. The remaining features, those not associated with permissibility of either kind, should be conceptualized some other way.

Given (A) and a grasp of the alternative conceptions, we can state Silva's argument against the received wisdom in this way:

Silva Against Wisdom (SAW):

1. MJ_b captures the notion of moral permissibility.
2. If MJ_b captures the notion of moral permissibility, then DJ_b captures the notion of doxastic permissibility.
3. So DJ_b captures the notion of doxastic permissibility.

Two defensive strategies come to mind almost immediately. One can reject premise 2 by denying (A), claiming that justification, in one or either case, is not essentially a notion of permissibility. Perhaps it is instead essentially the notion of requirement, or responsibility, or fittingness, or goodness, or something else. Alternatively, one can reject premise 1 by denying that MJ_b captures the notion of moral permissibility, claiming instead that doing the right thing for the wrong moral reasons is always impermissible. SAW is

defective if either of these defensive strategies works, though I here leave their details and plausibility to the side.

In what follows, I argue that premise 2 of SAW must be rejected for a different reason. First, I argue for a fundamental difference between moral and doxastic permissibility. Next, I argue that the corresponding notions of moral and doxastic justification reflect this difference in a way favorable to the received wisdom. Lastly, I argue that one of Silva’s key maneuvers in making plausible the rejection of the received wisdom falls apart.

7.2 Non-Agential Permissibility

Support for premise 2 comes from what I will call the *strong parity principle* (cf. Silva 2014, 6):

Strong Parity Principle: What is true of the structure of moral permissibility is true of the structure of doxastic permissibility.

If there is strong parity between the two notions of permissibility, then we have a good reason to accept that ‘ MJ_b iff DJ_b ’ and ‘ MJ_r iff DJ_r ’: if it is true that the moral notion has no basing requirement, then strong parity implies that the doxastic notion has no basing requirement either; conversely, if the doxastic notion has a basing requirement, then strong parity implies that the moral has a basing requirement as well. The strong parity principle is thus silent on what is characteristic of permissibility (on what *makes* an act or belief permissible), simply allowing us to infer the structure of one notion when having insight into the structure of the other—hence the need for premise 1.

But strong parity should be rejected. There is in fact a fundamental difference between the notions of moral and doxastic permissibility that removes our warrant for taking it that whatever goes for the moral notion also goes for the doxastic notion, and vice versa. The difference is this: one is a notion of *agential* permissibility, while the other is a notion of *non-agential* permissibility. We are agents with respect to our actions; that means that our actions are under our voluntary control. To say that certain actions are permitted, then, is to say that certain deployments of our agency do not violate the norms guiding our choice of available options. But we are not similarly agents with respect to our *beliefs*. That is, our beliefs are not similarly under our voluntary control. To say that certain beliefs are permitted, then, cannot be to say that certain deployments of our agency do not violate the norms guiding our choice of available options: we have no relevant agency in the doxastic case, and there are no corresponding available choices to be made.²

This fundamental difference between moral and doxastic permissibility—between agential and non-agential permissibility—is in fact a reflection of the deeper distinction between the *prescriptive* and the *evaluative* senses of

²We must be careful to distinguish between having control over *the kind of believer one is* and having control over *a particular token belief*. The orthodoxy since Alston (1989) is that we sometimes have the former but never the latter, and that the former by itself is not enough for agency with respect to a particular token belief. Recently, however, some have suggested that (in some sense or another) we sometimes have the latter kind of control (cf. Weatherson (2008) and Peels (2014)). I do not have space to address these challenges here. My view, at any rate, is that the empirical facts about the extent of our reflective agency are bleaker than even what is supposed by the Alstonian orthodoxy (cf. Kornblith (2012, 73-107)).

the English ‘ought’.³ In its prescriptive sense, the claim that ‘S ought to ϕ ’ expresses a requirement-relation between an agent and a course of action. This is the sense of ‘ought’ common from moral and prudential normativity. In its evaluative sense, however, the claim that ‘S ought to ϕ ’ expresses simply that, according to a relevant standard, the state of affairs of ‘S ϕ -ing’ is ideal. This is the sense of ‘ought’ familiar from claims such as ‘the world ought to be just’, where no agent and no action is involved. The agential notion of permissibility, then, is the notion of consistency with some relevant prescriptive ‘ought’: there is a requirement-relation between an agent and a course of action, and ϕ -ing does not violate that requirement-relation. A bit differently, the non-agential notion of permissibility is the notion of consistency with some relevant evaluative ‘ought’: there is a state of affairs that is ideal, according to a relevant standard, and ‘S ϕ -ing’ does not prevent that state of affairs from coming about.

The notion of moral permissibility is thus fundamentally different from the notion of doxastic permissibility. While the former is an agential notion, the latter is a non-agential notion; while the former is a claim about consistency with a certain requirement-relation, the latter is a claim about consistency with a relevant ideal. This is good reason to reject the strong parity principle: what is true of an agential notion about requirements may well differ from what is true about a non-agential notion about ideals.

³See Schroeder (2011) for detailed discussion and defense of this distinction, and see Chrisman (2008) for the view where ‘ought to believe’ always deploys the evaluative sense of the English ‘ought’.

7.3 Praise and Blame

Most will agree that ‘doing the right thing for the right reason’ is more ideal than simply ‘doing the right thing’. Both are positive evaluative features, but the former has an added good-making element. This suggests that the notion of doxastic permissibility, an evaluative notion about consistency with what is *ideal*, is best captured by DJ_r . At the same time, it seems plausible that what is required of us as agents often falls short of what is ideal. This suggests that the notion of moral permissibility, a prescriptive notion about consistency with a *requirement-relation*, is best captured by MJ_b . So we not only have reasons for rejecting the strong parity principle, we also have reasons for believing that something specific that is true of moral permissibility—no basing requirement—is not true of doxastic permissibility. This is just what the received wisdom would have us say.

In fact, the received wisdom gains even further support when we recognize that one of Silva’s (2014, 8-9) key maneuvers in making its rejection plausible falls apart. Since Silva accepts both MJ_b and DJ_b , and since he does not wish to deny that ‘doing the right thing for the right reason’ is a positive evaluative feature of some sort, he needs to tell a plausible story about which concept that happens to be. In the moral case, some such story seems readily available: cases of right actions based on bad moral reasons are naturally taken as cases of *blameworthy yet permissible* actions. Suppose you believe that pressing a certain button will cause incredible undeserved pain to thousands of people, when in fact it will only cause incredible deserved pleasure to them; suppose you press the button for that nasty reason; then it seems intuitive

that you are blameworthy for so doing, while doing it was a permissible action nonetheless (cf. Haji 1997). So MJ_r, we can say, fails to distinguish between the concepts of permissibility and praiseworthiness, running them together as if two necessary features of the former. Isn't there a similarly plausible doxastic story to be told? Just as in the moral case, perhaps cases of believing the right thing for the wrong doxastic reasons can be naturally taken as *blameworthy yet permissible* believing. Perhaps DJ_r also fails to distinguish between the concepts of permissibility and praiseworthiness, running them together as if two necessary features of the former.

Whatever we say about the plausibility of the moral story, clarity on the doxastic notion of permissibility reveals that the doxastic story is not at all forthcoming. This is because talk of praise and blame is appropriate only in the context of agents and requirement-relations. If S is an agent who is required to ϕ , then S can be praiseworthy and blameworthy for ϕ -ing or not. There was a certain requirement, after all, a requirement which S's action could flaunt or satisfy, and we rightly praise or blame S according to her voluntary choice. But if S is not an agent with respect to ϕ -ing, and if, consequently, 'S ϕ -ing' is simply an ideal state of affairs instead of a requirement, then there is no sense in which S can be praised or blamed: there are no requirements in this context, after all, no alternative actions that can flaunt or satisfy them, and no agent to voluntarily choose between them.⁴ A clock ought to strike every hour on the

⁴Fischer (2006, 24-25), for example, uses precisely this principle—that S is blameworthy for ϕ -ing only if S prescriptively ought to ϕ —to argue from Frankfurt-style cases to a rejection of the ought-implies-can principle.

hour, but there is no sense in claiming that a clock is blameworthy for being fast.

7.4 Conclusion

While the notions of praise and blame may provide us a plausible story regarding the concept associated with the positive evaluative feature of ‘doing the right thing for the right moral reasons’, there is no correlative plausible story regarding ‘believing the right thing for the right doxastic reasons’. Silva’s key maneuver in support of DJ_b thus falls apart. But we have already undermined his positive reasons for accepting DJ_b over DJ_r anyway. And if we take DJ_r over DJ_b instead, as reflection on the non-agential nature of doxastic permissibility suggests we do, then we simply have no need for an alternative story mirroring the role of praise and blame in the moral case. This is also what the received wisdom would have us say.⁵

⁵The fact that we are not agents with respect to our beliefs has direct and indirect implications for several alternative notions of doxastic justification. Notions of justification as praiseworthiness or blamelessness are most directly affected. But notions of doxastic justification as responsible belief may be indirectly affected as well. I cannot pursue those implications here.

BIBLIOGRAPHY

- Ahlstrom-Vij, Kristoffer (2013). "In Defense of Veritistic Value Monism," *Pacific Philosophical Quarterly* 94(1): 19-40.
- Alston, William (1988). "The Deontological Conception of Epistemic Justification," *Philosophical Perspectives* 2: 257-299.
- Alston, William (2005). *Beyond Justification: Dimensions of Epistemic Evaluation*. Cornell University Press.
- Altschul, Jon (2014). "Epistemic Deontology and Role-Oughts," *Logos and Episteme* 3: 245-263.
- Anomaly, Jonny (2008). "Internal Reasons and the Ought-Implies-Can Principle," *Philosophical Forum* 39(4): 469-483
- Arpaly, Nomy (2006). *Merit, Meaning, and Human Bondage: An Essay on Free Will*. Princeton University Press
- Blum, Alex (2000). "The Kantian Versus Frankfurt," *Analysis* 60: 287-288.
- BonJour, Lawrence (2010). *Epistemology: Classic Problems and Contemporary Responses*. Rowman & Littlefield Publishers, Inc.
- Booth, Anthony Robert (2012). "All Things Considered Duties to Believe," *Synthese* 187(2): 509-517.
- Brogaard, Berit (2014). "Wide-Scope Requirements and the Ethics of Belief," in *The Ethics of Belief*, Matheson, J. & Vitz, R. (eds). Oxford University Press: 130-146
- Broome, John (1999). "Normative Requirements." *Ratio* 12(4): 398-419.
- Broome, John (2007). "Wide or Narrow Scope?" *Mind* 116(462): 359-370.
- Chang, Ruth (2004). "All Things Considered". *Philosophical Perspectives* 18(1):1-22.
- Chrisman, Matthew (2008). "Ought to Believe." *Journal of Philosophy* 105(7): 346-370.

- Chrisman, Matthew (2012a). "Ought and Control". *Australasian Journal of Philosophy*, 90(3): 433-451.
- Chrisman, Matthew (2012b). "The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions". *Journal of Philosophy* 109(10): 588-612.
- Clifford, W.K. (1877/1999). "The Ethics of Belief," in *The Ethics of Belief and Other Essays*. Prometheus Books: 70-97
- Cohen, Stewart (1984). "Justification and Truth," *Philosophical Studies* 46(3): 279-95.
- Conee, Earl & Feldman, Richard (2011). "Replies," in *Evidentialism and its Discontents*, Dougherty, T. (ed.). Oxford University Press: 283-324
- Cowie, Christopher (2014). "In Defence of Instrumentalism About Epistemic Normativity," *Synthese* 191(16): 4003-4017.
- Davies, Martin (2004). "Epistemic Entitlement, Warrant Transmission and Easy Knowledge," *Aristotelian Society Supplementary Volume* 78(1): 213-245.
- Dougherty, Trent (2011). "In Defense of Propositionalism About Evidence," in *Evidentialism and Its Discontents*, Dougherty, T. (ed.). Oxford University Press: 226-35.
- Dowell, J. L. (2013). "Flexible Contextualism About Deontic Modals: A Puzzle About Information-Sensitivity". *Inquiry* 56(2-3): 149-178.
- Dworkin, Ronald (2011). *Justice for Hedgehogs*. Harvard University Press.
- Enoch, David (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press
- Eklund, Matti (2013). "Evaluative Language and Evaluative Reality," in *Thick Concepts*, Kirchin, S. (ed.). Oxford University Press: 161-182.
- Feldman, Fred (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*. D. Reidel Publishing Company.
- Feldman, Richard (1988). "Epistemic Obligations," *Philosophical Perspectives* 2: 235-256.
- Feldman, Richard (2000). "The Ethics of Belief," *Philosophy and Phenomenological Research* 60(3): 667-695.

- Feldman, Richard (2001). "Voluntary Belief and Epistemic Evaluation," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Steup M. (ed). Oxford University Press: 77-92.
- Feldman, Richard (2006). "Clifford's Principle, James' options," *Social Epistemology* 20(1): 19-33.
- Feldman, Richard (2008). "Modest Deontologism in Epistemology". *Synthese* 161: 339-355.
- Finlay, Stephen (2009). "Oughts and Ends". *Philosophical Studies* 143(3): 315-40.
- Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Finlay, Stephen & Snedegar, Justin (2014). "One 'Ought' Too Many". *Philosophy and Phenomenological Research* 86(1): 102-124.
- von Fintel, Kai & Gillies, Anthony S. (2011). "'Might' Made Right". In *Epistemic Modality*, Egan, A. & Weatherson, B. (eds.). Oxford University Press: 108-130.
- Fischer, John Martin & Ravizza, Mark (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Fischer, John Martin (2006). "Introduction: A Framework for Moral Responsibility," *My Way*: 1-37. Oxford University Press.
- Frankena, William. K. (1964). *Ethics*. Prentice Hall, Inc.
- Frankfurt, Harry (1969). "Moral Responsibility and the Principle of Alternative Possibilities," *Journal of Philosophy* 66: 829-839.
- Garcia, Ernesto V. (2004). "Value Realism and the Internalism/Externalism Debate," *Philosophical Studies* 117(1-2): 231-258.
- Geach, Peter. (1956). "Good and Evil," *Analysis* 17: 33-42.
- Geach, Peter (1982). "Whatever Happened to Deontic Logic?" *Philosophia* 11: 1-12.
- Gert, Joshua (2004). *Brute Rationality: Normativity and Human Action*. Cambridge University Press.
- Goldman, Alvin (2009). "Internalism, Externalism, and the Architecture of Justification," *Journal of Philosophy* 106(6): 309-338.

- Graham, Peter A. (2011a). "Ought and Ability," *Philosophical Review* 120(3): 337-382.
- Graham, Peter A. (2011b). "Fischer on Blameworthiness and 'Ought' Implies 'Can'," *Social Theory and Practice* 37(1): 63-80.
- Grimm, Stephen (2009). "Epistemic Normativity." In *Epistemic Value*, Pritchard, D., Haddock, A., and Millar, A. (eds). Oxford University Press: 243-264.
- Haack, Susan (2001). "The Ethics of Belief Reconsidered," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Steup M. (ed.). Oxford University Press: 12-54.
- Haji, Ishtiyaque (1997). "An Epistemic Dimension of Praiseworthiness," *Philosophy and Phenomenological Research* 57(3): 523-544.
- Harman, Gilbert (1986). *Change in View: Principles of Reasoning*. MIT Press.
- Henning, Tim (2014). "Normative Reasons Contextualism". *Philosophy and Phenomenological Research*. 88(3): 593-624.
- Hieronymi, Pamela (2008). "Responsibility for Believing," *Synthese* 161(3): 357-373.
- Horty, John F. (2001). *Agency and Deontic Logic*. Oxford University Press.
- Jackson, Frank (1991). "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection". *Ethics* 101: 461-82.
- Jones, Ward E. (2009). "The Goods and the Motivation of Believing," in *Epistemic Value*, Pritchard, D., Haddock, A., and Millar, A. (eds). Oxford University Press: 139-62.
- Kant, Immanuel (1788/1956). *Critique of Practical Reason*, trans. Lewis White Beck. University of Chicago Press.
- Kearns, Stephen & Star, Daniel (2009). "Reasons as Evidence" in *Oxford Studies in Metaethics* 4: 215-42.
- Kolodny, Niko (2005). "Why Be Rational?" *Mind* 114(455): 509-63.
- Kornblith, Hilary (1983). "Justified Belief and Epistemically Responsible Action," *Philosophical Review* 92(1): 33-48.

- Kornblith, Hilary (2001). "Epistemic Obligation and the Possibility of Internalism," in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, Fairweather, A. and Zagzebski L. (eds.). Oxford University Press: 231-248.
- Kornblith, Hilary (2002). *Knowledge And Its Place in Nature*. Oxford University Press.
- Kornblith, Hilary (2014). *On Reflection*. Oxford University Press.
- Korsgaard, Christine (1996). *The Sources of Normativity*. Cambridge University Press.
- Kratzer, Angelica (1977). "What Must and Can Must and Can Mean". *Linguistics and Philosophy* 1: 337-55.
- Kratzer, Angelica (1981). "The Notional Category of Modality" In *Semantics From Different Points of View*, Eikmeyer, H. and Rieser, H. (eds). Springer: 117-47.
- Lewis, David (1989). "Dispositional Theories of Value", *Proceedings of the Aristotelian Society* 63: 113-137.
- Littlejohn, Clayton (2012a). *Justification and the Truth-Connection*. Cambridge University Press
- Littlejohn, Clayton (2012b). "Does 'Ought' Still Imply 'Can'?" *Philosophia* 40(4): 821-828.
- Littlejohn, Clayton (2013). "Are Epistemic Reasons Ever Reasons to Promote?" *Logos and Episteme* 4(3): 353-360.
- Locke, John (1689/1979). *The Clarendon Edition of the Works of John Locke: An Essay Concerning Human Understanding*. Oxford: Clarendon Press.
- Locke, John (1706/1966). *John Locke's of the Conduct of the Understanding*. New York: Teachers College Press.
- Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. Penguin.
- McDowell, John (1995). "Might There Be External Reasons?" In *World, Mind and Ethics: Essays on the Ethical Philosophy of Bernard Williams*, Altham J.E.J. and Harrison, R. (eds.). Cambridge University Press.
- Mill, John Stuart. (1863/1962). *Utilitarianism*. Cleveland, World Pub. Co.
- Moore, G. E. (1903/2004). *Principia Ethica*. Dover Publications.

- Moran, Richard A. (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton University Press.
- Nisbett, Richard & Timothy Wilson (1977). "Telling More Than We Can Know: Verbal Reports on Mental Processes," *Psychological Review* 84: 231-59.
- Nolfi, Kate (2014). "Why is Epistemic Evaluation Prescriptive?" *Inquiry* 57(1): 97-121.
- Olsen, Kristian. (2014). "Ross and the Particularism/Generalism Divide". *Canadian Journal of Philosophy* 44(1): 56-75.
- Papineau, David (2014). "There Are No Norms of Belief," in *The Aim of Belief*, Chan, T. (ed.). Oxford University Press: 64-80.
- Parfit, Derek. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Parfit, Derek (1997). "Reasons and Motivation". *Aristotelian Society Supplementary Volume* 71(1):99-130.
- Parfit, Derek (2011). *On What Matters*. Oxford University Press.
- Pascal, B (1670/1966). *Pensées* (translated by A.J. Krailshiemer). Penguin Press.
- Peels, Rik (2014). "Against Doxastic Compatibilism," *Philosophy and Phenomenological Research* 89(3): 679-702.
- Peels, Rik (2014). "Believing at Will is Possible," *Australasian Journal of Philosophy* 93(3): 1-18.
- Petersen, Steve (2013). "Utilitarian Epistemology," *Synthese* 190(6): 1173-1184.
- Plantinga, Alvin and Wolterstorff, Nicholas (1983). *Faith and Rationality: Reason and Belief in God*. University of Notre Dame Press.
- Plantinga, Alvin (1993). *Warrant: The Current Debate*. Oxford University Press.
- Pryor, James (2001). "Highlights of Recent Epistemology," *British Journal for the Philosophy of Science* 52(1): 95-124.
- Railton, Peter (1986) "Moral Realism," *Philosophical Review* 95: 163-207.
- Railton, Peter (1999). "Normative Force and Normative Freedom: Hume and Kant, but Not Hume Versus Kant," *Ratio* 12(4): 320-353.

- Raz, Joseph (2001). *Value, Respect, and Attachment*. Cambridge University Press.
- Raz, Joseph (2002). *Practical Reasoning and Norms*. Oxford University Press.
- Raz, Joseph (2011). *From Normativity to Responsibility*. Oxford University Press.
- Ryan, Sharon (2003). "Doxastic Compatibilism and the Ethics of Belief." *Philosophical Studies* 114(1-2): 47-79.
- Ross, Jake (2010). "The Irreducibility of Personal Obligation". *Journal of Philosophical Logic* 39(3): 307-23.
- Ross, W. D. (1930/2002). *The Right and The Good*. Clarendon Press.
- Ross, W. D. (1939/2000). *Foundations of Ethics*. Oxford University Press.
- Saka, Paul (2000). "Ought Does Not Imply Can." *American Philosophical Quarterly* 37(2): 93-105.
- Skorupski, John (2010). *The Domain of Reasons*. Oxford University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Harvard University Press.
- Scanlon T. M. (2014). *Being Realistic About Reasons*. Oxford University Press.
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press.
- Schroeder, Mark (2009). "Means-End Coherence, Stringency, and Subjective Reasons," *Philosophical Studies* 143: 223-248.
- Schroeder, Mark (2011). "Ought, Agents, and Actions." *Philosophical Review* 120: 1-41.
- Schroeder, Mark (2012). "The Ubiquity of State-Given Reasons," *Ethics* 122(3): 457-488.
- Schroeder, Mark (2015). "Knowledge is Belief for Sufficient (Objective and Subjective) Reason," in *Oxford Studies in Metaphysics: Volume 5*, Szabo Gendler, T. and Hawthorne J. (eds.). Oxford University Press: 226-252.
- Silva, Paul (2015). Does Doxastic Justification Have a Basing Requirement?, *Australasian Journal of Philosophy* 93(2): 371-387
- Shah, Nishi (2006). "A New Argument for Evidentialism," *Philosophical Quarterly* 56(225): 481-498.

- Shah, Nishi & Velleman, J. David (2005). "Doxastic Deliberation," *Philosophical Review* 114(4): 497-534.
- Smith, Michael (1994) *The Moral Problem*. Oxford University Press.
- Smith, Michael (2013). "On the Nature and Significance of the Distinction Between Thick and Thin Ethical Concepts," in *Thick Concepts*, Kirchin, S. (ed.) Oxford University Press: 97-120.
- Smithies, Declan (forthcoming). "Why Justification Matters," In *Epistemic Evaluation: Point and Purpose in Epistemology* Henderson, D. & Greco, J. (eds.). Oxford University Press.
- Sosa, Ernest (2009). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Volume I. Oxford University Press.
- Stalnaker, Robert (2014). *Context*. Oxford University Press.
- Steup, Mathias (2000) "Doxastic Voluntarism and Epistemic Deontology," *Acta Analytica* 15(24): 25-56.
- Steup, Mathias (2008) "Doxastic Freedom," *Synthese* 161(3): 375-392.
- Steup, Matthias (2012). "Belief Control and Intentionality," *Synthese* 188(2): 145-163.
- Stich, Stephen P. (1990) *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. MIT Press.
- Stratton-Lake, Phillip. (2002). "Pleasure and Reflection in Ross's Intuitionism". In *Ethical Intuitionism: Re-Evaluations*, Stratton-Lake, P. (ed.). Oxford University Press: 113-36.
- Street, Sharon (2008). "Constructivism about Reasons," *Oxford Studies in Metaethics* 3: 208-245.
- Street, Sharon (2009). "Evolution and the Normativity of Epistemic Reasons.," *Canadian Journal of Philosophy* 39(sup1): 213-248
- Street, Sharon (2012). "Coming to Terms With Contingency: Humean Constructivism About Practical Reason," in *Constructivism in Practical Philosophy*, Lenman, J. and Shemmer, Y. (eds). Oxford University Press.
- Talbot, Brian (2014). "Truth Promoting Non-Evidential Reasons for Belief," *Philosophical Studies* 168(3): 599-618.
- Thomson, Judith J. (2008). *Normativity*. Open Court.

- Tucker, Chris (2010). "When Transmission Fails," *Philosophical Review* 119(4): 497-529.
- Turri, John (2009). "The Ontology of Epistemic Reasons," *Nous* 43: 490-512.
- Turri, John (2011). "Believing For a Reason," *Erkenntnis* 74(3): 383-397.
- Van Woudenberg, René (2009). "Responsible Belief and Our Social Institutions," *Philosophy* 84(327): 47-73.
- Way, Jonathan (2011). "The Symmetry of Rational Requirements," *Philosophical Studies* 155(2): 227-239.
- Weatherson, Brian (2008). "Deontology and Descartes's Demon," *Journal of Philosophy* 105(9): 540-569.
- Wedgwood, Ralph (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- Wedgwood, Ralph (2012). "Justified Inference," *Synthese* 189(2): 1-23.
- Wedgwood, Ralph (2013). "Rational 'Ought' Implies 'Can'," *Philosophical Issues* 23(1): 70-92.
- Williams, Bernard (1981a). "Ought and Moral Obligation". In *Moral Luck*, 114-23. Oxford University Press.
- Williams, Bernard (1981b). "Internal and External Reasons". In *Moral Luck*: 101-113. Cambridge University Press.
- Williams, Bernard (1995). "Internal Reasons and the Obscurity of Blame". In *Making Sense of Humanity*. Cambridge University Press: 35-45.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press.
- Wolterstorff, Nicholas (1996). *John Locke and the Ethics of Belief*. Cambridge University Press.
- Wright, Crispin (2003). "Some Reflections on the Acquisition of Warrant by Inference," In *New Essays on Semantic Externalism and Self-Knowledge*, Nuccetelli, S. (ed.). MIT Press: 57-78.
- Wright, Crispin (2004). "Warrant for Nothing (and Foundations for Free)?" *Aristotelian Society Supplementary Volume* 78(1): 167-212.
- Zimmerman, Michael. J. (1980). "On the Intrinsic Value of States of Pleasure," *Philosophy and Phenomenological Research* 41(1-2): 26-45.

Zimmerman, Michael (2006). "Is Moral Obligation Objective or Subjective?" *Utilitas* 18: 329-361.

Zimmerman, Michael (2008). *Living With Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.