Philos Stud (2017) 174:485–506
DOI 10.1007/s11098-016-0693-4

Deontological evidentialism, wide-scope, and privileged values

Luis R. G. Oliveira[1]

Published online: 31 May 2016

Abstract Deontological evidentialism (DE) is the claim that we ought to form and maintain our beliefs in accordance with our evidence. In this paper, I criticize two arguments in its defense. I begin by discussing Berit Broogard's use of the distinction between narrow-scope and wide-scope requirements against W.K. Clifford's moral defense of (DE). I then use this very distinction against a defense of (DE) inspired by Stephen Grimm's more recent claims about the moral source of epistemic normativity. I use this distinction once again to argue that Hilary Kornblith's criticism of Richard Feldman's defense of (DE) is incomplete. Finally, I argue that Feldman's defense is insensitive to the relation between normative requirements and privileged values: values that have normative authority over us.

Keywords Epistemology · Deontology · Evidentialism · Wide-scope · Normative authority

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)

✉ Luis R. G. Oliveira
luis_oliveira@uml.edu

[1] Department of Philosophy, University of Massachusetts, Lowell, USA

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, 44–46), 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]

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.[4]

[1] The 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.

[2] Some 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.

[3] See Plantinga and Wolterstorff (1983) for discussion of this point.

[4] Much of the literature in defense of *evidentialism* states it as an account of *epistemic justification*. In such cases, it is often unclear which kind of normative claim evidentialism is intended to be and, more to

Here is how I proceed. In the first section, I discuss 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 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 Feldman's (2000, 2001, 2008) defense of (DE) and I clarify the unstated commitments that it involves. In the fourth section, I discuss Kornblith's (2001) reply to Feldman's defense. I argue that Kornblith's use of the distinction between narrow-scope and wide-scope 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 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. In the sixth and final section, I say a bit more about privileged values and the nature of normativity.[5]

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

Footnote 4 continued
my present point, whether and how it is related to (DE). I suspect, at any rate, that my criticism of (DE) in this paper will be relevant to a good number of evidentialists about epistemic justification, but I will not examine or argue for this further claim here. See, for example, McCain (2014) and the essays collected in Dougherty (2011).

[5] Cowie (2014, 4003) argues that there is a presumptive case in favor of what he calls *intrumentalism* 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 *intrinsicalism* about epistemic normativity: the claim that "there is reason to believe in accordance with one's evidence in virtue of a brutely 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 intrinsicalism—but do not appeal in any way to our goals—unlike instrumentalism. This is another reason—besides their historical influence and plausibility—why they deserve the separate and careful treatment they receive here.

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) while 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.[6]

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 shopper 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, 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 in fact a dog, then the desire will not by itself produce an act of kicking the dog. Brogaard thus 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*.

[6] I read Clifford differently from Brogaard (2014, 130–135). She takes him as committed to the premise that "false beliefs *could have* morally harmful consequences" (my emphasis).

So, 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:

(**narrow$_1$**) 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.

(**wide$_1$**) 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 (narrow$_1$) is true, then the Clifford's consequentialist moral argument goes through. But if (narrow$_1$) is false and (wide$_1$) 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.[7]

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 believes 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 (wide$_1$) instead of (narrow$_1$). 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 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

[7] Talk of wide and narrow *scope* is here talk about the place of the deontic operator in the underlying logical structure of premise (2). Let '□' 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 (narrow$_1$) has the form $\forall x\,(Fx \rightarrow \Box Ex)$. 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 (wide$_1$) has instead the form $\forall x\,(Fx \rightarrow \Box(Bx \rightarrow Rx))$. For a broader discussion of wide and narrow scope requirements in connection to *rationality*, see Kolodny (2005), Broome (2007), Schroeder (2009), and Way (2011).

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 (narrow_1)? I will consider and reject one such attempt.

2 The Grimm-inspired moral argument

Grimm (2009, 253–254) 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, 258–259), 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.[8]

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

[8] Grimm (2009, 259 fn. 32) is in fact aware that his claims are, as he puts it, "Cliffordian."

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, 258–259) 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 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. But 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.[9]

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 claim 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 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 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

[9] There 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* 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.

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:

(**narrow$_2$**) 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.

(**wide$_2$**) 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 (narrow$_2$) is true, then the Grimm-inspired moral argument goes through. But if (narrow$_2$) is false and (wide$_2$) is true instead, then the argument is unsound. It does not follow from (wide$_2$) 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 (wide$_2$) 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.[10]

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, (narrow$_2$) 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.

[10] One 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, however, 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 it does arise due to the possibility that someone who cares about p might come to depend on us in the future. (See how both (narrow$_2$) and (wide$_2$) are formulated in terms of *possible* dependences.) On this matter, the Grimmean view and the proposed alternative are in agreement. Where they differ is in the nature—more exactly, the *structure*—of the requirement that they identify: the Grimmean view sees it as a requirement that can be satisfied in just one way, while the proposed alternative provides for a choice. I am grateful to an anonymous reviewer for bringing this worry to my attention.

3 The epistemic point-of-view argument

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.[11]

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, moreover, play the role of the *believer*. We are all engaged in the activity of forming, maintaining, and revising our beliefs. Accordingly, just as 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:

[11] My discussion here puts together as a unified picture the claims we find in Feldman (2000, 676), (2001, 87–89), (2008, 349–352).

(correctness) For each role, there is a correct way to perform that role.[12]

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 that role.

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 capture this with the *ought* principle:

[12] A more modest proposal could say that, though there are often more than just one correct way to perform a role, there are certain essential features that are shared by every correct way. I am grateful to Chris Meacham and Seth Cable for discussion of this point.

> **(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.[13] 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, premise (1) seems true. Premise (2), however, requires some careful defense. Feldman's defense of it turns on the claim

[13] In this way, (Feldman 1988, 240–243) rejects his earlier claim 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.

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.[14] 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.[15]

4 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. Yet Kornblith's explanation of why this is the case is mistaken. 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

[14] See Feldman (2008, 346–347).

[15] See 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.

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$_{\text{Narrow}}$**) 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$_{\text{Wide}}$**) 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$_{\text{Wide}}$) 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.

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$_{\textbf{Wide}}$**) 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:

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.[16]

5 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,

[16] In 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.

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.[17] 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 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. We can then 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.[18]

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 in the next section.) 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

[17] This is not to suggest these various senses of the English 'ought' betray a difference in semantic structure and/or syntactic behavior. See, for example, Finlay and Snedegar (2014).

[18] I 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. I am grateful to Miles Tucker for bringing this to my attention.

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* ($\text{ought}_{\text{NER}}$) 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 $\text{ought}_{\text{NER}}$ 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 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 robing 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 $\text{ought}_{\text{BRR}}$ 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.[19]

[19] Let 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 of epistemic role oughts that distinguishes it from more ordinary role oughts, such as the teacher role ought, and one might think that this distinguishing feature is enough to justify the claim that epistemic role oughts express normative requirements after all. But this is not the case. On this point, Kornblith (2001, 237–238) 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."

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.[20]

6 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. 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.[21]

[20] Feldman is not alone in grounding epistemic normativity on what is valuable from the epistemic point-of-view. For two influential further examples, see Alston (1993, 531) and Sosa (2009, 70). My argument in this section applies to these and similar views as well: authoritative requirements cannot be grounded *merely* on the claim that epistemic values issue from a certain point-of-view.

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

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.[22]

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.[23] Another option is to connect epistemic requirements to non-epistemic privileged values. We have seen 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.[24]

7 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

[22] For the constructivist approach to normativity, see Railton (1986), Lewis (1989), Smith (1994), Korsgaard (1996), and Street (2008).

[23] See Ahlstrom-Vij (2013) for discussion.

[24] For the instrumentalist approach to epistemic normativity, see Stich (1990), Kornblith (2002), and Cowie (2014). I do not assume that an account of epistemic normativity (instrumentalist or otherwise) is *ipso facto* an account of the notion of *epistemic justification*. For my views on justification, see Oliveira (2015) and Oliveira (forthcoming). For my views on epistemic normativity—and its relation to justification—see Oliveira (unpublished). I am grateful to an anonymous reviewer for the encouragement to add the brief positive remarks that compose this section.

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 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.

Acknowledgments I'm grateful to Hilary Kornblith, Christopher Meacham, Seth Cable, Timothy Perrine, Kristian Olsen, Scott Hill, Dennis Kavlakoglu, Ed Ferrier, Miles Tucker, Bob Gruber, and an anonymous referee for discussion and comments on previous drafts.

References

Ahlstrom-Vij, K. (2013). In defense of veritistic value monism. *Pacific Philosophical Quarterly*, *94*(1), 19–40.

Alston, W. (1993). Epistemic desiderata. *Synthese*, *53*(3), 527–551.

Altschul, J. (2014). Epistemic deontologism and role-oughts. *Logos and Episteme*, *3*, 245–263.

Brogaard, B. (2014). Wide-scope requirements and the ethics of belief. In J. Matheson & R. Vitz (Eds.), *The ethics of belief* (pp. 130–146). Oxford: Oxford University Press.

Broome, J. (1999). Normative requirements. *Ratio*, *12*(4), 398–419.

Broome, J. (2007). Wide or narrow scope? *Mind*, *116*(462), 359–370.

Clifford, W. K. (1877/1999). The ethics of belief. In *The ethics of belief and other essays* (pp. 70–97). Amherst: Prometheus Books.

Cowie, C. (2014). In defence of instrumentalism about epistemic normativity. *Synthese*, *191*(16), 4003–4017.

Dougherty, T. (2011). *Evidentialism and its discontents*. Oxford: Oxford University Press.

Feldman, R. (1988). Epistemic obligations. *Philosophical Perspectives*, *2*, 235–256.

Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, *60*(3), 667–695.

Feldman, R. (2001). Voluntary belief and epistemic evaluation. In M. Steup (Ed.), *Knowledge, truth, and duty: Essays on epistemic justification, responsibility, and virtue* (pp. 77–92). Oxford: Oxford University Press.

Feldman, R. (2006). Clifford's principle, James' options. *Social Epistemology*, *20*(1), 19–33.

Feldman, R. (2008). Modest deontologism in epistemology. *Synthese*, *161*, 339–355.

Finlay, S., & Snedegar, J. (2014). One ought too many. *Philosophy and Phenomenological Research*, *86*(1), 102–124.

Grimm, S. (2009). Epistemic normativity. In *Epistemic value* (pp. 243–264). Oxford: Oxford University Press.

Haack, S. (2001). The ethics of belief reconsidered. In S. Matthias (Ed.), *Knowledge, truth, and duty: Essays on epistemic justification, responsibility, and virtue* (pp. 12–54). Oxford: Oxford University Press.

Kolodny, N. (2005). Why be rational? *Mind*, *114*(455), 509–563.

Kornblith, H. (2001). Epistemic obligation and the possibility of internalism. In A. Fairweather & L. Zagzebski (Eds.), *Virtue epistemology: Essays on epistemic virtue and responsibility* (pp. 231–248). Oxford: Oxford University Press.

Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Oxford University Press.

Korsgaard, C. (1996). *The sources of normativity*. Cambridge: Cambridge University Press.

Lewis, D. (1989). Dispositional theories of value. *Proceedings of the Aristotelian Society*, *63*, 113–137.

Locke, J. (1689/1979). *The Clarendon edition of the works of John Locke: an essay concerning human understanding*. Oxford: Clarendon Press.

Locke, J. (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.

McCain, K. (2014). *Evidentialism and epistemic justification*. London: Routledge Press.

Moore, G. E. (1903/1993) *Principia Ethica*, T. Baldwin (ed.). Cambridge: Cambridge University Press.

Oliveira, L. R. G. (2015). Non-Agential permissibility in epistemology. *Australasian Journal of Philosophy, 93*(2), 389–394.
Oliveira, L. R. G. Ampliative transmission and deontological internalism. *Pacific Philosophical Quarterly*. (forthcoming).
Oliveira, L. R. G. Constructivist instrumentalism about epistemic normativity. (unpublished).
Plantinga, A., & Wolterstorff, N. (1983). *Faith and rationality: Reason and belief in god*. Notre Dame: University of Notre Dame Press.
Railton, P. (1986). Moral Realism. *Philosophical Review*, *95*, 163–207.
Schroeder, M. (2009). Means-end coherence, stringency, and subjective reasons. *Philosophical Studies*, *143*, 223–248.
Skorupski, J. (2010). *The domain of reasons*. Oxford: oxford University Press.
Smith, M. (1994). *The moral problem*. Oxford: Oxford University Press.
Sosa, E. (2009). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. I). Oxford: Oxford University Press.
Stich, S. P. (1990). *The fragmentation of reason: Preface to a pragmatic theory of cognitive evaluation*. Cambridge: MIT Press.
Street, S. (2008). Constructivism about reasons. *Oxford Studies in Metaethics*, *3*, 208–245.
Street, S. (2012). Coming to terms with contingency: Humean constructivism about practical reason. In J. Lenman & Y. Shemmer (Eds.), *Constructivism in practical philosophy*. Oxford: Oxford University Press.
Way, J. (2011). The symmetry of rational requirements. *Philosophical Studies*, *155*(2), 227–239.
Wolterstorff, N. (1996). *John Locke and the ethics of belief*. Cambridge: Cambridge University Press.