

DEONTOLOGICAL EPISTEMOLOGY

by

SEAN MESLAR

(Under the Direction of Sarah Wright)

ABSTRACT

The goal of my dissertation project is to defend two logically independent but related conclusions regarding epistemic justification. The first is that epistemic justification is correctly described only by a deontological theory (abbreviated hereafter as ‘DCEJ,’ for, ‘deontological conception of epistemic justification), meaning justification is analyzed in terms of duty fulfillment. The second is that we have at least one epistemic duty: an unconditional obligation to believe the truth. The first thesis provokes something of a ground-breaking in contemporary epistemology. While the DCEJ is controversial in itself, my review of the related literature suggests no one has correctly identified what, precisely, the DCEJ *is*. Once one resolves this latter issue, one is faced with a set of hitherto unconsidered questions. The second thesis finds support in the strong intuition that justification and truth are importantly connected, but stands opposed to almost every theorist on epistemic justification in that it denies the possibility of justified, but false belief. Finally, given the truth of these theses, epistemic norms are incommensurable with other types of norms, meaning the importance of obeying our obligation to believe the truth cannot be weighed against the importance of any moral obligations we might face.

INDEX WORDS: Epistemology, Meta-epistemology, Justification, Deontology, Doxastic
voluntarism

DEONTOLOGICAL EPISTEMOLOGY

by

SEAN MESLAR

B.A., Christopher Newport University, 2012

A Dissertation Submitted to the Graduate Faculty of The University of Georgia in Partial
Fulfillment of the Requirements for the Degree

DOCTOR OF PHILOSOPHY

ATHENS, GEORGIA

2017

© 2017

Sean Meslar

All Rights Reserved

DEONTOLOGICAL EPISTEMOLOGY

by

SEAN MESLAR

Major Professor:	Sarah Wright
Committee:	Melissa Seymour Fahmy
	Beth Preston

Electronic Version Approved:

Suzanne Barbour
Dean of the Graduate School
The University of Georgia
May 2017

ACKNOWLEDGEMENTS

Charles B. Cross served as a committee member prior to his retirement in 2016.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iv
CHAPTER	
1 INTRODUCTION	1
2 DOXASTIC VOLUNTARISM	11
2.1 Types of Control	13
2.2 Steup and Epistemic Compatibilism.....	16
2.3 Feldman and the Nature of ‘Epistemic Oughts’	19
2.4 Voluntariness and the Meaning of ‘Can’	23
3 INTERNALISM AND DEONTOLOGY	27
3.1 Goldman and Plantinga.....	29
3.2 Steup and Ginet.....	33
4 DEONTOLOGY AND BLAME	42
4.1 Weak Blame Arguments	43
4.2 Strong Blame Arguments.....	48
5 METAEPISTEMOLOGY.....	55
5.1 Deontology.....	57
5.2 Consequentialism.....	59
5.3 Virtue Epistemology	65
5.4 Proper Functionalism and Divine Command Theory	73

5.5 Applied Epistemology and Truth.....	77
6 THE TRUTH NORM OF BELIEF.....	80
6.1 Objects of Epistemic Evaluation.....	81
6.2 True Belief and Normative Epistemology	84
6.3 Formulating the Truth Norm.....	87
6.4 Objections to the Truth Norm	93
6.5 Derivative Norms.....	98
7 THE SPECIES OF NORMATIVITY	110
7.1 The Epistemology of Action.....	111
7.2 The Ethics of Belief	115
7.3 General Normativity	121
7.4 Objections and Replies	127
REFERENCES	131

CHAPTER 1

INTRODUCTION

Dating back to Plato's *Theaetetus*, the concept of epistemic justification was thought to derive its importance from the central topic of epistemology: knowledge. By giving a satisfactory definition of justification, it was assumed that the concept of knowledge would be consequently illuminated. That tradition was disrupted most visibly by Gettier's famous counter-examples to the justified true belief analysis of knowledge. Despite justification being incapable of turning true belief into knowledge, epistemologists have retained a keen interest in the subject. In this way, the present work is closer to its contemporaries; the central topic I shall consider herein is epistemic justification, leaving the question of knowledge largely untouched. Specifically, I wish to defend a deontological conception of epistemic justification (hereafter DCEJ). The theory has some prominent historical supporters, most famously Locke¹ and Descartes² but has fallen on hard times of late. Given the state of the literature, my first task will be to respond to the most pressing criticisms of the DCEJ; doing this will also accomplish the goal of detailing exactly what a deontological view of epistemic justification must say and, crucially, that about which it can defensibly remain silent. Once I have given my case for why the DCEJ is not beset with fatal difficulties from the start, I turn to the more positive task of defending a particular deontological theory of epistemic justification.

¹ See his *Essay Concerning Human Understanding*, Book IV, Chapter xvii, §24

² Most visibly in the fourth meditation

The second chapter seeks to respond to the most prominent objection faced by the DCEJ, namely, that it entails doxastic voluntarism, the thesis that our beliefs are under our voluntary control. Deontic formulations of epistemic duties will necessarily make use of ‘ought’ statements. ‘Ought’ implies ‘can’, so the DCEJ is committed to the view that we can control our beliefs, which is just the thesis of doxastic voluntarism. Because doxastic voluntarism is false, the argument goes, so too is any position that entails it. The canonical articulation of this objection comes from Alston (1988). I also consider the strategies of Steup (1988) and Feldman (1988) as responses to Alston. Finally, I develop my own defense on behalf of the DCEJ, which concludes that Alston’s argument is inapplicable to the DCEJ because ‘can,’ when predicated of an agent’s ability to believe, does not carry with it any implication of voluntariness; rather, it is best understood as ranging over the ways in which one can believe. Thus, the sort of ‘ought’ statements employed by the DCEJ will not carry the problematic entailment.

Chapter 3 continues the project of responding to objections to the DCEJ. It is often suggested that the DCEJ entails internalism about epistemic justification, the thesis that epistemic justification supervenes on the internal states of the evaluated agent. Since internalism is a highly contentious position, it would be in the interest of the DCEJ not to be committed to it. The chapter consists of a review of arguments in support of this entailment claim and my arguments against them. I challenge the pro-entailment arguments with the observation that the only essential commitment of the DCEJ is merely that epistemic justification is a matter of duty fulfillment. Since this is so, the DCEJ only entails internalism if all epistemic duties are concerned only with internal conditions. But, internalism just is the claim that all epistemic duties are concerned only with internal conditions. Therefore, the DCEJ, taken with the truth of

internalism, entails the truth of internalism, but so does everything. No extant argument suffices to show that the DCEJ is independently indicative of the truth of internalism.

The fourth chapter concludes the first part of the dissertation, and so continues with the same general project of the previous two chapters. Another concept commonly associated with the DCEJ is that of blame, such that deontic evaluations should be understood in terms thereof, e.g., if a subject believes without regard to her evidence, then she is blameworthy. I show that, in most cases, the association between the DCEJ and blame appeals to the arguments I dismiss in the previous chapter. Nottelmann (2013) argues in favor of the entailment without appealing to a connection between the DCEJ and internalism, and so merits special consideration. His central claim is that it is only through an appeal to the notion of blame that the DCEJ can avoid explaining justification in non-deontic terms. The DCEJ must seek to avoid this, for if it allows that there are true statements about epistemic justification that do not use any deontic terms, then it is not a truly deontological account. I argue for the following alternative view: the DCEJ is not committed to the truth of *any* statements about epistemic justification. Rather, what makes a theory deontological is how it explains the relation in which we must stand with regard to the true epistemic norms, whatever they may be. Establishing this point suffices to show that there are few theorists who have correctly identified *what* the DCEJ is, despite its acceptance by historical luminaries like Descartes and Locke.

Having dismissed some of the more popular objections to the DCEJ, I turn to motivating my preferred position in chapter five. As elaborated below, I take as a premise the aptness of an analogy between epistemology and ethics, which indicates that there are distinct areas of inquiry within epistemology corresponding to the popular division from ethics between meta, normative and applied ethics. In these terms, the DCEJ is best classified as a position within meta-

epistemology. I consider the prospects of epistemic consequentialism and virtue epistemology, which correspond to the two most popular alternatives to deontology in the metaethical sphere. I argue that the former faces difficulties in accommodating the unconditional badness of contradictory beliefs, and that the latter, while a commendable area of inquiry, is too broad in scope to serve as the meta-theoretical grounding for a theory of specifically epistemic justification.

The task of chapter six is to give a positive, albeit indirect, argument in favor of the DCEJ. I argue that an unconditional requirement of true belief is a fundamental epistemic norm. Since the DCEJ is best able to provide the meta-epistemological endorsement of this norm, the DCEJ should be accepted. An unconditional truth norm runs against many intuitions, and has to face many puzzling cases if it is to be endorsed wholeheartedly. The majority of this chapter is devoted to considering probable objections. Ultimately, however, the truth norm should be endorsed because beliefs are legitimate objects of epistemic evaluation, and beliefs can only be the immediate object of normative evaluation in terms of their correctness. Since the aim at truth is constitutive of belief, correctness of a belief and the truth of its content are identical.

If the arguments from all the previous chapters succeed, then the project approaches a complete account of epistemic justification. However, one significant issue remains if we are fully to understand the nature of our epistemic obligations. Most will accept that we have not only epistemic obligations, but moral, and perhaps prudential obligations as well. How do these obligations relate to each other? The first task of chapter seven is to show that our epistemic obligations as described in the previous chapters can come into conflict with any plausible account of our moral obligations. In order to resolve this conflict it seems we need to appeal to some more general theory of normativity. Any such higher-order theory is faced with a dilemma,

however. Either the general theory will accommodate the specific theories by deeming valuable the attainment of the specific values targeted by the lower-order theories, or it will seek to accommodate all of them through the attainment of a more general value. If the former route is taken, then we are at a loss to explain why one specific value is preferred to another. If the latter route is taken, then the lower-order theories will conflict with the higher-order theory, since the specific norms have their own non-normative grounding independent of concerns about a general value. As such, it may be impossible to satisfy all the norms to which we are subject, and in such cases, there is no principled way of deciding which norms we ought to respect. I consider the Aristotelian notion of *eudaimonia* as a candidate for a higher-order value, but conclude that it succumbs to the same problems described in outline above.

With the project thus outlined, I want to make a few disclaimers, then a few assumptions. My concern in this dissertation is limited to *epistemic* justification. Therefore, in cases lacking further specification, the reader can infer that when I mention justification, I mean epistemic justification in particular. The term carries with it what might be, given its current usage, misleading connotations. Epistemic justification is typically assumed merely to be that which makes beliefs better, epistemically, than unjustified beliefs.³ There are presumably many different types of justification, though legal, moral and epistemic justification are perhaps the most prominent of the bunch. There need not be any correlation between different species of justification. For instance, I might be justified according to rules of poker to bluff about the contents of my hand, even if I lack the moral justification to lie for a morally insignificant end. Only in the seventh chapter do I consider any other species of justification at any length, thus the reader ought to take extra care in interpreting the term there.

³ Admittedly, this is something of an unhelpful definition. As will become apparent later, however, it is difficult to say more about the meaning of the term without demonstrating a bias for a particular view concerning the nature of epistemic justification.

Additionally, the project treats epistemic justification as a normative concept, such that one should have justified beliefs, and one should avoid having unjustified beliefs. It is worth noting that accepting the sense of justification discussed in the previous paragraph is separable from the normative element I indicate here. If knowledge is justified, true belief, plus some additional element, it does not follow in the absence of further argument that we ought be justified in our beliefs. Nonetheless, it is the more involved, normative understanding of the concept that shall be the focus of this project.

In making my project concerned with epistemic justification, I put it at odds with ‘naturalized epistemology’ of the sort advocated by Quine.⁴ By way of a (very) brief summary, naturalized epistemology is the view that the study of beliefs and knowledge is best accomplished through empirical research of the sort produced by psychology and cognitive science. Perhaps unsurprisingly, epistemologists have not been convinced by Quine’s arguments.⁵ It is my intention neither to recapitulate, nor to advance any arguments against naturalized epistemology— that is a significant enough project in its own right. Rather, I recognize that the truth of the theory would be fatal to normative epistemology in general, perhaps particularly so to the DCEJ. Nonetheless, my position faces criticism from much closer sources. Arguing for a specific view within normative epistemology, I assume, rather than argue for, the viability of the traditional epistemological project.

My next assumption is the principle that ‘ought implies can’, or ‘Kant’s law’, after its namesake’s frequent appeal to the principle.⁶ This rule is straightforward and intuitive: if one

⁴ This is a project that can be seen consistently in his writings, but most prominently in “Epistemology Naturalized”(1969).

⁵ I say unsurprisingly since, if he were right, philosophers claiming to be epistemologists would be out of a job.

⁶ His notable appeals to this principle include arguments for belief in God, the immortality of the soul, and the possibility of property.

ought to do something, then one can do that thing, and, contra-positively, if one cannot do something, it is not the case that one ought to do that thing. Reflecting on a few examples shows the appeal of the principle. At least one of the reasons ex-post facto laws are commonly thought to be unjust is that they create obligations that one cannot fulfill due to temporal constraints— an ‘ought’ without a ‘can’, as it were. Of course, even the simplest rules can hide subtleties. Exactly how one should understand both ‘ought’ and ‘can’ merits some additional consideration. I shall address the meaning of ‘ought’ later in the introduction and that of ‘can’ in section 2.4, making sure to reflect on the stipulated definitions’ ability to cohere with the basic intuition behind this premise.

The next is a two-part observation about the so-called ‘direction of fit’ of beliefs and actions. Beliefs are said to succeed when they fit the mind to the world, that is, when they represent to the subject the truth through their propositional content.⁷ Actions are said to succeed when they fit the world to the mind, that is, when they make it the case that the intention motivating the action becomes true of the world. Naturally, my concern is more directly with the observation about beliefs. It follows that true beliefs are good, in some sense of the word, even before establishing any normative epistemological principles. The goodness that true beliefs have merely in virtue of being true is intrinsic. The goodness does not derive from the content of the proposition that is believed, the agent that believes it, nor from the causal history of the belief’s formation. One can push the observation further: the only potential intrinsic good available to a belief is in truthfully representing the content that it carries; nothing could make a true belief better *qua* belief. While this is a principle upon which I shall rely heavily, I list it here as a

⁷ This rule becomes a little more complicated in thinking about *reflective* beliefs, inasmuch as the direction of fit in at least some such beliefs seems to be one of mind to mind. We can say of these cases that the ‘world’ referred to is less a matter of externality as opposed to internality and more a matter of objectivity and subjectivity; reflective beliefs take the subject as their object, making the subject itself a part of the rest of the cognizable world, so to speak.

premise, and therefore assume its uncontroversiality, because the goodness it discusses is fully non-normative. Chapter 6 makes the much more controversial move of attempting to derive a norm from this initial observation. We can note that this additional inference requires argumentative support since, for instance, few, if any normative ethical theories suggest that an action is morally good (merely) in virtue of its successfully realizing its intention.

Of less immediate interest, though still relevant to the project in general, is the observation concerning the success of actions. If their criteria for success are distinct, then it follows that beliefs are not a specific type of action, but rather closer to what philosophers have heretofore called passions—ways in which subjects are affected.⁸ Whether this is problematic for those who, with Clifford, want to develop an ethics of belief, will be explored later, in chapter 7.

A different point in support of the same conclusion is the phenomenon called the transparency of belief. For a fixed subject S and proposition P, there is no semantic difference between the statements ‘S believes that P’ and ‘S believes that P is true’. As such, false beliefs can be said to be faulty in a more fundamental way than beliefs that are imprudent, at odds with one’s religious or ethical obligations, epistemically lucky, etc. From this phenomenon, it is often observed that the aim at truth is constitutive of belief.⁹ It is my estimation that chapter 6 argues to a similar conclusion through different means.

I also assume that doxastic involuntarism is true. The beliefs that we adopt are not the direct result of a voluntary action on our part, nor is it the case that we can suspend belief, withhold belief or change our attitudes toward a given proposition at will. Once more, the explanation seems to be found in the centrality of truth to belief. We believe in accordance with

⁸ One thinks most readily of Descartes’ *On the Passions of the Soul*.

⁹ Shah and Velleman (2005) seems to be the landmark work in support of this claim.

how the truth appears to us, even if we want to believe otherwise. This premise is readily supported by a simple exercise. Consider a proposition that you believe to be false, or one to which you have no particular doxastic attitude, and try to will yourself to believe the proposition to be true. Presumably, nothing will happen. The most obvious explanation is that your willing has no bearing on the truth of the proposition, and hence does not influence that at which the belief aims.¹⁰ The popular opinion is that this assumption is incompatible with my thesis, hence I make my attempt at demonstrating their consistency the topic of the second chapter. I say that I accept it as a premise, however, because I appeal to the truth of the phenomenon later in defending deontology against its metaepistemological competitors.

Last, I want to assume the legitimacy of a widely observed parallel between ethics and epistemology. This is a significantly vaguer notion than the premises previously endorsed and will therefore serve less as a specific datum and more as a source of direction. Most prominently, I shall develop a distinction that mirrors the three-tiered system of ethical inquiry: meta-ethics, normative ethics and applied ethics. More specific claims can garner motivation from the analogy; e.g., if rule consequentialism faces a difficulty in distinguishing itself from act consequentialism in ethics, we can expect that the epistemological analogs will face a similar challenge. However, the inferences made from applications of this parallel are best understood as being merely *prima facie*. I have already noted an importance difference between the criteria of success for beliefs and for actions. In general, my policy will be to let stand these sorts of analogous inferences only in the absence of substantial independent arguments.

¹⁰ Here reflective beliefs may serve as a legitimate counter-example. It seems one really can will oneself to believe that one is willing oneself to believe. Even in this case, however, we can plausibly say that the connection between will and belief is accidental; the preponderance of non-voluntary belief formations suggest that it is the truth, rather than the willing, that gives rise to the belief.

With these preliminary remarks made, it will also be of use to consider, what, exactly, I mean by a deontological conception of epistemic justification. A theory is deontological iff its evaluations are fundamentally a matter of the evaluated's standing with regard to some proposition that employs deontic operators.¹¹ There are three deontic operators: 'permission', 'forbiddance' and 'obligation'. Their relations are as follows: some proposition P is permissible iff it is not forbidden; P is obligatory iff $\sim$ P is not permissible. By one's duty, I mean one's respecting the rules of the theory. Though my focus is on the epistemological version, respecting the ethics/epistemology analogy suggests that this definition should suffice for ethical deontology as well.

From this definition, we can also define how the DCEJ analyzes the 'ought' operator when used in epistemological contexts. Any true 'ought' claim will observe some truth regarding an individual's epistemic duty. For instance, if it is true that one ought to believe in proportion to one's evidence, then it is also true that we are obliged to do so. 'Ought' claims seem applicable only to the stronger two operators; it cannot rightly be said that, if one is merely permitted to believe that it will rain tomorrow, then one ought or ought not to do so.

Chapters two through four consist, in principle, of an application of this definition. It is not until the fifth chapter, having considered the main objections it faces, that we arrive at a final account of the DCEJ's essential commitments. Finally, the sixth and seventh chapters adopt the task of filling in the requisite details for a theory of epistemic justification.

¹¹ I specify 'fundamentally' since, for instance, a consequentialist could formulate a norm that uses deontological terms, but this would only follow due some other value-theoretic principles.

CHAPTER 2

DOXASTIC VOLUNTARISM

I have already claimed that a deontological account of epistemic justification will define justification in terms of adherence to epistemic duty. Having not yet arrived at a refined understanding of what that duty is, two broad characterizations of justification remain live options. First, we can think of justification in positive terms, such that one is justified in one's belief if and only if it is one's epistemic duty to have that belief. Alternatively, a negative characterization will deem that one is justified in one's belief if and only if one has not failed to do one's epistemic duty in so believing. A number of advocates of a deontological view have ruled in favor of the latter.¹² Assuming paradigmatic instances of justified belief are preserved by a deontological view, the former generate epistemic duties that require us to be implausibly opinionated. If we consider an evidentialist formulation of an epistemic duty, something like 'believe all the propositions for which you have adequate evidence', we will be obliged to have an infinite number of beliefs affirming logic theorems of ever-increasing complexity.¹³ The negative characterization requires only that, when we do form beliefs, we do so in a way that conforms to epistemic duty; as such, the demands are much less radical.

On this view, beliefs that involve a transgression of one's epistemic duty are forbidden. From how I have defined epistemic duties above, this suggests that, for forbidden beliefs, one

¹² See (Steup, 2000), (Ginet 1975).

¹³ Admittedly, the second half of the dissertation is devoted to defending such a view, but that I feel the need to defend the view suggests I recognize the attractiveness of the reasoning given above.

ought (not to believe them). From Kant's law, it follows that one can (not believe them).¹⁴ But this last implication is false, since we cannot believe and disbelieve at will. As such, the deontological view asserts a contradiction: it both is and is not the case that we ought not to believe in accordance with our epistemic duty. Let us call this the 'argument from doxastic voluntarism.' Indeed, if my preliminary definition of deontology is correct, the argument from doxastic voluntarism poses a unique threat to the DCEJ. Since it is necessary and sufficient for a view's being deontological that it make its evaluations essentially a matter of duty fulfillment (and hence essentially about epistemic 'oughts') the applicability of Kant's law to deontological evaluations is inescapable.¹⁵ This argument is most famously championed by William Alston in his article "The Deontological Conception of Epistemic Justification," (1988) thus, I shall turn to his article to give a more complete account of the argument.

First, however, we might note an easy, if somewhat unsatisfying response to this argument. One might accept the premises that we cannot voluntarily control our beliefs and that if we cannot do so, then we cannot have any obligation to believe. One could consistently accept both of the premises along with the DCEJ, though the result would be something like justificational nihilism. I say nihilism, rather than skepticism, because not only would claims to having justification admit of doubt, but also they could be proven to be universally false. Unsurprisingly, I do not wish to examine this line of thinking at any length. Since my project is to defend a particular account of epistemic justification, it would be a rather futile exercise to defend a theory that delivered such a destructive result. Of course, this is not meant to prove

¹⁴ I have added parentheses to clarify an otherwise irremediable scope ambiguity

¹⁵ For instance, a consequentialist view could analyze justification merely in terms of maximizing epistemic value, so, while it would be best if an agent were to believe justifiedly, a consequentialist is not committed to an accompanying 'ought' statement.

justificational nihilism false, but only to support preferring more optimistic strategies, at least until it is shown that they are fundamentally flawed.

2.1 Types of Control

As noted in the introduction, I think that doxastic voluntarism is false, so, I shall ultimately agree with Alston that we lack the sort of control over our beliefs that he considers. Nonetheless, his observations concerning the extent of our control over our doxastic attitudes are useful in helping to determine what an adequate theory of justification must say, and why a deontological view is at particular risk of error with regard to this question.

Let us begin with the simplest sort of control one might have, what Alston calls “basic voluntary control.” (260) Examples of this sort of control are found in our ability to take a breath and to blink¹⁶ It seems obvious that we have no such control with regard to beliefs. No one would ever commit perjury, for instance, if he or she could simply change his or her beliefs when asked about them. As Alston notes (261), this inability extends to the three basic doxastic attitudes: believing, disbelieving and withholding; moving ourselves to believe contrarily to our standing beliefs is just as much beyond the reach of our will as is making ourselves adopt a belief toward a proposition the truth or falsity for which we have absolutely no evidence.

Presumably, Alston will face little opposition on this point. At least somewhat more controversial is his taking a similar position on the existence of what he calls “non-basic voluntary control.” (268) Actions over which we have non-basic voluntary control are those that require the mediate impact of basic voluntary actions for their accomplishment. Speaking is a fine example of such actions; it is one that we can perform at will, but only through the use of the more basic actions of making the sounds of individual syllables. There appear to be some

¹⁶ We have positive control over these actions even if we lack negative control— e.g. we can make ourselves blink at will, but we may not be able not to blink through sheer willpower.

instances of beliefs that we can cause ourselves to have by virtue of actions over which we have basic voluntary control. The clearest examples of such beliefs will be reflections of our performances of basic voluntary actions, e.g., the proposition that one is exhaling.

There are at least two difficulties with this line of thinking, however. In the first place, it seems that any theory of justification will address more beliefs than those over which we have direct control. As Alston notes, even if we take the broad category of 'looking for evidence' as a voluntary action, we do not always succeed in finding evidence, and thus cannot make ourselves adopt a belief on the matter. If there are any epistemically significant propositions for which it is not in our (immediate) power to find evidence, then the problem is not solved. (271) Secondly, even in the most favorable of cases, it is not clear that the sort of control we have pertains to any specific proposition. (271) What we really can control is not the belief that we adopt, but the proposition to which we adopt the doxastic attitude. However, evaluations regarding epistemic justification are about specific attitudes, not simply having an attitude toward a proposition. Thus, while we have a certain kind of non-basic voluntary control over our beliefs, it is not sufficient for the purposes of the DCEJ.¹⁷

The last sort of control Alston considers is "long range voluntary control." (474) The case for a degree of control sufficient for deontological evaluation is more favorable here as it seems we can control more of our beliefs in this way. We can make ourselves more likely to believe a proposition by repeatedly presenting evidence in favor of that claim and by consistently avoiding evidence against that claim. Obviously, even this sort of control can only range so far; perhaps we only have this sort of control over propositions that admit of reasonable doubt. This limitation

¹⁷ This point should also indicate, if it was not already clear, why I stick so closely to Alston as a source for an argument against doxastic voluntarism; while the above point may cause problems for some arguments against the thesis, since Alston's argument is concerned with the DCEJ, it is especially qualified only to concern itself with the relevant degree of voluntary control considering the more general goal of this project.

should not be seen as fatal, however, since the more interesting cases will fall into just this class.¹⁸ More discouraging for the prospects of doxastic voluntarism is that all cases of long-range voluntary control require a great deal of effort on behalf of the epistemic agent. In many cases, even attempting to exercise this sort of control over our beliefs will prove ineffective. Thus, although our attitude toward a given proposition is possibly under our control, it does not follow that it is actually under our control, and it is this latter condition that Kant's law appears to require. If we have the opportunity to exert ourselves and, by so doing, bring some condition under our control, this indirect control seems to suffice for Kant's law, but if we exert ourselves to the utmost and are still unable to bring the condition under our control, it seems there can be no corresponding obligation.¹⁹ If this is the case, then it seems we are left once more with insufficient ground for deontological evaluations.

Having made little progress in finding doxastic control, Alston makes what he feels to be a charitable shift in requirement for a deontological view, "one that identifies being justified in believing that p with not being intellectually to blame for believing that p..." (284) This final condition is subject to more or less the same difficulties as would beset any deontological view, if only less directly. One can come to hold unjustified doxastic attitudes as a result of no personal failings. Since these attitudes are involuntary, one cannot be blamed for the unjustified nature of one's beliefs, so even the most favorable case, that of merely influencing beliefs through habits, fails to provide the deontological view that which it needs.

¹⁸ I suppose my reasoning behind this claim is vaguely Cartesian; if there is no room for reasonable doubt in considering a proposition, then we will believe in accord with the evidence. Since we have already endorsed the negative understanding of justification, beliefs of this sort will not transgress any epistemic duties, so they are justified. 'Ought' may imply 'can', but 'must' implies 'can' as well. Therefore, these sorts of insurmountable beliefs do not threaten to trigger the problematic application of Kant's law.

¹⁹ It is somewhat difficult to think of an analogous case pertaining to actions. Perhaps some feat requiring exceptional fitness might be morally required of a person who is exceptionally fit. Less plausible, however, is the claim that anyone who, contrary to fact, could have brought himself or herself to such fitness levels faces the same obligation.

2.2 Steup and Epistemic Compatibilism

Given that we have already endorsed a general parallel between ethics and epistemology, the discussion to this point naturally suggests the epistemic analog of compatibilism. If compatibilism (about action) is just the thesis that it is logically possible that there can be free acts that are fully and antecedently causally determined, the epistemic analog will state that it is logically possible that doxastic attitudes are adopted freely, even if those attitudes are fully and antecedently determined. Traditionally, compatibilists have maintained that the necessary and sufficient condition for an action's being free is that it is 'internally' caused, i.e. caused by some psychological influence proper to the agent.²⁰ The task of determining the viability of an epistemic compatibilism will, therefore center on what it means for doxastic attitudes to be caused internally.

Steup (2008) offers one such account. He seems to settle on the following analysis of doxastic freedom:

S's attitude A toward p is free iff (i) S has attitude A toward p, and (ii) S's attitude A is weakly intentional; (iii) S's having taken attitude A toward p is the causal outcome of a reason-responsive mental process. (385)

While "S's Φ ing is weakly intentional iff S Φ s in a way that is... non-accidental and involves a pro-attitude [an unconscious, but retrospectively endorsed intention]." (385) On this analysis, one may come to have beliefs even in the absence of a conscious intention to do so, and yet still be free in so believing. Considering the conditions individually suggests that most of our doxastic attitudes are free. (i) follows trivially, and, in general, we believe in accordance with how things seem to us, which suffices for both (ii) and (iii). In a similar vein to compatibilism

²⁰ The canonical source for this view is Stace (1980), though Frankfurt (1969) is also seen an important contributor to the literature.

about action, it is urged that, even if this view of doxastic freedom is different from the traditional conception thereof, it secures everything that is theoretically desirable about doxastic voluntarism with none of the dubious metaphysical (or in this case, psychological) assumptions.

The inclusion of the 'pro-attitude' seems to secure the viability of Steup's analysis. Since free doxastic attitudes will involve a retrospective endorsement, there is a convincing sense in which one becomes responsible for the attitude in question, even if that attitude did not result directly from an exercise of the will. Presumably, if one is responsible for an attitude, then that attitude is a legitimate subject of deontological evaluation.

Nonetheless, I think this analysis only displaces the problem, rather than solving it. If one endorses an attitude retrospectively, one does so on the basis of a standing attitude. For instance, one might believe that one has hands on the basis that it so seems to him or her, and one might have a pro-attitude that seemings should be believed in the absence of defeating evidence. That pro-attitude is a substantive epistemic thesis, one that must be believed if it is to be used to make retrospective endorsements. However, if the pro-attitude was not freely adopted, then there is little reason to think that the beliefs endorsed thereby are free either. Thus, a vicious regress ensues; in order for any belief to be free, one must already have a standing, freely adopted belief to ground one's pro-attitude.²¹ While less controversial than a bold-faced assertion of doxastic voluntarism, grounding the DCEJ in a compatibilist view of doxastic freedom leaves it in a dialectically undesirable position. I admit that, in some respects, I have not given doxastic compatibilism a fair trial here. The motivating idea behind the theory is less that we can use it to explain how doxastic attitudes can be free in the sense required by the sort of argument Alston gives and more about how properly to define doxastic freedom. The epistemic compatibilist need

²¹ Traditional compatibilism about action faces a similar concern. Even if I do something because I want to do it, it is argued, if my so wanting is not free, then neither is that action that results therefrom (see Frankfurt, 1971).

not dread the regress argument I give if he or she is willing to say that the sort of condition described by Steup is all we should require from an analysis of doxastic freedom. However, I do think it fair to say that it cannot provide the sort of *sui generis* attitudes that Alston evidently thinks are necessary for epistemic responsibility. Because of the explicit function of the pro-attitude in the aforementioned analysis, beliefs adhering to this schema will be antecedently determined, and hence not ones over which we have voluntary control.

Fortunately, while the case for doxastic voluntarism might be quite bad, as Steup (1988) argues, it might not need to be very good to meet the requirements of the DCEJ. He appeals to two examples from an earlier paper by Alston (1985). In the first, a member of an isolated community believes in witchcraft on objectively poor epistemic grounds. Alston thinks that the belief will be justified according the DCEJ, since there is no basis for blaming the agent; there is nothing that he or she could have done such that, if he or she had done it, he or she would have believed differently (67-68). On the other hand, Steup rejects this verdict. Working from an evidentialist perspective, he argues that all epistemic agents are capable of realizing the truth of an evidentialist norm (something like 'believe to and only to the extent that your beliefs are supported by your evidence) and that all epistemic agents are able to know, upon reflection, what their evidence is. (1988:78) From this, it follows that the agent in question, even as a part of an isolated, epistemically vicious community, is to blame for his or her unjustified belief.

The two authors consider another example that focuses on an individual's intellectual limitations, but to largely the same opposed conclusions. Alston think that individuals cannot be blamed (and hence the DCEJ cannot deem them unjustified) for the intuitively unjustified beliefs that form as a result of intellectual limitations (68), while Steup argues that the bare criteria for epistemic agency demand an ability to be aware of what is epistemically required of one. If an

agent is unable to understand that he or she should not believe that for which he or she has no evidence, then the agent is “epistemically, beyond good and evil.” (1998:80) It would perhaps be hasty to conclude that Steup has thoroughly refuted Alston's examples, but at the very least he raises points that merit consideration. Less controversially, we can observe that Steup relies on some plainly internalist intuitions about the nature of epistemic justification: most notably, that one is always in a position to determine that on which (given his additional endorsement of evidentialism) the justificatory status of one's beliefs supervenes. Without these intuitions, it seems that Steup would be left with little means to make a stand against Alston's objection. Thus, we can see that, at least on Steup's interpretation, the DCEJ entails epistemic internalism. The legitimacy of this claim is the subject of the next chapter. For now, however, our focus remains on the argument that the DCEJ entails doxastic voluntarism. At present, I shall turn to considering an objection of a different style.

2.3 Feldman and the Nature of 'Epistemic Oughts'

If we think of Alston's argument as essentially an application of Kant's law to the DCEJ, then there are two basic strategies one could adopt as a response. In the first place, one could affirm that one can do as the DCEJ indicates; this is roughly the strategy of Steup. On the other hand, one might deny that Kant's law applies to all uses of 'ought.' Feldman (1988) develops this strategy.

Some species of obligation are not bound by Kant's law. As Feldman notes (241), a student enrolled in a course who is unable to fulfill the requirements thereof (say, because the student is seriously ill) is not absolved of these responsibilities. Indeed, in such a case, the sensible thing to do would be to withdraw the student from the class, precisely because he or she cannot fulfill the academic obligations imposed by the course. Similarly, financial obligations are

not discharged by an inability to meet the obligatory conditions. Even if one cannot afford a payment due on one's student loans, one is still financially obliged to do so. If epistemic obligations are more like financial and academic obligations than they are like moral obligations (at least with respect to the question of entailing an ability to be fulfilled), then Alston's argument is disarmed.

One response that Feldman considers is that financial and academic obligations are not stand-alone species of obligations at all, but rather merely conditional prudential obligations. (242-243) On this view, the sorts of claims made above are better understood as saying 'if you want to pass the course, you should...' or 'if you want not to go to jail, you should pay your student loans.' Epistemic obligations, it is urged, are unconditional, and therefore relevantly unlike these latter sorts of obligations. However, we might just as easily turn talk of epistemic justification into a conditional obligation as well— "to say that you epistemically ought to believe *p* is to say that given that you want to achieve epistemic excellence, you ought to believe *p*." (243)²² The reduction of other species of obligations to merely conditional obligations therefore fails to exclude epistemic obligations, and amounts to little as a result.

This, however, is merely a negative point in Feldman's favor. The question of whether epistemic obligations actually are more like ethical obligations than non-'can' entailing obligations remains unanswered. Indeed, I have already noted the aptness of a general analogy between ethics and epistemology, which gives a *prima facie* consideration against Feldman's point. In order to give a definitive answer, we need to have at least an example of an epistemic obligation. Feldman, like Steup, holds an evidentialist view on justification. (1988: 254) If justification is analyzed in terms of an evidentialist principle, then it is unclear what, if anything,

²² And, in a positive indication for Feldman's point here, this strategy does not seem to be applicable in the case of actions, since, at least according to the most famous deontological view in ethics, moral obligations are categorical.

of value an element of control or voluntariness adds to the justification of our beliefs. Suppose there is an epistemic agent that can believe voluntarily, but is otherwise identical to us in its doxastic nature.²³ By hypothesis, the being under consideration and a typical human are capable of the same beliefs, held to the same degree of confidence (even if we are more limited in how those beliefs might originate). Both also can possibly hold all the same evidence. Now think of a case where the two agents hold the same beliefs (to the same extent) on the basis of the same evidence; suppose also that the human believes involuntarily and that the other agent does so voluntarily. If an evidentialist analysis of justification is true, then the two beings have exactly the same justificatory status. We can maintain this result because the *cause* of belief is not evaluated in the evidentialist analysis, but only the fit between belief and evidence. Therefore, whether or not our beliefs are voluntary makes no difference to the justification of our beliefs. If this is so, then the claim that voluntary control is a necessary condition for a belief's being evaluated in deontological terms seems entirely unmotivated.

Two concerns might be raised about this argument. In the first place, one might wonder if the evidentialist aspect has overpowered the deontological element. After all, the case was presented (intentionally) without using 'ought' or any other deontological terms. This worry, as I shall explain in chapter four, stems largely from a category error. Let the evidentialist rule be expressed as follows: it is obligatory that one believes to and only to the extent that is supported by one's evidence. So, one ought to believe in accordance with just that rule. If one does believe in accordance with this rule, then one is justified in so believing, while one is not justified when one breaks the rule. Given how we have defined a deontological view, all the elements are in place, despite the affirmation of an evidentialist norm.

²³ There is some debate among those in favor of doxastic involuntarism as to whether voluntary belief is conceptually or merely psychologically impossible. This thought-experiment is obviously unlikely to impress those who subscribe to the former view. Fortunately for its applicability here, Alston is in the latter camp (1988: 263)

Second, one could simply argue that evidentialism (when combined with deontology) is wrong, precisely because it obligates us to do something that is not under our voluntary control. I shall not pursue the strength of this argument, since, mercifully, it is beside the point. We are no longer considering the question of whether deontology can consistently deem obligatory beliefs over which we have no control; we have already seen that it can. One can consistently adopt a deontological view that endorses an evidentialist rule as its basis— that was shown in the previous paragraph. If this evidentialist rule makes obligatory beliefs over which we have no control, then the DCEJ does not entail doxastic voluntarism. Accepting such a view, one would be tempted to say that Kant's law is false; we have 'oughts' but no 'cans.' Since I have already indicated that I disagree with this conclusion, I do not wish to endorse it here. Rather, the point to be observed is a logical one: the DCEJ does not entail doxastic voluntarism, since the former can be true while, in some cases, the latter might be false. Of course, if doxastic voluntarism is a necessary truth, or at least a necessary truth about human psychology, then pointing out this logical gap will accomplish little, and there will remain a problem for the DCEJ to address. Since I accept doxastic voluntarism, I will address the problem in the next section. This paragraph, however, suffices to observe that, if one does reject doxastic voluntarism, there need not be fears of internal inconsistency on the part of the DCEJ.

If it is this easy to construct counter-examples, one might wonder why entailment arguments of the sort Alston gives have gained so much traction. The problem, as is usually the case, lies in a disagreement over definitions. Recalling Alston's analysis of deontological evaluations of justification in general, we can note that the concept of blame plays a prominent role. It is just this prominence that I think is to blame for the confusion at hand. I shall address the connection between blame and the DCEJ at length in the fourth chapter.

2.4 Voluntariness and the Meaning of 'Can'

I have given what I think serves as a counter-example to Alston's argument in noting the possible combination of deontology and evidentialism. Nonetheless, the case presented does nothing to disarm the original intuition that gave rise to his argument, an intuition that I have accepted as a premise: 'ought implies can.' My proposal in this section is that the claim that one can Φ does not entail, and should not be equated with, the claim that one can Φ *voluntarily*. The first part of this claim is easy to support through common usage. Things obviously not in possession of a will can do things. My car can drive faster than 50 mph, but it cannot do anything voluntarily.

Making a case that this sort of 'can' statement does not entail voluntariness is easy enough, but it does not obviously solve the problem. It would be just as misguided to say of my car that it ought to do anything as it would be to say that it willed anything; indeed, presumably the former would be absurd precisely because the latter is.²⁴ So, the likely response on the part of Alston would be to say that deontic 'oughts' imply voluntary 'cans'; both words can be used appropriately in other contexts, but these two usages are conceptually connected.

This position seems to be supported by thinking of cases of obligatory actions. If (for whatever reason), I have a moral obligation to run a mile, the voluntary sense of 'can' follows. The obligation could not sensibly be attributed to me while I was skydiving, even if I could run a mile, in the sense that I am in good enough cardiovascular shape to do so. Perhaps there are clever counter-examples that I am overlooking, but I shall grant (in what strikes me as a move generous to my opposition) that every *moral* 'ought' implies a voluntary 'can.' Given the general analogy between ethics and epistemology, we therefore have a *prima facie* reason to think that

²⁴ There is a certain usage of 'ought' that might sensibly be applied to a car, but certainly not one that is even in the same family as the epistemic oughts that are my central concern.

the same will hold true of epistemic 'oughts.' If this is the case, then we shall have difficulty explaining why the argument from doxastic voluntarism does not work.

Recall the premise concerning the different directions of fit of belief and action. Actions succeed when they fit the world to the mind, while beliefs succeed when they fit the mind to the world. Specifically, actions involve an exercise of the will, but as we have noted, beliefs are better described as ways of being affected than as a subspecies of action; there is no epistemic will, and even if there was, using it could only get in the way of the success of our beliefs. By considering direction of fit, we can see why it is essential to 'oughts' concerning actions that they entail voluntariness. A good action, at least on the deontological view, will be one that involves fitting the world to the mind in the right way. One fits the world to one's mind by exercising one's will. A good belief, according to the DCEJ, will be one that fits the mind to the world in the right way. It still follows that we cannot be epistemically obliged to be affected in a way that we cannot, but the demand that it is within our voluntary control to be affected in a certain way is clearly mistaken; the two are mutually exclusive categories.²⁵

One disclaimer is needed as an accompaniment to this argument. I noted in introducing the varying directions of fit between belief and action that the sort of success that can be attributed to a belief and action that fit in the right direction is not obviously identical with epistemic or ethical justification. There are, on all accounts, actions that can succeed without being ethically justified. However, my claim above seems to conflate the two, such that epistemic justification is identical with, or at least requires, having the right fit. I believe this is true of epistemic justification, as I shall argue in chapter 6. Thus, the success of my argument here relies on my success there.

²⁵ We may have an obligation to take an action that results in our being affected in a certain way, but this is only possible through an intermediate element. We cannot take action to make ourselves passive.

If the sort of 'can' that is implied by epistemic 'oughts' does not imply voluntariness, then how ought we to interpret it? This question will merit additional consideration when we move to considering more substantive epistemic norms, but I should like to give some initial thoughts on the topic. Using the ethical requirement as a guide, my suggestion is 'can' in the sense implied by epistemic 'oughts' means that the mind can represent the proposition about the world targeted by the belief. An example will help to shed some of the terminological density. The proposition that I have hands is eligible to serve as epistemically obligatory, since my mind has command of the relevant concepts and categories. On the other hand, I cannot be obliged to believe a contradiction, because my mind is simply not constituted so as to be (possibly) affected in that way.²⁶ Analogously, if the world is such that it cannot (physically, logically, etc.) be shaped by the will to fit the mind, we determine that there are no ethical obligations that it be so shaped.

Conclusion to Chapter 2

We have reviewed and dismissed one of the prominent objections to the deontological conception of epistemic justification: the argument from doxastic voluntarism. The argument claims that the DCEJ entails that beliefs are voluntary, but since beliefs are not voluntary, the DCEJ is false. William Alston considered a variety of types of control that one might claim we have over our beliefs, and found them all to be insufficient for the demands of the DCEJ. Even the most favorable deontic analysis of justification, which is centrally concerned with belief-influencing actions, is still insufficient to preserve the blamelessness (or blameworthiness) that Alston feels is essential to deontic evaluations.

There are two basic strategies of objecting to the argument from doxastic voluntarism. In the first place, one might claim that our beliefs are voluntary, and accept the entailment claim, or

²⁶ We might say that this 'because' is over-determined, since a contradiction will never be true of the world, and hence never the object of a successful belief, I am simply looking for the least controversial instance of an impossible belief.

one might deny the entailment claim. The first strategy is endorsed by Matthias Steup, who argued that we wield a sufficient degree of indirect control over our beliefs to merit their being evaluated. The second strategy is endorsed by Richard Feldman, who argues that epistemic 'oughts', unlike their moral counter-parts, do not entail 'cans.' My own argument reflects elements of both strategies. We must be able to have the beliefs demanded of us by a deontological theory, but that ability does not need to be, nor should it be, one that can be exercised voluntarily.

One may notice a general trend in all the pro-deontological authors I have mentioned thus far: they all endorse internalism. Indeed, we saw that Steup and Feldman both appealed to an evidentialist norm in responding to the argument from doxastic voluntarism. Much of the literature on the DCEJ has associated it with internalism. The goal of the next chapter is to explore the relation between the two theories.

CHAPTER 3

INTERNALISM AND DEONTOLOGY

Perhaps the most prominent debate in contemporary epistemology is between internalists and externalists. In keeping with the central interest of this project, when I speak of both, I shall be referring to the respective stances on justification, rather than, say, knowledge. Giving a definition to which all parties of the debate would agree is difficult for both, perhaps more so for internalism than for externalism. Nevertheless, it will help to have a general idea of both terms, even if one could disagree over some of the finer details of what I propose. Let us take internalism to be the view that the justificatory status of some agent A's beliefs supervenes on states internal to A.²⁷ Most take this requirement to imply that everything relevant to a given belief's justification is something to which the believer has access. Feldman and Connee have described the core thesis of internalism as 'mentalism', where the internal states mentioned in the original definition are members of the believer's 'mental life.' (2001: 2) Evidentialists like Steup endorse internalism by analyzing justification in terms of fit to evidence, and in turn requiring that whatever is evidence is something that is in the believer's cognitive possession. (Steup, 2001: 237)²⁸ Though the distinction has certainly been developed long past his time, we can say with confidence that Descartes would have had an internalist view of justification, since clarity and distinctness are both mental terms.

²⁷ 'internal' in the sense in which it is used in philosophy of mind, not (at least not directly) physiological facts about the agent's body.

²⁸ It is perhaps relevant to note that one could be an evidentialist and not an internalist, depending on how one defines 'evidence.'

Externalism, on the other hand, is essentially the view that external factors can at least play a partial role in determining epistemic justification.²⁹ The earliest major externalist arguments were on the topic of knowledge, coming from Dretske (1971), Armstrong (1973) and Nozick (1981). The first significant externalist position on epistemic justification comes from Goldman (1979). Without going into any of the finer details of these views, all of them require (for either knowledge or justification) some strong counter-factual connection with truth.³⁰ Truth, or some connection thereto, is the most prominent externalist criterion, since most assume truth to be of significant epistemic value, though in principle, anything could count as contributing to epistemic justification on an externalist view.

With no end in sight, nor even a pronounced willingness to move on from the topic, getting involved in the discussion could only hinder a complete defense of the DCEJ. If it is committed to either internalism or externalism, then it shall face the stock objections that each side aims at the other. Typically, arguments in support of a commitment to either side have linked deontology with internalism.³¹ The goal of this chapter shall be to present some of the more notable arguments that have been given to the conclusion that the DCEJ entails internalism, to argue against this entailment claim and to offer an explanation of why the association between the two views is so strong. It is worth noting, however, that I shall not try to show that the DCEJ is incompatible with internalism, nor that it entails externalism; the middle path is preferable here.

²⁹ Many externalists will accommodate internal states in their analyses of justification, the idea is merely that they cannot make up the whole story.

³⁰ One might contest this point on Goldman's analysis, which requires that a belief is the result of a reliable cognitive process (1979: 20), depending on how one interprets 'reliable', but this attribution is at least not far from Goldman's position.

³¹ I am not aware of any argument to the conclusion that the DCEJ *entails* externalism, though I shall present a view that is both deontological and externalist in chapter 6.

3.1 Goldman and Plantinga

There are two main motivations for why one would want to connect deontology and internalism. In the first place, one might think that deontology commands a certain independent degree of plausibility, and wish to advance the entailment claim as a means to supporting internalism. Alternatively, one could think that internalism is false and that the entailment claim is true, meaning the DCEJ is false.³²

Goldman addresses attempts to support internalism by appeal to the DCEJ. (1999) He begins with the observation that, on the deontological view, justification is a matter of believing as one is permitted, required or forbidden to believe as determined by one's epistemic duty. (273) This is a slightly more specific definition than what I have given earlier, but is similar enough to conclude that we are talking about the same thing. He continues on to note that there is an intimate relation between the DCEJ and what he calls the guidance idea, which is roughly the claim that whatever theory of epistemic justification we accept ought to guide our epistemic undertakings. Together, the two deliver the result that it is "a person's epistemic duty to guide his doxastic attitudes by his evidence, or by whatever factors determine the justificational status of a proposition at a given time." (273) From here, the argument to internalism is clear. If we are to be guided by our theory of justification, then factors to which we have no access cannot be relevant to justification. Counter-factual profiles of the sort used in notable externalist theories are especially problematic (assuming a standard possible worlds semantics), since facts about other worlds are particularly hard to come by. More strongly, we can say that, in principle, nothing external can be relevant to justification, since anything external is, by hypothesis,

³² In much the same style as the argument from doxastic voluntarism

something to which we lack (immediate) access.³³ The final step is an observation that all of that to which we have immediate access is internal to us. Thus, we are left with internalism. Goldman offers an extended critique of internalism as motivated by the guidance-deontological conception, concluding that, at best, internalism needs to use a different theoretical foundation.³⁴ (293)

If Goldman's argument were sound, it would be problematic for both internalism and the DCEJ, but probably more so for the DCEJ, since it entails a deeply flawed version of internalism, while internalism at least admits of different theoretical foundations. However, exactly why Goldman thinks deontology is connected to the guidance desideratum is unclear. He says both that "the guidance and deontological conceptions of epistemic justification are intimately related" and that "[i]t is possible to separate the deontological conception from the guidance idea." (273) Of course, in order to connect deontology with internalism, something like the guidance idea is required, but in considering whether deontology entails internalism, that would be a question-begging reason for thinking the two are connected. Perhaps he means only that the two views (guidance and deontology) have had a close historical connection. I would accept this much, and shall address this association later in the chapter, but the historical connection obviously provides no indication of a connection in terms of the contents of the views. Even as Goldman has defined the DCEJ, it is clear enough that nothing about guidance or access is required in an attribution of epistemic justification. Without a reason to connect guidance and deontology, therefore, the DCEJ need not commit itself to internalism.

³³ I specify 'immediate' since, as Goldman notes, there are external facts to which we have access. Concerns over this specificity motivate one of the main argumentative points of the article.

³⁴ Which is precisely the strategy adopted by Connee and Feldman (2001) in their well-known response to Goldman's article.

Alvin Plantinga also offers an argument for deontology's entailing internalism, though he does so with less argumentative intent, and more so as a means to a more complete understanding of internalism. To understand his argument, we must first introduce the distinction between objective and subjective duty. Fortunately, as philosophical distinctions go, this is a straightforward one. One's subjective duties are the duties that one believes oneself to have, while one's objective duties are the duties that one in fact does have. I shall illustrate the difference by thinking of a case from ethics, so as to remain as uncontroversial as possible on the topic of epistemic duties. Unfortunately, the most prominent ethical deontologist, Kant, sees objective and subjective ethical duties as identical, which makes his theory an impractical choice for as an example here.³⁵ More amenable to the present task is the work of W.D. Ross.

According to Ross, possessing conclusive knowledge of our moral duties is impossible; rather, we must be content with probable opinion. (2002, 17) A duty that might hold without our (possibly) knowing it must be objective. Similarly, if the view is to maintain any normative recommendations in the absence of knowing one's objective duty, then our subjective duties will correspond to the most probable opinions we can attain. Since merely probable opinion is not certain opinion, there will be possible instances in which our subjective duties diverge from our objective duties.

We can say of theories that equate subjective and objective duties that they are internalist. Since one's subjective duties are a function of one's beliefs, then if subjective and objective duties are the same, objective duties will also be a function of one's beliefs. But it follows from this that there is nothing external to the agent that is relevant to the justificatory status of his or

³⁵ Because the capacity to understand the moral law is a precondition for being obligated by it, all moral agents will be able to determine their subjective duties, and have no duties other than to follow the moral law. (See the introduction to the *Doctrine of Virtue*, 6:380)

her beliefs, which just is the central thesis of internalism.³⁶ Plantinga's argument proceeds by attempting to show that such an overlap is entailed by the DCEJ.

Here is a simplified version of his reasoning. We can tell, in a large and important number of cases, what our objective duty is. (1993:21) If we can do this, then we must be able to recognize the property of our objective duty that confers justification. (21) But whatever property we recognize will also be the property that we recognize in our own beliefs, the property that determines our subjective duty. (22) Therefore, whatever confers justification objectively and subjectively is the same. (23)

This argument is deeply flawed. Most obviously, there is little reason to think that the first premise is true, unless we have already assumed the truth of internalism. A counter-factual analysis of justification might even deliver the result that it is impossible to know if we have fulfilled our subjective duty, since that will depend on facts that are in principle beyond our access.³⁷ Furthermore, even if we could recognize our objective duty and our subjective duty, there is no reason to think that the two kinds of duties need to be generated from the same properties simply because we can determine what each requires. Additionally, the first premise is too weak for the argument's conclusion. Even if we can tell what our objective duty is in a large and important number of cases, internalism is committed to the view that we can do so in all cases. The possible counter-examples might belong to a small and unimportant class, but they would be counter-examples nonetheless. The most charitable interpretation, and, likely, the most plausible, is that Plantinga does not endorse this argument, but rather sees it as illustrative of the

³⁶ We can note that uniquely to epistemology, there is the threat of something resembling the problem of the criterion in this internalist position. If justification is explained in terms of subjective duties, and subjective duties in terms of beliefs, then it is unclear how (or if) we can say whether any beliefs about subjective duties are justified.

³⁷ Maybe the view could be rejected for this very reason, but other, less extreme externalist views have little reason to accept this conclusion.

sort of reasoning used by early deontologists like Locke and Descartes.³⁸ Even so, if this is indeed a fair reading of these two, then this sort of argument has been used by prominent philosophers, and therefore merits a response, given the goal of this chapter.

3.2 Steup and Ginet

Another entailment argument comes from Matthias Steup. Since, as already noted, Steup is an internalist and a defender of the DCEJ. He goes as far as to suggest that internalism in general is committed to thinking of justification deontologically. (2001:135) I do not think this is necessarily true, but only Steup's commitments are at issue here. The argument moves from deontology to internalism, so we can assume that he takes the former to be the less controversial position. His argument adds a degree of complexity to the sort of argument given by Goldman, since he moves from deontology to evidentialism, and from evidentialism to internalism. As a final point of contrast with Goldman's argument, Steup specifically dismisses the guidance desideratum that Goldman attributes to internalists (147, endnote 7), meaning if I am to disagree with Steup's argument, it must be on different grounds than those on which I objected to that of Goldman.

In something of a recurring theme of the project thus far, Steup has his own view on what, exactly, the DCEJ is. His preferred analysis is as follows: "S is justified in believing that p at t if and only if it is epistemically responsible of (or permissible for) S to believe that p at t ." (136) Reflecting his particular strategy in responding to the argument from doxastic voluntarism, he thinks that deontology can, by itself, generate an accessibility constraint on justification, such that only that which is accessible is relevant to the justificatory status of one's beliefs. (136) This premise is motivated largely by Laurence Bonjour's famous counter-examples to reliabilism.

³⁸ Bergmann (2006:81) advocates for a similar interpretation.

Simplifying the cases a bit, the suggestion is that, even if one is in possession of a perfectly reliable belief-forming faculty, if one lacks (or is in possession of defeaters for) good reason to think that this faculty is reliable, it is not epistemically responsible for one to believe in the manner produced by said faculty (1980). Since deontological justification is a matter of epistemic responsibility, that to which one has access is what determines the justificatory status of one's beliefs. Conceding Goldman's aforementioned point³⁹ about the possible divergence of the internal and that to which we have access, his next move is to argue from deontology to evidentialism. If this argument succeeds, then, even if the accessibility constraint does not entail internalism by itself, its conjunction with evidentialism does.

In truth, Steup's argument from deontology to evidentialism seems more like an argument to the unconditional truth of evidentialism. Reflecting once more on BonJour's examples, what seems to make the beliefs irresponsible is their lack of support by something in the believer's cognitive possession, but all and only that which is in our cognitive possession and contributes to epistemic justification is our evidence. (Steup, 2001:137)⁴⁰ But since evidential states are internal, we have the elements in place to show that the DCEJ entails internalism.⁴¹ In sum, deontology entails an accessibility constraint, which in turn entails evidentialism, which lastly suffices for internalism.

The first major issue with Steup's argument is his definition of deontology. By itself, defining justification in terms of responsibility is perhaps not so problematic, but by making

³⁹See footnote 6.

⁴⁰ Admittedly, this is a premise for which Steup recognizes a need for support, but he notes that such an argument would be too large an undertaking for his purposes. If my arguments work as I intend, we can grant this premise and still find fault in the overall argument.

⁴¹ Steup notes that this is not a "full-fledged" internalist view, since it will allow factors like reliability to play a justifying role, so long as there is a mediating evidential state. (138) This distinction should not prove relevant for our purposes.

responsibility a function of that to which one has cognitive access, he assumes that which he wishes to prove. In the first place, note that he takes BonJour's intuitions about epistemic responsibility as a premise. From here, there are two ways of interpreting how Steup's argument proceeds, since, as mentioned in footnote 13, we do not know how Steup sees deontology entailing internalism, we need to consider a few options to cover all possible strategies. In the first place, let us assume that evidentialism is a general truth about epistemology, and does not require the DCEJ for its truth. But if this is the case, then deontology only entails internalism in the most trivial of senses; since evidentialism is a necessary truth about epistemology, everything will entail it. So according to this option, deontology bears no significant relation to internalism. Alternatively, and perhaps more generously, we can assume that the argument in favor of evidentialism will appeal to some facet of the DCEJ, which, on Steup's view, is essentially about epistemic responsibility, which in turn is determined by that to which one has cognitive access. But, as we have already noted, according to many, the access requirement is the essential feature of internalism. On this interpretation, evidentialism plays no real role in the entailment. Of course, even if evidentialism were left out, the argument, if sound, would still show that the DCEJ entails internalism. By now, the trouble with Steup's original analysis of deontology should be clear. If taken in a very broad sense, responsibility could be treated as more or less the same thing as adhering to one's duty, but that is not how Steup uses the term. Rather, he takes it to be a function of the internal, so deontology entails internalism by definition.

I have shown that Steup's argument perhaps does not work exactly as he attests, but not that the argument is fatally flawed. If deontology entails internalism by definition, so much the better for those in support of the entailment claim. We may at last note that BonJour's argument

is about the inadequacy of externalism.⁴² While well received, the examples used in his paper have not ended the debate between internalists and externalists. An externalist could deny the legitimacy of Bonjour's examples and maintain that we only have an epistemic responsibility to hold beliefs that are the result of reliable cognitive processes. Steup could object to this claim, but the debate would no longer be about deontology. Instead, it would simply be a regression to the original debate between internalists and externalists. Obviously, I prefer the analysis of deontology I give in the introduction; by starting neutral on the question of internalism, the analysis remains neutral in the absence of a definitive argument in favor thereof.

I think we can trace the sort of reasoning Steup uses back to the argument from doxastic involuntarism. We saw in the second chapter that his strategy was not to deny that the 'oughts' of deontology imply that one can do as one is obligated, but rather to accept that implication, and to claim that we have sufficient control to meet these obligations. Since Steup persists in thinking of beliefs as actions, (1988:65) even on an epistemic compatibilist view, an externalist deontology would make obligatory beliefs such that there would be no way one could act to make oneself so believe, or forbidden some beliefs such that there would be no way one could act to make oneself not so believe. This would, of course, break Kant's law, so we reject externalist deontology. If I have succeeded in chapter 2, then the reader will understand this argument to be flawed. Beliefs are not actions, and thus do not imply the sense of 'can' that Steup's reasoning employs. While the shift to epistemic compatibilism perhaps allows deontology to bypass the argument from doxastic voluntarism, the move commits the DCEJ to internalism, which, for reasons I have already mentioned, is to be avoided if possible. If there is

⁴² Specifically, about the results they deliver on empirical, or foundational, knowledge (1980:56)

another way of showing that the DCEJ does not succumb to the argument from doxastic voluntarism, as I have argued there is, then the DCEJ need not commit itself to internalism.

Carl Ginet's argument for an awareness requirement on justification also begins from the assumption of deontology. To begin, let us note that Ginet understands justification to be a matter of justifying a specific degree of confidence. (1975:12) This allows his theory to be more versatile than one that considers merely whether instances of full-fledged belief should qualify as justified, but should not interfere with its applicability here. I shall present his argument structurally for clarity's sake.

(1) "Assuming that *S* has the concept of justification for being confident that *p*, *S* *ought* always to possess or lack confidence that *p* according to whether or not he has such justification."⁴³

(2) What *S* ought to do, *S* can do.

(3) *S* can always tell whether or not he has justification for being confident that *p*. [from 1 & 2]

(4) *S* would not always be able to tell whether or not he has justification for being confident that *p* unless the difference between having such justification and not having it were always directly recognizable to *S*.

(5) If any fact that was not directly recognizable (or entailed by something that was directly recognizable) contributed to *S*'s justification for being confident that *p*, then the difference between having such justification and not having it would not always be directly recognizable to *S*. (36)

Therefore, we conclude that no fact that is not directly recognizable by a given subject contributes to the subject's justification to be confident in a given proposition. This latter claim suffices for the truth of some variety of internalism.

⁴³ Ginet uses the term justification as something that one has, rather than as a property of beliefs (or doxastic attitudes in general). I shall continue to think of the term in the latter sense, since if we assert that justification is necessarily *had*, then we have made an unfair assumption in favor of internalism. Ginet equates the two senses of the term (28) so I take it that this move on my part is not an equally unfair assumption *against* internalism.

A preliminary concern, similar to one raised against one interpretation of Steup's argument, is that the argument is really just an independent claim to the truth of internalism, rather than an argument from deontology to internalism. It is at least clear, however, that Ginet adopts a deontological view of epistemic justification: "Circumstances that justify a person in being confident that something is the case are generally also enough to oblige him to be confident..." (28) The modality of the claim is not quite strong enough to make his analysis of justification identical to my own, but it is close. Of the premises given, (2) seems the most likely to be indicative of Ginet's accepting the DCEJ. At this early stage in the project, it is still easiest to illustrate why this is so by thinking of an example from ethics. On some versions of consequentialism, what one ought to do is that which maximizes utility. It may be the case that one cannot perform the action that will do so. One might not be subject to blame as a result, but there is still a significant sense in which the 'ought' can hold without an accompanying 'can.' Deontological ethical views are committed to an unconditional applicability of Kant's law; I have also endorsed a similar understanding of the rule as part of the DCEJ in the introduction. Since Ginet's second premise is also unconditional, we can plausibly interpret the whole argument as making a non-trivial claim that deontology entails internalism.⁴⁴

There are a few ways of objecting to Ginet's argument, depending on where one wishes to start. My particular suspicion is that (3) is false, though roughly the same argument I shall give below could be used to suggest that (2) is false. It need not be the case that we can always tell whether we would be justified in a particular belief. One can possess or lack confidence that *p* according to whether belief that *p* would be justified without being able to do so at will. Indeed, for reasons discussed in chapter 2, we have good reason to think that we would not be

⁴⁴ I specify 'non-trivial' since, as noted before, if internalism is necessarily true, then everything will entail it.

able to do so at will. If it need not be the case that we can believe at will, then the requirement stipulated by (3) is unmotivated. Indeed, the argument from doxastic voluntarism gives us good reason to think that the inference from the first two premises to the third is faulty.⁴⁵ With both Steup and Ginet's arguments shown to be flawed, we are now in a position to make some general observations on the relation between internalism and deontology.

Conclusion to Chapter 3

My reason for thinking there is no necessary connection between internalism and deontology is fairly simple, and has been reflected in my criticisms of the preceding arguments. Deontology is about fulfilling one's epistemic duty; some deontological theories recognize the existence, and possible divergence, of both subjective and objective duties. Internalism entails that our subjective and objective duties are the same. Without further support for internalism or externalism, there is no reason to think deontology suggests either. Indeed, as I shall suggest in chapter 5, on the deontological view, internalism and externalism should be thought of as explanations of what our duties are, not as providing a case for why we should do our epistemic duty.

I confess that I find the proliferation of arguments connecting deontology and internalism perplexing. I shall use the remainder of this chapter to diagnose why this connection has been made, and to see if the underlying reasons for this connection can play any role in a continued defense of the DCEJ.

One reason to consider is that both deontology and internalism are integral to what we might call traditional epistemology. I have already noted that two of the luminaries of early-modern epistemology, Locke and Descartes, apparently adopt some sort of deontological

⁴⁵ In fact, as we shall note in the next chapter, Alston's model of deontology on which he makes his attack is this very passage from Ginet.

position. Of course, these two were also (even if applying the term is something of an anachronism) internalists, since everyone was through at least the first half of the twentieth century. Another strong representative of both views, Chisholm, carries the association into contemporary philosophy.⁴⁶ Perhaps externalists have simply taken aim at 'traditional epistemology' and, in the process, conflated two independent parts of that tradition. In response, we could explain the tendency of internalists to make the same mistake as motivated by a desire to defend the traditional project as a whole. This is perhaps an uncharitable accusation to make of many brilliant philosophers, but the other likely explanation, that no one has seriously reflected on what, exactly, makes a view deontological, seems just as much so.

Referring once more to the notion of direction of fit may yield a more satisfactory explanation. Kantian ethics is clearly internalist with regard to its criteria of moral rightness; that is to say, on this view, ethical justification supervenes on factors internal to the relevant agent. This must be so because the success of an action involves fitting the world to one's will. If ethical justification is a matter of having the right state of will, then internalism in ethics works well because the will itself is internal to the agent. Internalism in deontological epistemology works comparatively poorly, since beliefs will succeed in virtue of their accurately representing the world, which is (except in a narrow set of exceptions) not internal to the agent. It follows that we have little reason to think that the proper account of epistemic duties will only make use of that which is internal to the subject of epistemic evaluations. One potential explanation, therefore, of the standard but mistaken association between epistemic deontology and internalism is that the latter is sensibly associated with the ethical counter-part of the former. Since deontological ethics are comparatively much better developed, a natural move might be to incorporate as much as one

⁴⁶ Both theories are prominent elements of the epistemological system he outlines in his classic *Theory of Knowledge* (1966).

can from the ethical into the epistemological. On the topic of internalism, however, I have shown why this should not be done.

I have to this point failed to mention one of the more promising arguments to the conclusion that deontology entails internalism. It begins with the assertion that deontology is essentially about blame, blameworthiness or blamelessness. Standard externalist criteria for epistemic justification are unable to equate justification with blameworthiness (or blamelessness). Therefore, deontology can only appeal to internalist criteria in analyzing justification. I have omitted this argument because it has been the subject of much independent attention, and merits a degree of consideration sufficient to constitute a chapter in its own right. As such, the conclusion of this chapter should be treated as, at best, provisionally successful. Only if I can show that the argument from blame is also flawed will I have shown with any degree of finality that deontology does not entail internalism.

CHAPTER 4

DEONTOLOGY AND BLAME

A very prevalent theme in the literature on deontology in epistemology is the centrality of the concept of blame to a deontological analysis of justification. The following authors exemplify this theme: Alston (1988), Audi (2001), Bergmann (2006), Booth (2008), Feldman (2008), Ginet (1975), Nottelmann (2013), Plantinga (1993), Russell (2001) and Steup (1988). The goals of this chapter are to investigate the reasoning behind this association, to consider whether this association is likely to be of use in a defense of the DCEJ and, ultimately, to argue that there is no essential commitment on the part of the DCEJ to an analysis of justification in terms of blame, or any related concept.

In the conclusion to the previous chapter, we noted that one possible motivation for thinking that deontology entails internalism is that it is committed to analyzing justification in terms of blameworthiness, but since blameworthiness supervenes on the internal, this commitment suffices for the entailment argument. Following the considerations given in the introduction to the previous chapter, this gives a *prima facie* reason for the DCEJ to avoid commitment to a blame-based analysis. Nonetheless, there may yet prove to be an argument in favor of such a commitment the denial of which would harm the prospects of the DCEJ more than avoiding a commitment to internalism would help it. There are two types of arguments that draw on this observation. First, one might claim that the DCEJ is obligated on pain of inconsistency to include blame in its analysis of justification. Alternatively, one might say that blame is a favorable inclusion in such an analysis, even if there is no aspect of the DCEJ that

necessarily commits it to such an inclusion. I shall call the former strong blame arguments and the latter weak blame arguments.

4.1 Strong Blame Arguments

Despite the widespread association between deontology and blame, there are comparatively few arguments that motivate this connection; in most instances, it is merely assumed. One of the more overt associations of deontology and blame is found in William Alston's article "The Deontological Conception of Epistemic Justification" (1988), which was discussed at length in the second chapter. While the argument from doxastic voluntarism is the aspect of the article that has drawn the most attention, the article also addresses the concerns of the present chapter. Alston, quoting Ginet (1975), gives the following as a first approximation of an analysis of deontology: "One is *justified* in being confident that p if and only if it is not the case that one ought not to be confident that p; one could not be justly reproached for being confident that p." (256) In addition to concerns that Alston raises in the rest of his article, this clearly will not do as an analysis since it introduces a form of the analysandum in the appeal to just reproach. Alston's subsequent remarks fit better with what I have identified as the mark of the deontological: "this conception of epistemic justification is viable only if beliefs are sufficiently under voluntary control to render such concepts as *requirement*, *permission*, *obligation*, *reproach*, and *blame* applicable to them." (256) I did not identify blame and reproach as central deontic concepts, but Alston's requirement could be made to work with my own if we think of (epistemic) blame and reproach as a function of the degree to which one fulfills one's obligations, which would be expressed in terms of the former three terms. Halfway through his argument, however, the relation has been inverted: "[t]he upshot of the paper thus far is that the only viable deontological conception of justification is the one that identifies being justified in

believing that p with not being intellectually to blame for believing that p...” (284) I shall attempt to discern the reasoning behind Alston's switch to espousing the centrality of blamelessness.

To be sure, part of the explanation is Alston's acceptance of the argument from doxastic voluntarism. Since we cannot determine our doxastic attitudes at will, some varieties of deontology will have no prospect of success, for reasons already discussed at length. However, since he is willing to admit that we have control over our actions, some of which determine our doxastic attitudes, deontology improves its prospects if it posits indirect epistemic duties: duties concerning actions that have epistemic consequences. (280) Since actions involve intentions, and one is responsible for one's intentional actions, epistemic justification, on this view, becomes a matter of responsibility. From here, the move to the concept of blame is simple enough; one can be blamed for intentional violations of epistemic obligations. The preference of blamelessness to blameworthiness derives from the aforementioned aversion to asserting positive epistemic obligations. (283) Why blamelessness becomes the analysans rather than responsibility is not clear to me, but I think one can give a more definitive objection to Alston's argument.

The reader will likely have anticipated my concerns about the above argument. Since the argument from doxastic voluntarism is unsound, the argument to the centrality of blamelessness, which takes the success of the former as a premise, is also unsound. Specifically, we have no reason to think that justification must be understood indirectly, in terms of action. Absent any such reason, there is equally little reason to think that justification is about responsibility, and, hence, blamelessness.

Alston's argument does illustrate an important lesson for one of the topics of enquiry for this chapter. If justification were a matter of blame (or a related concept), then it would also be true that justification supervenes on the internal. But if we consider why Alston thinks

justification is a matter of blamelessness, it is likely because he has already assumed that deontology entails internalism. That he makes such an assumption is reflected in his attribution to deontology that it requires our doxastic attitudes to be voluntary. I have shown that, without a commitment to internalism, deontology need not commit itself to the view that doxastic attitudes are voluntary. This sort of attempt to derive internalism from the importance of blame to deontology is, therefore, circular. In sum, the argument runs as follows: deontology is essentially concerned with blame, or some closely related concept, because its evaluations are essentially a matter of epistemic responsibility. Its evaluations are essentially a matter of epistemic responsibility because it needs to define justification in terms of intentional actions. It needs to define justification in terms of intentional actions because its evaluations require voluntary control, which we lack with regard to our doxastic attitudes. Its evaluations require voluntary control, however, because ought implies can, and one can only control that which is internal to one. I have questioned the legitimacy of this application of Kant's law, but it should be clear that the reasoning behind it is grounded in some kind of internalism. It seems unlikely, therefore, that an argument grounded in the traditional argument from doxastic voluntarism will convincingly show that deontology entails internalism through its reliance on the concept of blame.

In an extended attack on epistemic internalism, Michael Bergmann (2006) considers a similar question to that asked by the previous chapter: does deontology entail internalism? In particular, his concern is with those who attempt to provide independent support for internalism by showing that it is entailed by the DCEJ. (77) In keeping with the concern of this chapter, he notes that one of the more common attempts to bridge the two concepts is by an appeal to the concept of blamelessness. (90) Thus, while his argument is relevant to the questions at hand, it is not concerned with precisely the same issue. It does not address whether the DCEJ must or

should appeal to the concept of blame. Rather, it attempts to show a disconnection between blame and internalism. Therefore, even if the strong blame argument succeeds, if Bergmann's argument also succeeds, the DCEJ is not committed to internalism. While I have argued that the strong blame argument fails, it is salubrious to my overall position that I need not have succeeded in my earlier claim.

To get a good understanding of Bergmann's point, we must refer back to a distinction made earlier: the difference between objective and subjective duties. Clearly, one's subjective duties will be better connected to the concept of blame. Think of a case wherein one's subjective and objective duties diverge; one thinks one ought to believe something, when in fact one ought not so to believe. If one adheres to one's subjective duty, it is clear that one deserves no blame for doing so, even if, objectively speaking, the belief is unjustified.⁴⁷ Rather generously to the internalist, Bergmann's argument proceeds on the assumption that the DCEJ is concerned with subjective duties. (90) With the centrality of subjective duties in mind, we can introduce the following analyses of blameworthiness (EBW) and blamelessness (EBL):

EBW: S's believing p at t is something for which she is epistemically blameworthy if and only if either (i) S believes at t that she ought not to believe p or (ii) S's failure to believe at t that she ought not to believe p is relevantly due to some other doing or failure of hers for which she is epistemically blameworthy...

EBL: S's believing p at t is something for which she is epistemically blameless if and only if (i) S doesn't believe at t that she ought not to believe p and (ii) S's failure to believe at t that she ought not to believe p is not relevantly due to some other doing or failure of hers for which she is epistemically blameworthy. (92, emphases removed)

We are now in a position to consider an argument from blamelessness to an awareness requirement, which Bergmann identifies as the key aspect of internalism. (9)

⁴⁷ We have already noted that an internalist will identify one's subjective and objective duties, so I might be suspected of beginning with an unfair hypothesis. I only need the example to function conditionally: if there are such instances of divergence, blameworthiness is more clearly attached to subjective duty.

He considers five versions of the same basic argument, but I shall try instead to give an inclusive summary for brevity's sake. The essential claim behind the blamelessness argument is that, if justification is a matter of not breaking one's subjective duties, then one will be justified only if one does not believe that, for a given belief, one ought not so to believe. This, however, does not suffice for an awareness requirement; it only suffices for an absence of awareness requirement (specifically, the absence of a defeater). (94) Shifting the argument to blameworthiness does no better, since one can be derivatively blameworthy for a belief without (necessarily) having access thereto, as specified in clause (ii) of EBW. Therefore, even if the DCEJ is committed to an analysis of justification in terms of blame, or some closely related concept, it does not follow that it is also committed to internalism.

It seems likely that much of the tendency to associate the DCEJ and blame is derived from thinking that the DCEJ is importantly related to internalism as well, since there is no obvious connection between blameworthiness/blamelessness and objective duties. If objective and subjective epistemic duties can come apart, then blame only pertains to the latter.⁴⁸ It follows, then, that our concern with internalism is concluded; there is no positive reason to think that the DCEJ is inconsistent with externalism, nor that it entails internalism. Nonetheless, there may be more of value in thinking about the relation between deontology and blame. It remains to be seen whether it would not be best for the viability of the DCEJ for it to embrace a blame-based analysis of justification. If this were so, then it would prove to be an invaluable datum in gaining a complete understanding of what the DCEJ is, and what it must say.

⁴⁸ This is perhaps a controversial assertion; I shall defend it in the last section of this chapter.

4.2 Weak Blame Arguments

Nikolaj Nottelmann's article *The deontological conception of epistemic justification: a reassessment*, is a direct attack against the DCEJ.⁴⁹ My treatment of the article will not involve the majority of the argument, since it will become clear that he and I disagree about a fundamental premise in his argument. Importantly for our purposes here, however, Nottelmann reaches the conclusion that the DCEJ, while deeply flawed in any instantiation, fares best when it defines justification in terms of blamelessness. I shall presently outline his reasoning in support of this point, and shall respond on behalf of the DCEJ, indicating the locus of our fundamental disagreement.

As has become something of a theme for this project, we must begin with Nottelmann's understanding of what, exactly, the DCEJ is. He articulates two necessary criteria:

No non-deontic sufficiency condition (NNDSC): No statement containing only non-deontic terms logically entails a statement affirming a subject's EJ [epistemic justification]. (2013:2222)

Deontological entailment (DE): Any statement concerning an agent's EJ for a belief logically entails a statement applying deontic predicates to the agent... (2223)⁵⁰

Given these two clauses' criteria, he argues, the typical analysis offered by the DCEJ, one that appeals to the concept of permissible belief, will not do. In the first place, there are apparently beliefs that, while epistemically justified, are not permissible, due to, say, overriding ethical concerns. (2226) The natural response on behalf of the DCEJ is to claim that the sense in which 'permissible' is employed in the analysis is only concerned with *epistemic* permissibility. The move to this revision, however, leaves the analysis at a loss to accommodate the cacophonously abbreviated NNDSC. Saying that some belief is epistemically permissible for some agent A and

⁴⁹ In fact, the abbreviation itself is Nottelmann's idea.

⁵⁰ He later suggests that these two conditions suffice for the DCEJ, (2237) so we can interpret these as a classical analysis.

proposition *p* is equivalent to saying “A’s epistemic reasons, *pro tanto*, tell in favour of believing *p*.” (2229) But we can invariably explain one’s epistemic reasons in non-deontic terms (something like ‘*p* was best supported by A’s evidence.’) If this is the case, then we can give a complete account of a belief’s permissibility and, hence, its justification, without using any deontic terms, meaning NNDSC is not met. (2229) As such, on Nottelmann’s analysis, the DCEJ cannot analyze justification in terms of mere permissibility, since any theory that gives such an analysis is not the DCEJ.

Apparently in keeping with the list of deontic operators given by Alston, (1988:257) it is at this point that Nottelmann argues for the favorability of a blamelessness-based analysis. Since blamelessness is a deontic term, an analysis of justification in terms thereof will necessarily meet the NNDSC. Similarly, any statement that affirms of some agent that he or she is justified in a belief will entail that he or she is blameless in so believing, thereby satisfying DE. At the very least, then, using the concept of blamelessness allows for a conception of epistemic justification to be deontic. Nottelmann continues his argument by arguing that even this preferable understanding of the DCEJ fails, but what has been said to this point shall suffice for my response to his general position.

I shall attempt to illustrate what I see as the problem with Nottelmann’s view by carrying out a *reductio* of sorts with his blamelessness-based version of the DCEJ. When we are blameless or blameworthy, there is inevitably some reason *why* we are so. We deserve blame because we broke a rule that we ought not to have broken, for instance. If this is true, however, then we can give a statement affirming a subject’s epistemic justification without using any deontic terms:⁵¹ (1) Necessarily, if one breaks a rule that one ought not to break, then one is

⁵¹ Or, more specifically in this case, *denying* a subject’s epistemic justification, but I trust that an example resulting in an affirmation can be constructed just as easily.

blameworthy. (2) Necessarily, one is justified in believing *p* if and only if one is blameless in believing *p*. (3) If one is blameworthy in believing *p*, one is not blameless in believing *p*. (4) In believing *p*, *A* has broken a rule that *A* ought not to have broken. Given the previous three premises, (4) is a statement containing only non-deontic terms logically that entails a statement denying a subject's epistemic justification. It follows that, if Nottelmann's criteria are correct, the blamelessness-based analysis is no more deontological than the analysis in terms of permissibility.

Nottelmann may not be bothered by this result; after all, his general conclusion is that the DCEJ is deeply flawed, and the considerations given above only seem to make the case worse. Obviously, I am inclined to interpret the result in another way: Nottelmann is mistaken in requiring the NNDSC of the deontological conception of epistemic justification. I shall now explain why I see the NNDSC as an unreasonable requirement.

The only way the blamelessness analysis could work would be if we took blamelessness (or perhaps blame, if that is a sufficiently deontic term) to be a conceptually primitive term. However, it is clearly not such a term, or at least is no better a candidate than 'justification' itself. It should be clear, then, that the deck has been stacked, so to speak, against the DCEJ. Assuming the NNDSC is true, unless there is a deontic term that carries a set of epistemic rules unto itself, then there is no way for *any* theory to qualify as legitimately deontic. But this surely misconstrues what deontic terms are. As noted in the introduction, the central deontic terms are operators, not norms. The task of deontic terms in a theory of justification is not to tell us what the true epistemic rules are, but rather to tell us in what relation we are to stand with regard to those rules. Since justification will be a function of this standing, the content of those rules matters nothing to whether or not the theory is deontological. To require of a deontological

theory that it not allow any statement that does not use deontic terms to suffice for an affirmation of justification is to require that it deem *that* we have epistemic duties, but also to require that it remain silent as to *what* our epistemic duties are. Worse still, if there are any epistemic duties, then there must be some statement that does not use any deontic terms that suffices for an affirmation of epistemic justification, namely, a description of a doxastic profile that is obligated and/or permitted by one's epistemic duties. Hence, the conclusion of the weak blame argument is not accomplished. On Nottelmann's analysis of the DCEJ, blame-based accounts of justification fare no better than any other accounts, because his analysis discounts *any* account of justification. If one agrees with my rejection of the NNDSC, one is left without a reason to prefer blame-based accounts of justification, because there is no need to include deontic terms in the correct set of epistemic norms.

Even if my more ambitious arguments above fail, Brian Weatherson gives a more moderate argument to the conclusion that praise is a more useful concept to employ in explaining epistemic justification. The argument begins by considering a modification of the Cartesian 'Evil Demon' thought-experiment. Consider my epistemic counter-part, who, while in possession of the same doxastic profile as I am, is the victim of the demon's machinations. My belief that I have hands is (presumably) true, and grounded in the fact that I am appeared to handily, that this belief coheres well with other beliefs and experiences I have, etc. My counter-part's same belief is false, but inferred from the same reasons. As such, if I do not deserve any blame for my beliefs, then neither should he. After all, there is nothing that he has done wrong in coming to have the beliefs he does; in fact, he generally reasons quite well. If justification is merely a matter of blamelessness, then my epistemic counter-part is fully justified. This, however, seems to be the wrong result. While neither of us deserves blame, my beliefs merit a higher standing

than his do, since mine are true. Appealing to praise allows us to solve this problem; although my counter-part does not deserve blame, he also deserves less praise than I do, thus my beliefs are afforded the superior status. (Weatherson, 2008: 565)

A defender of a blame-based analysis has a plausible response to this objection available. Following Gettier (1963), justification is understood to be insufficient to make true belief into knowledge. One could explain the additional positive standing enjoyed by my beliefs in terms of a fourth condition (the condition that, if met, turns justified true belief into knowledge) while still holding that my counter-part is fully justified, thereby maintaining consistency in analyzing justification in terms of blamelessness.

Weatherson has a response in anticipation of such an objection. He begins with two assumptions, abbreviated as A1 and A2, respectively, below. First, it is possible that there are justified, false beliefs. (567) Second, blamelessness is transitive; that is to say, if one believes something on the basis of a belief for which he or she is blameless, then this new belief is also blameless. (568) Weatherson gives a proof of the inconsistency of these two assumptions with an equation of justification and blamelessness (abbreviated below as $j=b$). For some subject S, and some proposition p:

- (1) S justifiedly, but falsely, believes that she is justified in believing p. (Assumption, A1)
- (2) On the basis of this belief, S comes to believe that p. (Assumption)
- (3) S blamelessly believes that she is justified in believing that p (1, $j=b$)
- (4) S blamelessly believes that p (2, 3, A2)
- (5) S is justified in believing that p. (4, $j=b$)
- (6) It is false that S is justified in believing that p. (1) (568-569, emphases removed)

Neither of the assumptions are plausible candidates for rejecting in response to the *reductio*, at least not if one is committed to $j=b$. The initial considerations about blame in response to

Cartesian Evil Demon cases suggest that one can be blameless, and hence justified, even in cases

of false beliefs. Similarly, if blamelessness is not transitive across inferences, then it is unclear why it should serve as such a key concept. In other words, if our concept of justification is to indicate positive status, then fully justified beliefs should be safe bases for inference. One might plausibly deny A2 from an externalist perspective, but, as I have argued throughout this chapter, if one is an externalist, there is little reason to accept (and much reason to reject) the claim that justification is merely blamelessness. It follows, then, that we ought to reject $j=b$.

Conclusion to Chapter 4

Combining the conclusions of the two sections, we are left with the observation that the DCEJ neither must nor should analyze justification in terms of blame or any closely related concept.⁵² Also, as was noted at the end of the first section, we have seen that there is no positive reason to think that the DCEJ entails internalism, since the argument from the DCEJ to blame and from blame to internalism fails. Given the widespread acceptance of these associations, it merits some consideration why, if I am correct, so many philosophers have proven to be so mistaken about what the DCEJ is.

Reflecting on the three central topics of chapters two through four will help to illuminate the issue. On the one hand, the DCEJ is saddled with implausible applications of a deontological *ethical* theory. The most famous deontological ethical theory is committed to voluntarism about actions and internalism about justification.⁵³ It is therefore natural to think that the same might be true of deontology in epistemology. However, this line of thought does not survive critical reflection. Since beliefs have a different direction of fit than actions, the inclusion of a voluntaristic element and an internalist criterion of justification, while natural in the ethical

⁵² To clarify a possible scope ambiguity in this sentence: I claim to have shown that it is not the case that the DCEJ should analyze justification in terms of blame, not that the DCEJ should analyze justification in terms of blame.

⁵³ Although we have already noted that not all ethical deontologists must do this.

sphere, is only deleterious to a well-conceived epistemological theory. Without a necessary connection to internalism, the connection to blame is also implausible. That for which we deserve blame is a subset of that over which we have control. If the DCEJ can require for the justification of a belief something that is external to, and not under the voluntary control of, an agent, then there can obviously be cases of blameless, unjustified beliefs.

Even if both concepts are excluded from the final analysis of justification, their treatment thus far has indicated that internalism and blame are importantly connected. We might, therefore, explain one in terms of the other, such that internal justification is equivalent to blamelessness. If internalism is true, then it will follow that the DCEJ should commit itself to a blame based analysis. It will be the goal of chapter 6 to argue that the DCEJ is best served adopting a variety of externalism about epistemic justification.

There is another widespread misunderstanding about the DCEJ, which is best exemplified by Nottelmann's argument. As I understand it, given only the information that a theory is deontological, one is not in a position to determine whether or not any belief is justified. Beliefs that satisfy one's epistemic duties are justified, but there need not be only one account of what those epistemic duties are. Clarifying this claim is of the utmost importance to the completion of this project. To this point, the work has been largely negative, i.e. rejecting claims about what the DCEJ must say. Now, it is time to turn our attention to what the DCEJ actually does say. In doing so, we will come to have a better understanding of its proper opposition, and the more serious set of challenges it faces.

CHAPTER 5

METAEPISTEMOLOGY

There is a popular notion in the study of ethics according to which the discipline can be divided into three main areas of inquiry: metaethics, normative ethics and applied ethics. In all cases, the names are admirably informative. Metaethics is, generally, the investigation of that which is presupposed by the study of ethics. Normative ethics is concerned with the identification of moral norms, or rules. Applied ethics is the study of how to apply ethical theories to particular situations.⁵⁴ We can also observe that a complete ethical theory, even if only an idealization, would cover all three of these areas. Despite the general acceptance of the aptitude of an over-arching analogy between ethics and epistemology, there has been little use of a division in epistemology analogous to the one described above. I shall defend three main claims in this chapter. First, and least controversially, these level-divisions could prove to be a useful way of thinking for epistemology in general. Second, adopting these divisions allows us to reach a clearer understanding of various positions on the topic of epistemic justification. Third, and most controversially, the deontological conception of epistemic justification is the best meta-epistemological position of the prominent contenders.⁵⁵

Before beginning in earnest, the claims I have made thus far may face an immediate objection. I have stated my intention to defend deontology as a position in meta-epistemology.

⁵⁴ None of these is meant to be a contentious claim; I am confident that, even if I am mistaken here, my reference to these concepts in the rest of the project will cohere with any normal understanding of these terms.

⁵⁵ Here I encounter something of a problem since, while the epistemological analog of metaethics has not been an area of research, there is still a significant literature on meta-epistemology. The way the term has been used is closer to what we might call 'the epistemology of epistemology' rather than an attempt to investigate that that is presupposed by epistemological inquiry. Nonetheless, I shall make use of the term as I describe it in the body of the text. With hope, this footnote shall be sufficient to ward off any confusion resulting from the semantic ambiguity.

This would suggest, by analogy, that deontological ethical theories are properly conceived as metaethical views. This assertion runs contrary to the accepted view, according to which deontology is a position within *normative* ethics.⁵⁶ I think this view is mistaken in such a way as to be explained easily.

Any particular deontological theory one considers will also have a normative element, since a theory will necessarily be incomplete if it omits the central element.⁵⁷ One might have an ethical theory that is difficult to apply (i.e. one that lacks a developed applied ethical view) or that relies on implausible assumptions about the nature of moral goodness (e.g. one that makes a controversial assumption about moral skepticism). However, it seems that an ethical theory, in order to count as such, must address the questions of normative ethics. If we then consider what all deontological theories have in common, it is not to be found at the normative level. Since deontology is essentially committed to explaining justification (both epistemological and ethical) in terms of duty fulfillment, rather than to a specific account of what one's duties are, it is a meta-theoretical position.⁵⁸ Thus, even if all deontological views give an expression of ethical norms in terms of fulfilling duties, their theoretical motivation for doing so is to be found in the metaethical view that moral goodness consists in doing one's duty.

It is also worth noting that there are different questions that can be considered at the meta level. It will help to note that I am only concerned with one such question here, namely, the question of the nature of epistemic justification. A deontologist and a consequentialist will disagree about the nature of epistemic justification, but could agree on some other meta-

⁵⁶ As suggested by the Stanford Encyclopedia of Philosophy's entry on deontological ethics (Alexande, 2007).

⁵⁷ 'Central' in terms of position, since the meta level is "higher" and the applied level is "lower."

⁵⁸ E.g., Kant and Ross are both committed to the view that we have ethical duties, and that fulfilling these duties is constitutive of ethical goodness. They disagree about what duties we have, and the modality of our obligations to these duties. Their agreements suffice to make them both deontologists; their disagreements are over topics addressed in normative ethics.

epistemological issue. Furthermore, two deontologists would agree about epistemic justification, but could disagree over other meta-epistemological issues.

5.1 Deontology

Having argued that the DCEJ is a position properly understood as situated within meta-epistemology, we are now in a position to give a complete explanation of what the view is. Specifically, to the meta-epistemological question ‘what is epistemic justification?’ it is the answer ‘the degree of similarity between an agent’s actual epistemic properties and those suggested by the agent’s epistemic duty.’ As is typical for such a bold articulation of a key concept, a few points of clarification are called for immediately. In the question asked above, I am thinking of a broad sense of the concept, one that incorporates the presence and absence of epistemic justification. That is to say, actually having epistemic justification (or being epistemically justified, I make no distinction between these) is a matter of having a high degree of similarity between one’s actual epistemic properties and those suggested by the agent’s epistemic duty. A lack of epistemic justification, on the other hand, is a matter of having a low degree of similarity therebetween. Phrasing the analysis in terms of similarity with the ideal state allows for a broad evaluation of an agent’s justificatory status. We can put the analysis in terms of beliefs, or other epistemic states as a particular response to a particular epistemic duty. The broad formulation allows for a simple expression that accounts for all of our epistemic duties. Additionally, one can plausibly think of justification binarily or non-binarily. In the binary sense, justification would be the state of a perfect fit between the two (or maybe meeting a particular benchmark of agreement), while the original formulation expresses justification understood non-binarily. Lastly, I opt for the most general account of what is subject to epistemic evaluation that also remains exclusively a matter of epistemology: ‘epistemic properties.’ I give this account in

order to avoid committing the DCEJ to controversial theses concerning the constituent elements of one's epistemic life. 'Properties,' as I use the term, can incorporate far more than merely beliefs; for instance, knowledge, understanding, virtues and perhaps even certain brain states advocated by an eliminativist philosophy of mind's acceptance of normative epistemology are all epistemic properties.⁵⁹ In the next chapter, I shall argue that beliefs merit special concern as epistemic properties, but there I present a normative ethical theory. The DCEJ, *qua* metaepistemology, need not commit itself to this specific position. I do specify 'epistemic properties,' however, to emphasize that the DCEJ is only a view on epistemic justification, rather than an 'all-things-considered' sort of justification.⁶⁰

Given its rather scant expression, it can be difficult to feel much intuitive pull either for or against deontology. Without an account of what one's epistemic duties are, there is very little of substance to the view. To illustrate this, consider a few prominent views on epistemic justification: foundationalism, coherentism and process reliabilism. The fundamental claims of each of these can be expressed as an epistemic duty (e.g. one is permitted to believe some proposition *p* if and only if the belief that *p* is the result of a reliable belief forming process). As long as justification accords essentially with the dictates of one's epistemic duty, the theory is deontological.

One common objection levied against deontological ethical theories is that they are too inflexible. Since duties tend toward absolute formulations, (e.g. thou shalt not kill) they will not admit of obviously needed exceptions in exceptional circumstances. There may be merit to this claim, but it is difficult to apply when thinking only at the meta level, since we cannot deduce

⁵⁹ Even if this last one seems, *prima facie*, like an uneasy conceptual alliance

⁶⁰ Ultimately, I will argue in chapter 7 that there is no sense to claims about all-things-considered justification, but I will avoid that point of contention here.

exactly in what ways the theory in question might prove inflexible. One could certainly contrive a set of duties to accommodate one's intuitions for every exceptional scenario; in this case, the theory would be inflexible, but unobjectionably so. Obviously, one might have other concerns about such a theory, most obviously that it is painfully *ad hoc*, but the lesson here is that inflexibility is not, in itself, a bad thing.

The remainder of the chapter is devoted to considering some of the most prominent alternatives to the DCEJ. Since the idea of meta-epistemology (as I use the term) has lacked adequate definition, there is not a very well developed literature on the topic. Here again it will be helpful to appeal to metaethics. I shall consider the epistemological versions of consequentialism, virtue theory and, most imaginatively, divine command theory. By observing that it can resolve the major problems faced by its competitors, I show that the DCEJ is the best of the notable candidates offering an account of epistemic justification.

5.2 Consequentialism

Since the goal of this section is to argue against epistemic consequentialism, it will help to give a minimal statement of the position first. Arguing against any particular position ascribing to epistemic consequentialism runs the risk of being insufficiently general. Accordingly, let us take epistemic consequentialism to be the view that the justificatory status of an agent's epistemic properties (in the same sense described in the previous section) is determined by the consequences thereof.⁶¹ Although it is a key term, it is unlikely that we can give a confident definition for what is to count as a consequence of an epistemic property, and, in some cases, that of which epistemic properties are the consequences. To do so would require a knowledge of human psychology that surpasses extant scholarship. I should like to avoid this

⁶¹ Percival and Stalnaker (2002) provide a helpful explanation of the minimal commitments of epistemic consequentialism.

problem as much as possible; my arguments below rely on plausible stipulated cases of epistemic properties serving as causes or consequences.

Without further premises, it seems there can be no non-question-begging objection to such a minimal articulation. I should therefore like to take as uncontroversial premises two claims: first, that true belief is epistemically good and that false belief is epistemically bad, and second, that it is an acceptable demand of a normative epistemic theory that it recognize believing contradictorily as bad. I shall attempt to show that epistemic consequentialism is incompatible with the conjunction of these premises; for if this is so, and the premises are true, then epistemic consequentialism is false.

Making use of the first premise with regard to epistemic consequentialism is more difficult than the analogous application regarding bad actions, moral badness and ethical consequentialism. According to ethical consequentialism, actions are the bearers of goodness and badness, and they are so according to their consequences. Crucially, actions and consequences are mutually exclusive categories. This does not hold in the case of epistemic consequentialism, since beliefs can plausibly be treated as consequences or as the originator of consequences.⁶² Thus, an adequate consideration of epistemic consequentialism is likely to prove more complicated than one concerning ethical consequentialism. Applying these considerations to the particular terms of the premise, we are left with the following result: either beliefs are consequences of other epistemic properties, such that their truth counts (*prima facie*) in favor of

⁶² This seems to be so as a result of the difference between the voluntariness of belief and action, respectively. Incompatibilism is widely accepted as a necessary condition for moral evaluation, but doxastic voluntarism is a comparatively unpopular stance, despite the prevalence of epistemic evaluations. Because beliefs are not voluntary, and actions are, the former can be treated as consequences while the latter cannot. A compatibilist view of action treats actions as consequences, and so would not generate this sharp distinction.

said property,⁶³ or that true beliefs are epistemic properties that have (prima facie) good consequences. With this established, let us consider some arguments against epistemic consequentialism.

The most prominent route of argumentation against epistemic consequentialism is that it is apparently committed to so-called ‘trade-off’ cases.⁶⁴ Consider the following: as a result of some research, I come to have a false belief, which, if true, would be very surprising. Being thus surprised, I share my belief with many of my friends. My friends, surprised by the content of my testimony, do their own, more effective research and come to believe the truth on the matter. Suppose further that if I had come to have the true belief, I would have thought it insufficiently interesting to mention in conversation, meaning my friends would never have researched the matter, and consequently, would not have come to believe the truth thereabout. The epistemic consequences of my belief are undeniably positive, so it is justified, yet there is obviously *something* wrong with my belief.⁶⁵ If epistemic consequentialism is the whole story about epistemic justification, then something of epistemic significance is left out.

First, we should note that, although the example deals with the specific consequence of true belief, our minimal interpretation of epistemic consequentialism is not committed to such a narrow axiology. I assume, however, that the case above is indicative of the viability of the following recipe: take some *ceteris paribus* bad epistemic property held by an individual, and

⁶³ Let us also make the obvious opposed remarks regarding false belief, that it is normatively bad, or that it must have some bad consequence.

⁶⁴ For similar examples to the one considered here, see Firth (1981), Littlejohn (2012) and Berker (2013).

⁶⁵ It will also help to note that the example presupposes a sort of veristic value monism, such that true beliefs are the sole factor in determining justification. Any such example will require some normative epistemological theory for its formulation, but each normative epistemological theory should suffice for generating an analogous example.

suppose that, by some quirk of circumstances, the total consequences of this property are epistemically good.

I do not think trade-off cases provide a compelling case against epistemic consequentialism. An epistemic consequentialist can still hold that there is a sense in which false belief is a bad epistemic property, yet one the badness of which can be offset by special circumstances, such as the ones described in the example. Although the result is counter-intuitive to many,⁶⁶ the same can be said of analogous cases regarding ethical consequentialism. As most will agree it is too hasty to dismiss ethical consequentialism on the results of such cases, so too should trade-off cases not be considered fatal to epistemic consequentialism.⁶⁷

I think a more compelling case can be made against epistemic consequentialism through considering instances of inconsequential beliefs. Presumably, there are some true beliefs that, if held, would not have any further consequences with regard to our epistemic properties. If there are such beliefs, then it is clear that consequentialism must treat them as valuable consequences, since their goodness cannot be explained in terms of positive consequences if they have no consequences at all. This would mean that such beliefs are not justified; rather, they justify other properties.⁶⁸

We need not get overly specific about what sort of property has true beliefs for consequences; a detailed investigation of this question would likely be more psychological than epistemological. Instead, I wish to challenge this move by proposing a dilemma. If (some) true beliefs improve the normative standing of other epistemic properties, then they must bear some

⁶⁶ See Andow (2016) for a survey of the extent of these intuitions.

⁶⁷ Alhstrom-Vij and Dunn (2014) provide a more detailed response on behalf of epistemic consequentialism to trade-off objections.

⁶⁸ Or, at least, some other epistemic property has its normative standing improved by having said beliefs as consequences. One might think of justification strictly in terms of belief, and so not accept the exact wording of this dilemma, even accepting the general point.

value; this value is either intrinsic or extrinsic. Intrinsic values cannot be evaluated comparatively, while extrinsic values can be so evaluated. It is natural to associate consequentialism with treating values as extrinsic, but given the minimal formulation above, there is no reason to think this must be the case.⁶⁹

Let us first consider the case in which true beliefs are deemed intrinsically valuable. Suppose P and Q are true propositions. Now take two epistemic properties, EP1 and EP2, such that the consequence of EP1 is the belief that P, and that the consequences of EP2 are the beliefs that Q, some other proposition R and $\sim R$.⁷⁰ If the normative status of EP1 and EP2 are determined by the respective values of their consequences, and the value stemming from the beliefs is intrinsic, then consequentialism lacks the means to deliver the verdict that EP1 is normatively better than EP2. One set of consequences has no normatively bad elements, and the other does, but since both have intrinsically valuable elements, they cannot be put into terms of each other. Hence, treating true belief as intrinsically valuable leaves consequentialism unable to affirm the platitude that, *ceteris paribus*, it is normatively bad to believe contradictorily.

Consider now the case in which true beliefs are deemed extrinsically valuable. Since this is, analogously, the closest position to utilitarianism, a popular ethical consequentialist view, this seems the stance most likely to be adopted by actual epistemic consequentialists. Consider again the case presented in the previous paragraph, only now such that EP2 results in the beliefs that P, Q and $\sim Q$. Since the natural move is to assign equal weight to avoiding error and believing the

⁶⁹ Let us again suppose the opposites of these descriptions when applied to false belief.

⁷⁰ Suppose also, for simplicity's sake, that these the consequences named are the only ones resulting from these properties.

truth, this interpretation needs to offer additional explanation if it is to deliver the result that EP1 is normatively better than EP2.⁷¹

It is important to recall at this point that, in principle, consequentialism is not committed to any sort of value monism. One could explain the preferability of EP1 to EP2 by positing a new value, aimed at accounting for the normatively bad status of believing contradictorily. Thus, the application would be that EP1 is normatively better than EP2 because EP2 is inconsistent with attaining this new value. However, on such a view, one is at a loss to explain the relation between the value associated with believing truly and that associated with avoiding contradiction. There is no principled way to weigh the normative status of sets of consequences when these values are at odds. In such a situation, epistemic consequentialism fails to hold to its definitive thesis, since the normative status of an agent's epistemic properties is not determined by the consequences of these states—the normative status is not determined at all.

While objections to a meta-epistemic theory will typically have some metaethical analog, this does not appear to be true in the present case. Believing contradictorily is absolutely epistemically bad, but no action is absolutely morally bad according to ethical consequentialists. It is impossible for a contradiction to be true, but, at least in terms of bare logical possibility, we cannot say that any action will necessarily have morally bad consequences. As such, the objection presented above should not trouble ethical consequentialists.⁷² The epistemic consequentialist might try to adopt the same response, and claim that believing contradictorily has no normative status in its own right. This is precisely the claim I think epistemic

⁷¹ Since beliefs are not closed by deduction, one could believe Q and ~Q without also believing (Q & ~Q). Thus, we cannot appeal to other epistemic risks associated with believing contradictions in trying to settle the matter of the value of this set of consequences. The example should also apply *mutatis mutandis*, to any different schema for valuing true beliefs and disvaluing false beliefs, e.g., if, with James, one wanted to place more value on believing the truth than on avoiding falsity.

⁷² At least not *qua* ethical consequentialist; if he or she were also an epistemic consequentialist, then he or she ought to be troubled.

consequentialists should make, but I think her or his commitment to this position is the best indicator we could wish for that epistemic consequentialism is simply the wrong approach to meta-epistemology.

Finally, let us note that a deontological theory does not face the same problems in dealing with such cases. One might simply determine that it is part of one's epistemic duty not to believe contradictorily, so any failures of this duty are epistemically unjustified. If we also have a duty to believe the truth, then the justification ensured by adhering to this norm can be explained in terms of the same value, namely, the value of adhering to one's epistemic duty.

5.3 Virtue Epistemology

While it is squarely a third-party candidate in the metaethical sphere, virtue theory may enjoy the most widespread acceptance in epistemology.⁷³ This is due in part to the fact that, as discussed in the introduction, meta-epistemology is a largely unexplored area, so there are few who have endorsed opposing meta-epistemological views. Put differently, since virtue epistemology is one of the few schools that has consistently asked the relevant meta-epistemological questions, it is natural that it should present one of the more accepted answers to those questions.

One is faced with a certain conceptual difficulty in comparing virtue epistemology and the DCEJ. According to the latter, beliefs are typically the subject of epistemic evaluation, so when one talks of epistemic justification, the property is ascribed to or withheld from beliefs. This is not so according to virtue theory; this view posits that epistemic agents are the proper object of epistemic evaluation. As such, it is more accurate to call virtue epistemology a theory about epistemic goodness than one about epistemic justification. One might claim on these

⁷³ The two metaethical theories that enjoy greater acceptance are, of course, deontology and consequentialism.

grounds that justification is a property of agents, but it seems a more natural move to dispense with the term ‘justification’ altogether; agents are virtuous or vicious, and that is what is important from an epistemic perspective. Given this fundamental difference, comparing the results of the DCEJ and virtue theory will be more difficult than was the case in comparing the former to epistemic consequentialism. My argument against virtue theory will thus need to be aimed at a more fundamental aspect thereof, rather than in merely pointing out that it delivers the ‘wrong’ result regarding epistemic justification.

Let us take virtue epistemology to be the view that epistemic goodness is a function of epistemic virtues and vices such that epistemic goodness is essentially a matter of exemplifying said virtues and epistemic badness is essentially a matter of exemplifying said vices. The view’s normative decrees will therefore consist in a demand for virtuousness.⁷⁴ The normative level of the theory will be concerned with defining what the various epistemic virtues and vices are, and perhaps how to cultivate the former and to eradicate the latter.

My challenge to virtue ethics begins with the following dilemma: either epistemic virtues are intrinsically valuable or they are extrinsically valuable. The two possible answers correspond to one of the main divisions within virtue epistemology: responsibilism and reliabilism, respectively.⁷⁵ I shall explore the second horn first.

I shall take Ernest Sosa as my exemplary virtue reliabilist. A brief recapitulation of some of Sosa’s more basic views should suffice to show why I anticipate his adopting the second horn of my dilemma. According to Sosa, apt, true belief suffices for knowledge.⁷⁶ (2007, 40) A belief

⁷⁴ For simplicity’s sake, I shall speak principally of virtue from here on, though most mentions of virtue should be understood to carry an additional, opposite claim concerning vice, e.g., virtue is epistemically desirable (and vice is epistemically undesirable).

⁷⁵ Virtue reliabilism is not to be confused with process reliabilism, another notable position in epistemology.

⁷⁶ Specifically, apt belief suffices for *animal* knowledge; this is contrasted with reflective knowledge, which is roughly knowing that one knows.

is apt if and only if the belief results from a manifestation of a competence of the believer. (29)

The competences that result in apt beliefs are epistemic virtues. Since epistemic virtues are defined in terms of their conduciveness to knowledge, it seems the value of the former must be derived from the latter. Knowledge is not virtue, so virtue is extrinsically valuable.

My objection to this view is that it is not, properly speaking, virtue epistemology. While this may sound like a rather drastic objection, one can clarify the grounds for this objection by appealing once more to level distinctions in epistemology. An example should suffice to illustrate my point. Consider a deontological epistemological theory according to which it is obligatory that one know that one has hands. Suppose further than Sosa's analysis of knowledge is correct. It follows that one is obligated to exemplify epistemic virtues, since without the virtue, we lack the apt belief necessary for the obligatory knowledge. Nonetheless, this is not a virtue epistemology in the sense with which we are concerned in the present section, since the theory is, by hypothesis, deontological. The problem is that virtue only becomes essential to the theory at the normative level. One is justified (or attains epistemic goodness) in virtue of fulfilling one's epistemic duties; that these duties involve virtues matters not at all to the question of whether the view is deontological. So, despite any plausibility that Sosa's views on the value of virtue might have, the sort of value, namely extrinsic value, that is attributed to virtue is insufficient for the central claim of virtue (meta-) epistemology. The point here is subtle; on the case considered, an exhaustive account of epistemic goodness will include the concept of virtue, since virtue is necessary for knowledge, and knowledge suffices for epistemic goodness, but to answer the

question “what constitutes epistemic goodness?” we need not mention virtue. As such, this is not, or at least not at the meta-epistemological level, a virtue theory.⁷⁷

Let us now consider the answer corresponding to the other horn of the dilemma: the claim that virtue is of intrinsic epistemic value. As noted, this claim will correspond with what has been dubbed the responsibilist camp within virtue epistemology. Obviously, responsibilism will not succumb to the same difficulty that faced reliabilism. If we think of epistemic virtues as excellences of epistemic agents, and we affirm that agents are the proper object of epistemic evaluation, then virtues are importantly involved, even at the meta level.⁷⁸ It seems, therefore, that responsibilism is the ‘true’ virtue theory, at least in limiting our concerns to meta-epistemology.⁷⁹

It will not do to contest the claim that epistemic virtues are excellences of agents; this much seems true by definition. It follows that, in order to object to responsibilism, I must take aim at the second claim, that agents are the proper object of epistemic evaluation. Unfortunately, there is not an obvious methodology for such a strategy.⁸⁰ Following my approach in objecting to epistemic consequentialism, I shall therefore present what I take to be some uncontroversial observations about the nature of epistemology in general and argue that the DCEJ is better equipped to accommodate the truth of these assumptions.

⁷⁷ We might clear things up by specifying at which level a theory can be said to involve the concept of virtue. Having done this, there is some sense in which Sosa’s view is a virtue epistemology, but my interest in this chapter is only with virtue theory at the meta level.

⁷⁸ In other words, one cannot, as was the case with reliabilism, affirm that epistemic goodness is merely a matter of fulfilling one’s epistemic duty.

⁷⁹ To clarify, the fact that responsibilism and reliabilism operate at different levels is not an indication of an advantage that one might hold over the other. If anything, the reasoning of this section suggests that the two are not really opposing views. One might adopt both horns of the dilemma I posed, and endorse virtue theory at the meta and normative level.

⁸⁰ I argue in favor of the viability of treating beliefs as one such object, but not to the extent that we should exclude other objects from possible evaluation.

I shall begin with a few points from the virtue ethics of Aristotle. According to Aristotle, all (human) actions aim, either directly or indirectly, at a flourishing human life. (NE, 1097b20) The flourishing life is the life that consists in performing the human function well. (NE, 1098a15) Having virtues is constitutive of being a functionally good human. Thus, our concern with actions is essentially a concern with virtues: actions succeed in their ultimate aim only when they are performed by one with a virtuous character.

To consider the case of virtue epistemology, we need to identify the analogous goal of the theoretical life.⁸¹ Following the analogy with ethics, the goal of the theoretical life will be that toward which all beliefs aim—the terminus to questions about why one believes what one believes. Any satisfying answer to these questions must involve, at least in part, some reference to the truth.⁸² Just as one acts for the sake of the good life, one believes for the sake of the truth. Thus, the goal of the theoretical life is to believe the truth, or at least those truths appropriate to human consideration.⁸³

It is logically possible that one could hold any given belief without exemplifying any epistemic virtues. Young children are presumably capable of beliefs, but also plausibly lack the degree of development needed to exemplify epistemic virtues or vices. Epistemically vicious agents are capable of representing complex propositions to themselves, and so capable of believing truths without virtues. So, no particular belief requires virtue for its success. However, since neither deontologists nor virtue theorists limit their assessments to a single belief, we should consider a more expanded case as well.

⁸¹ I use ‘theoretical’ since ‘theoretical reason’ is used more frequently than ‘epistemic reason,’ the latter sounding redundant. However, when I talk about the theoretical life, I am referring to the realm to which epistemic normativity applies.

⁸² Or, to put it in Aristotelian terms, the function of theoretical reason is to represent the truth.

⁸³ Presumably, there are some truths that are not relevant to the human condition, so it seems that it may not be for the sake of *those* truths that we believe what we do.

There is good reason to think that a moderate degree of skepticism is indicative of epistemic virtue, perhaps of open-mindedness. We might also think that an enthusiasm for truth is an epistemic virtue. Neither of these is of use to a being who believes all truths. Even if such a being were thoroughly dogmatic, it would detract nothing from her or his success in believing, and therefore nothing from his or her epistemic goodness. The significance of this case can be explained in terms of the distinction between responsibilists and reliabilists. It might be the case that the hypothesized being possesses cognitive abilities beyond the realm of what is possible for humans, and so achieves a goal that humans could not in the absence of virtue.⁸⁴ However, this response does more to motivate reliabilism than responsibilism. Just as was the case in considering Sosa's analysis of knowledge, when considering the state prescribed by the aim of our theoretical reason, we do not need to mention virtue. It may be a psychological fact of humans that virtue is needed to realize this state, but all that fact suffices to establish is that virtue is valuable as a means to an end. If believing the truth is the ultimate goal associated with the capacity to believe, then the realization of this good will not be solely a matter of instantiating epistemic virtues.

A deontological view does not face this problem. If epistemic goodness is a matter of believing the truth, then a theory that makes believing thusly obligatory will suffice for the highest epistemic good. Departures from the best doxastic profile can be assessed in terms of particular doxastic attitudes (and not necessarily in terms of character traits). It follows, therefore, that we are better served using beliefs, or epistemic states more generally, as the object of epistemic evaluation.

⁸⁴ For instance, such a being might have reached a state that no human could reach without a strong love of truth.

Obviously, my assumption about the goal of our theoretical lives is unlikely to sit well with the responsibilist. The likely objection is that this is an overly narrow conception of epistemic goodness. Significantly valuable epistemic ends like understanding, wisdom and know-how are left out entirely. Thus, even a being that believes the truth, if it lacks these other states, could live an impoverished intellectual existence. If epistemic virtues are necessary to any of these,⁸⁵ then we can still claim that agents are the proper object of normative epistemological evaluations and, consequently, that virtues are essential to an accurate meta-epistemological theory.

Beyond this point, a decisive advantage for either side seems unlikely. The claim I endorse, that knowledge and, derivatively, true belief are the proper concerns of epistemic evaluation has the backing of the majority of work in the discipline. Of course, this counts for little; I have already disputed claims supported by much of contemporary epistemology, so it will not do merely to appeal to the (philosophical) crowd. This illustrates a general difficulty in resolving disputes at the meta-normative level. Since one makes claims about fundamental value in this sphere, one has little recourse in trying to win over opponents since the claims for which one advocates will not be seen as valuable by hypothesis.⁸⁶

However, there are some good prospects for a pluralistic solution if both sides agree that they are simply interested in different questions. The deontologist might be interested in normative epistemology in a narrow sense, while the virtue theorist might be interested in how

⁸⁵ One might challenge this antecedent from the deontological perspective, arguing that a duty-based approach can accommodate all of these values, but I am not inclined to pursue this strategy here, as I think the prospects for such an attempt are poor.

⁸⁶ In other words, if I think there is no epistemic value in understanding something that is not possessed by knowing that same thing, there seems to be no non-question begging way to argue to the contrary.

one's intellectual life relates to the good life as a whole. This distinction made, both sides can recognize the methodology of the other as appropriate to the questions they ask.

One can push this happy pluralism back to a tense disagreement if one considers which set of questions is more important: the sphere of epistemic normativity, narrowly construed, or that of the intellectual life as part of the complete, flourishing human life. I do not see any way to resolve this disagreement with any finality.⁸⁷ Fortunately, I do not think the success of my project depends on answering in favor of the narrow construal of epistemic normativity. Even if epistemic justification turns out to be something less than the most important topic, there are still right answers to less-than-maximally-important questions. While all philosophers might think they are concerned with the most important questions, not all of them can be right about this. That the concept of epistemic justification has such a significant literature is a positive indication that it is at least worth continuing to think about.

To the dilemma regarding the type of epistemic value had by virtue, I identified the responsibilist answer, that it is intrinsic, as the true meta-epistemological virtue theory. I challenged the intrinsic worth of virtue by presenting a case where a being is maximally epistemically good, but lacks epistemic virtues and, less extremely, particular instances of beings that have acquired some degree of epistemic goodness without any epistemic virtues. The ensuing discussion yielded the following lesson: deontology looks to be the superior view if epistemology is essentially concerned with the value of true belief while virtue epistemology enjoys the advantage if epistemology is essentially concerned with how epistemological

⁸⁷ If virtue epistemologists are content to categorize their criteria of evaluation as pertaining to a holistic value, rather than a narrow epistemic value, with which the DCEJ is concerned, then the problem is resolved, at least in the immediate sense. The potential conflict of these two species of evaluation (all-things-considered and narrowly epistemic) is one of the subjects broached by the seventh chapter.

normativity relates to the good life. Both pursuits are valuable, if not central, so I shall continue with the deontological view with a clarified conception of the scope of epistemic normativity.

5.4 Proper Functionalism and Divine Command Theory

There remains a prominent view on epistemic justification that does not fit neatly into any of the meta-epistemological views we have described thus far: proper functionalism. The view was first articulated by Alvin Plantinga as an alternative to the concept of epistemic justification. Since the view is externalist, and the term ‘justified’ brings with it such a strong internalist, and, particularly, deontological, connotation, Plantinga urges the adoption of the more neutral term ‘positive epistemic status.’ (1988: 3) Despite his suggestion, the move has not really succeeded; ‘justification’ is still by far the more widely used term, among both internalists and externalists. If I have argued as I intended, then my case for the DCEJ has not taken advantage of this connotation. Thus, I mean the same thing by ‘justification’ and the related terms as Plantinga means by ‘positive epistemic status.’

Plantinga’s tentative analysis of positive epistemic status is as follows: “a belief B has positive epistemic status for S if and only if that belief is produced in S by his epistemic faculties working properly...”⁸⁸ (34) In terms of the level divisions developed in this chapter, there are two ways of assessing proper functionalism. In the first place, we might think of the view as making an assertion at the meta level. On such a view, epistemic justification just is a property of beliefs that are the result of the proper function of one’s epistemic faculties. This interpretation is saddled with the immediate difficulty of explaining what distinguishes proper from improper function. The natural move would be to give some non-normative explanation of what constitutes proper function. However, if the initial analysis were understood to be a meta-

⁸⁸ Plantinga continues the analysis to address cases of having more than one belief that results from the proper functioning of one’s epistemic faculties, but we need not concern ourselves with that degree of complexity here.

epistemological claim, then we would be left without an explanation of why we should value one particular non-normative means of function to another.

I assert that the more natural interpretation is to see Plantinga's claim as a normative one. If we can give a non-normative explanation of what constitutes proper function, then we can bolster the analysis with a meta-epistemological theory to yield a complete account of justification. This interpretation also fits better with the facts of Plantinga's argument, since he does indeed give a non-normative account of what constitutes proper function. Proper function is so called because it is supposed to accord with the design plan of our epistemic faculties. (36-37)

Curiously, none of the meta-epistemological theories we have described to this point can accommodate proper functionalism particularly well. The explanation open to a deontological conception seems impoverished. The account would need to say that believing in accordance with the design plan of one's epistemic faculties is one's epistemic duty, and therefore constitutive of epistemic duty. Still, if this is a full account of justification, it remains to be explained why the design plan itself is epistemically significant.

Consequentialism also delivers poor results when combined with proper functionalism. For whatever epistemic value the promotion of which one deems essential to justification, it seems we will be able to construct counter-examples in which said value is maximized by deviating from our proper function.⁸⁹ Plantinga's own counter-examples to reliabilism would suffice for showing the combination under consideration to be problematic.⁹⁰ Of course, if one were to say that beliefs that result from the proper function of one's epistemic faculties are the

⁸⁹ Knowledge as the epistemic value to be promoted may not be susceptible to this problem (i.e. there will be no case where knowledge is maximized by deviating from our epistemic design plan), but only if we include proper function in an analysis of knowledge. But if we include knowledge in our analysis of justification, then we cannot explain knowledge in terms of justification, which is at least inconsistent with Plantinga's aim in the paper under consideration.

⁹⁰ All of these examples involve a reliable mechanism giving rise to a true belief through a malfunction of the normal (proper) function of the believer's cognitive processes. (Plantinga, 1988 21, 23, 31)

consequences the maximization of which suffice for epistemic justification then the problem would not emerge, but this is a painfully ad hoc determination, and offers little explanatory value.

Of the three mentioned, virtue theory is probably the best match with proper functionalism, but it is still not an ideal union. Since epistemic virtues are just excellences of epistemic agents, then there is a sense in which being virtuous will consist in adhering to the design plan of one's epistemic faculties. However, in the case of virtue theory, the source of epistemic good seems to be self-contained. There is no further explanation offered for why one should want to be an excellent epistemic agent; indeed, the question fails to make much sense. Since proper functionalism is committed to the notion of a design plan, then a natural move in explaining the epistemic worth to adhering thereto would be to reflect on the properties of the designer. The strategy of explaining epistemic value in terms of the properties of an agent's designer leads us to consider a final position in meta-epistemology.

In the strictest sense, there is no designer of our epistemic faculties from the perspective of Darwinian evolution. Humans have the epistemic faculties that they do because they inherited them, and they inherited them because the same epistemic faculties have proven to be reproductively advantageous. Plantinga is non-committal over this claim, but maintains that, at the very least, we can use a sort of fictionalist description of proper function without committing ourselves to theism. (45-46) I think that this claim is mistaken. From a purely naturalistic perspective, there is little reason to think that our epistemic faculties will result in true beliefs when functioning properly. What matters with regard to reproductive fitness is not truth, but a tendency to survival and reproduction. Setting aside cases where empirical research suggests that the two can actually come apart, there are easily conceivable cases that show proper function is

not, by itself, sufficient for any epistemic value. Suppose life is meaningless, and there is no good reason not to kill ourselves.⁹¹ It would be terribly uncondusive to the perpetuation of a species' genetic information to realize this truth, so we have reason to think that there are no species the proper function of whose epistemic faculties would yield such a belief. This example shows that, absent a designer, proper function is at least insufficient for any epistemic value.

Let us ignore the question of whether the fictionalist route is a viable one; it appears that proper functionalism cannot do without appealing to the notion of a designer. Of course, not just any designer will do. A designer's plan is only as valuable as the designer is qualified; therefore, if the design plan is to serve as the basis for all epistemic justification, an omniscient God is the obvious choice for a conceptual appeal to a designer. Once this appeal is made, I think it becomes clear that the best meta-epistemological option available to the proper functionalist is the epistemic analog of divine command theory. The essential claim of this position would be that beliefs that result from properly functioning faculties are justified simply because that is how God has deemed we ought epistemically to function.

The combination of these views approaches a complete theory. Beliefs are justified when they are the result of the proper function of the believer's epistemic faculties. This proper function is normatively important because justification is merely a matter of adhering to God's decrees. God has decreed that we ought so to function, and we ought to respect God's decrees, even on purely epistemic grounds, since God is omniscient.⁹² Indeed, the epistemic version of divine command theory handles better the analogous version of the Euthyphro dilemma.⁹³ It is

⁹¹ I take it that this is at least a live option, even if an unpleasant one.

⁹² There lingers the prudential threat of eternal damnation should we break God's decrees, but those considerations are not obviously related to *epistemic* justification.

⁹³ Put in terms of the view under consideration, the dilemma would be as follows: "Is the true epistemically valuable because God's design plan has deemed it so, or has God's design plan deemed it so because the true is epistemically valuable."

not obviously contrary to the core of theism to say that there is independent epistemic value had by truth.⁹⁴

Despite the promising prospects for this view, I think it suffers from an ineliminable problem. God is defined as omniscient, so we have introduced the concept of knowledge into our analysis for justification. Since Plantinga is interested in using proper function as an element of warrant (the sort of thing that, if added to true belief, suffices for knowledge),⁹⁵ the proper functionalist view is committed to explaining knowledge circularly. Crucially, omniscience cannot be removed from the predicates assigned to the designer, since it is only by virtue of this predicate that adhering to the design plan can be plausibly said to be of epistemic value.⁹⁶

5.5 Applied Epistemology and Truth

Despite having made a division between three levels of epistemological inquiry, to this point I have only discussed two of these levels: meta and normative epistemology. Referring once again to the analogy, applied epistemology seems to be the discipline concerned with determining what epistemic states one should hold in response to particular situations. Despite the immense popularity of applied ethics, the analogous epistemic enterprise has not been widely practiced. Work in applied ethics typically proceeds by applying some moral theory to a particular action or practice (e.g. is killing in self-defense morally justified?). Rarely do philosophers write about some particular epistemically controversial belief, arguing that it ought or ought not to be accepted in virtue of its resulting from a reliable process, cohering with large sets of other beliefs, or stemming from an epistemically virtuous character. Doubtless, examples

⁹⁴ This might turn on distinguishing between ‘Truth’ and ‘truth’, with the former being more strongly associated with God, but there is on all accounts at least some epistemic value to be found in the latter.

⁹⁵ This is suggested most explicitly by the title of his later work *Warrant and Proper Function* (1993).

⁹⁶ I take it that this also suffices to discount any appeals to mystery associated with the divine. If God’s knowledge is not the sort of knowledge humans can have, then it is at best mysterious how the former sort of knowledge could have any significance in determining the latter sort of knowledge.

are given in the literature in epistemology, but these works tend to start with the theory and arrive at the examples, suggesting they belong to the normative epistemic sphere. There are two apparently viable explanations for this phenomenon.

On the one hand, one might suggest that there is little point to beginning the project of applied epistemology without answering the questions invoked by meta and normative epistemology. Since applied epistemology will consist in applying the true normative epistemological principles, we need first to know what principles we are to apply before we can apply them. Of course, there is nothing close to a consensus on the nature of epistemic justification, so, this line of reasoning concludes, there is nothing close to the foundations required for an applied epistemology. Still, similar concerns apply in the case of ethical justification, but this has not stopped applied ethics from becoming a thriving discipline. The applied sphere can cope with the incompleteness of the higher levels in two ways. First, it can assume that a particular normative theory is true, and attempt to determine what would be true of applied ethics if this assumption were correct. Second, it might apply principles that are held commonly by all normative theories (e.g. torturing for fun is morally wrong). If such options were available, it would seem odd if there were truly no applied epistemologists; after all, philosophers tend to philosophize about whatever they can. The other explanation, alluded to in the previous paragraph, explains this phenomenon better.

Applied epistemology is the discipline concerned with finding out what to believe. There are many people engaged in this process, even if they do not adopt the label. There is a sense in which all philosophy, and all scholarship in general, is an attempt to determine what we ought to believe. Working within this widened view of applied epistemology, we can note that truth seems to be the central concern; other values, like clarity, predictiveness and simplicity are all

theoretically valuable only to the extent that they are consistent with (or constitutive of) truth. It follows that, whatever is said at the level of meta and normative epistemology, we ought to expect that justification is importantly connected to truth. In the next chapter, I shall argue that truth is a necessary condition for a belief's being justified. This obligation would accommodate the lesson we can learn from applied epistemology well, but it should be observed that most views on justification have some means of accommodating the centrality of truth.

Prominent externalist views on justification are unified in their requirement of some non-accidental connection between beliefs and truth. Internalist views, while not explicitly connected to truth, would assure a strong connection between truth and belief if one were willing to accept a few basic anti-skeptical assumptions. If the observation about applied epistemology affords any advantage to a strong truth norm of belief, it is in its simplicity. In considering what to believe, we typically consider only if the belief is true, not whether it is the result of a reliable process, etc. Therefore, we must consider at greater length the viability of and motivations for a truth norm of belief.

CHAPTER 6

THE TRUTH NORM OF BELIEF

The preceding chapters were devoted to a defense of the deontological conception of epistemic justification. If I have succeeded there, then I have shown that the DCEJ is the correct view concerning justification. However, as the previous chapter explained, the DCEJ is best understood as a metaepistemological view, which means that it offers only a partial explanation of the abbreviated concept. The DCEJ states that justification is essentially a matter of duty fulfillment, but it is not committed to any particular view concerning what those duties are. In order for the DCEJ to be the correct metaepistemological view, it must also be able to accommodate the correct normative epistemology. The goal of the present chapter is to advocate for the prospects of the DCEJ to accomplish that goal, at least in part. I argue that we have an epistemic duty to believe the truth, and that, consequently, adherence to this duty is a necessary condition for epistemic justification.

It should be noted, however, that the conclusions of the first part of the project and the present chapter are independent of each other; one could accept the DCEJ and reject a truth norm of belief, and one could accept a truth norm of belief while rejecting the DCEJ. I am more strongly committed to accepting the DCEJ, but the central idea to this chapter is that the two views work well in conjunction with each other, and that the truth norm has some degree of independent plausibility.

Before beginning a defense of a truth norm, I must give some preliminary arguments. In the previous chapter, I noted possible disagreements as to the loci of epistemic evaluation— the

bearers of epistemic goodness and badness. One might reasonably say that agents, beliefs, sets of beliefs, or some other epistemic properties are the loci of epistemic evaluations. At the very least, it is clear that if agents, and not the beliefs of those agents, are to be evaluated, then it will be more difficult to defend the position that we are obliged by a simple truth norm. Therefore, I must consider which views can accommodate the thesis for which I argue in this chapter, and show that these views are independently acceptable.

6.1 Objects of Epistemic Evaluation

To begin, note that beliefs are the basic elements of one's epistemic agency. Even if, in keeping with the majority of traditional epistemology, we are primarily interested in knowledge, since there can be beliefs that are not knowledge, but no instances of knowledge that are not believed, beliefs are more basic than knowledge. Similar considerations apply in cases of understanding. Beliefs are more basic than epistemic character traits because it is through the traits' connection to beliefs that they are epistemic at all. Consider the difference between a practical virtue, say, bravery, and an epistemic virtue like open-mindedness. The latter is epistemic because it pertains, essentially, to one's attitudes concerning truth and falsity; to be open-minded is to countenance the possibility of error in one's standing doxastic attitudes.⁹⁷ However, it is beliefs that admit of the possibility of error. As such, one is capable of having epistemic character traits only if one is capable of believing, so beliefs are more basic than epistemic character traits. Beliefs are more basic than the processes that form them, since the latter are defined by reference to the former. A more general, and perhaps more controversial way of putting the point is as follows: epistemology is essentially (if not exclusively) concerned

⁹⁷ This is likely an overly simplified definition of the virtue, but I think any definition will include some reference to truth or falsity.

with truth and falsity, and it is only through beliefs that we are able to relate our cognitive lives thereto.

But the mere fact that one thing is more basic than another does not mean that the former is a more proper object of evaluation. A painting may be composed of brush strokes, but that does not mean that we should evaluate brush strokes and not paintings.⁹⁸ We require additional motivation to think that beliefs have any advanced standing as the proper objects of epistemic evaluation. My next point is to observe that having beliefs is necessary and sufficient for being the subject of epistemic evaluations. The previous paragraph should suffice to show that it is necessary; since beliefs are the most basic elements of epistemic agency, there will be no agent with epistemically evaluable properties that lacks beliefs. An example should suffice to show that it is also sufficient. Consider some epistemic agent about which the only information we have is that it holds a true belief regarding some proposition P. Compare that agent to another agent about which we have no information other than that it falsely believes the negation of P. The agent with the true belief should receive the more favorable epistemic evaluation. For almost any view of epistemic justification, this initial evaluation can be overthrown in light of further information concerning the two agents, but we are still in a position to make the provisional evaluation, so having beliefs is sufficient for epistemic evaluation. One's reasons for making the determination in question may vary: true beliefs are more likely to be the output of reliable belief-forming processes; they are more likely to cohere with other truths; true beliefs are more readily consistent with virtuous characters, etc. Only the most extreme internalists would deny the legitimacy of even so provisional an evaluation, since, in not knowing anything about the

⁹⁸ One might object that in evaluating the composite, we evaluate the more basic elements, but if our interest only emerges at a certain level of complexity, then the point is lost. For instance, it suffices to defend the view that diverse epistemic profiles are the proper object of epistemic evaluation to observe that we only care about the success or failures of beliefs when one has them in sufficiently high quantities.

agents' other internal states, we are not in a position to make any determination about their respective justificatory statuses.⁹⁹

If merely having beliefs is both necessary and sufficient for being the object of epistemic evaluations, then the beliefs themselves are a viable object for epistemic evaluations. If we pick any other property as the subject of these evaluations, then there could be cases in which an epistemic agent has beliefs, and is therefore eligible for evaluation, but cannot be evaluated, as it could lack the proposed alternative property. Since beliefs are also necessary for epistemic evaluations, no such problematic cases will emerge if we adopt beliefs as the object of epistemic evaluation. To contrast this point with the example of evaluating a painting in terms of its brush strokes, we can see that the end of a brush stroke can only be understood with reference to the whole of which it is a part, so brush strokes are not sufficient for evaluating a painting. Since beliefs have an intrinsic goal,¹⁰⁰ we have a means for evaluating them in their own right. Since we risk incomplete evaluation if we depart from beliefs to considering more complex epistemic properties, we should consider the prospects of developing a normative system in terms of the intrinsic goal of belief.

This argument suffices to show the legitimacy, and perhaps the fundamentality of a theory of epistemic justification centered on beliefs, but not necessarily that any other normative system, one that takes aim at a more complex epistemic property, is illegitimate. There is more to our epistemic agency than beliefs, and we might have more complex or ambitious goals than the intrinsic goal of truth held by beliefs. If these goals are sufficient to generate epistemic obligations, then we may have more duties than the truth norm, but, if the argument of this

⁹⁹ I address such a position when I discuss foundationalism in section 6.5.1

¹⁰⁰ That they have the goal they do, the aim at truth, is an intrinsic fact, even if the realization or failure to realize that goal needs to be explained in terms of another concept, i.e., correspondence with reality.

chapter succeeds, any other duties we have must be at least consistent with our obligation to believe the truth.

Having established the thesis of this section, we must consider how it supports the general thesis of this chapter: that we have an epistemic duty to believe the truth. To do this, I must return to a promise made in the introduction. I have already made heavy use of my premise concerning the direction of fit of beliefs, and the sense in which true beliefs can be said to have succeeded. However, this direction of fit does not provide the grounds for an epistemic norm immediately. Actions succeeded in the same sense when they realize the intention that motivated them, but few, if any ethical theories ground their norms in this fact. Asserting a truth norm of belief would amount to making the analogous claim in epistemology; therefore, I must give a positive argument as to why the direction of fit is given more weight in the case of epistemology.

6.2 True Belief and Normative Epistemology

I have two arguments, one negative and one positive in strategy, in support of the claim given at the end of the previous section. I shall begin with the more positive argument.

The positive argument appeals to doxastic involuntarism, and the voluntariness of action. Actions succeed just in case they bring about the intended result, but we can also evaluate the intentions of the agent in the case of actions. The same action might be afforded a different normative status, depending on the intentions that motivate it. This cannot hold with regard to beliefs, since there is only one reason one can have for believing: that the proposition in question is true, or at least seems so.¹⁰¹ Since there is only one possible motivation for belief, they cannot be evaluated in terms of this motivation, lest all evaluation deliver the same result, rendering the

¹⁰¹ One might describe evidence as being a ‘reason’ for belief, but in a sense clearly different from the reason mentioned in this sentence. We might describe evidence as a reason for something’s seeming true.

evaluation rather pointless.¹⁰² All that remains, if we adhere to what we have identified as the basic target of epistemic evaluation, is to consider whether they satisfy the end aimed at by the motivation. Thus, if there is to be any direct normative evaluation of beliefs, truth is the only applicable standard therefor.¹⁰³

Recall that, rather than being attributable to both beliefs and actions, voluntariness is a property of actions alone. Trivially, therefore, we hold a degree of control over reality in acting that is absent in believing. We have stated that actions succeed when they realize their motivating intention, and that beliefs succeed when their propositional content accurately reflects the state of affairs. However, there is a range of possible motivating intentions that might bring about an action, whereas there is only one state of affairs that is to be represented by our beliefs. Since, in the case of actions, we have a range of options that we lack in the case of beliefs, we can be stricter in formulating norms of action.¹⁰⁴ There is only one type of belief (namely, true ones) that can succeed, so there is (at least) one norm that can always be used in assessing them— the norm of truth.

The second argument of this section appeals to the success of the previous section. If we take beliefs as objects of normative evaluation, there is no plausible candidate for a normatively valuable epistemic property that they can have other than truth. Relations between beliefs are not beliefs, so the relational properties of beliefs are not the proper object of epistemic evaluation. It

¹⁰² It is worth noting that, if one accepts some variety of psychological egoism, the same might be true of ethical evaluations. If our intentional structures are such that, no matter what actions we take, we take them for the sole reason that we anticipate doing so will yield the most utility to us out of the options of which we are aware, then it would make no sense to evaluate actions in terms of motivations, since, from a formal standpoint, the motivation for all actions would be identical.

¹⁰³ I specify ‘direct’ because beliefs can be evaluated in terms of a wide range of alternative properties, as the history of epistemology has demonstrated. But the aim at the represented proposition’s truth is the only fact analytic to the belief itself.

¹⁰⁴ Again, I assume that some sort of principle of alternative possibilities must be true if there is to be any sense to normative evaluation of actions.

follows, at least when evaluating beliefs, most of the traditional positions in epistemology are of no help to us.¹⁰⁵ Internalists tend to equate justification with some fit between belief and evidence, or accessible mental state. Foundationalists view justification as a matter of a belief having the right inferential or causal support. All but the most extreme externalists think of justification as being a matter of holding some non-accidental, non-identical connection with the truth.¹⁰⁶ Coherentists are straightforwardly concerned with relations between beliefs, rather than directly with the beliefs themselves.

Since beliefs, qua intentional state, aim essentially at the truth of their represented proposition, we do not need to look beyond the belief itself to find a standard by which to evaluate it. Once we have identified beliefs as proper objects of normative evaluation, the truth norm is the only one to which we can appeal without incorporating more complex concepts, which would remove us from what we have identified as the basic locus. To *apply* the standard we would doubtless need to look beyond the belief itself, namely, to the state of affairs corresponding to the represented proposition; however, for any other norm, we would need to look beyond the belief both to formulate the norm, and to apply it. Because of their peculiar nature as intentional states, the standard of truth is uniquely essential to beliefs.

The last preliminary step is to consider how the DCEJ can accommodate the principles argued for above. There is at least a suggestion of inconsistency in that I have already described the core proposition of the DCEJ as saying that one is epistemically justified to the extent that one respects one's epistemic duty. However, this indicates that the individual or at least the individual's collective epistemic properties relevant to his or her epistemic duties are that which

¹⁰⁵ At least immediately, though I argue in section 6.5 that these views can be accommodated as derivative norms.

¹⁰⁶ The most extreme form of externalism would make truth a necessary condition for justification, which is the position I endorse by adopting a truth norm of belief.

is evaluated. One can put deontology into terms of beliefs as follows: one's epistemic duty is to have true beliefs; therefore, only true beliefs are justified and all false beliefs are unjustified. In the absence of any other epistemic duties, agents are justified to the extent that their beliefs are true and unjustified to the extent that their beliefs are false.

6.3 Formulating the Truth Norm

Even if my argument for a truth norm is wholly accepted, it remains to be established exactly what the norm demands. A natural formulation of the norm as motivated by the arguments of the previous section would be something like the following: "for any subject S, and any proposition P, if S forms a belief about P, then it is S's epistemic duty to believe P iff P is true and to believe $\sim$ P iff P is false." This handles instances of error well; we have the normative mechanism in place to give the correct, reprimanding judgment in cases of mistaken belief. However, there are other epistemological failings that this formulation seems incapable of diagnosing. Suppose I am deeply troubled by evidence that suggests the deterioration of the Earth's ability to support human life. I am not mistaken about the relevant facts, I simply do not like thinking about it, so I make a conscious effort to avoid information on the subject. As a result, I do not form any beliefs on the matter, and I meet my epistemic obligation, at least as expressed by the formulation under consideration. Let us call such cases instances of culpable ignorance. There is an obvious sense in which I merely am exploiting a failure of wording in such a case (rather than legitimately meeting my epistemic duty); I show a clear and conscious disregard for the truth, which has already been established as epistemically valuable. Furthermore, any attempt at formulating the truth norm that makes evaluation conditional on one's having a belief about a proposition will be susceptible to instances of culpable ignorance. It

seems, therefore, that we need a stronger iteration of the truth norm, one that will not allow anyone willfully to excuse himself or herself from evaluation.

Strengthening the truth norm generates difficulties with the contrapositive of the ‘ought implies can’ rule. Consider the simplest and strongest version of the truth norm: “One ought to believe that p if and only if p .” In the first place, it seems that there are many propositions that, while true, lack any sort of evidential support. Obscure and insignificant facts about the past provide accessible examples. Socrates had a certain number of hairs on his head at the time of his death, but this detail was not recorded; therefore, even if I fulfill the norm with regard to the corresponding proposition, I would only be able to do so by virtue of a lucky guess.¹⁰⁷ Secondly, it seems there are truths that are simply unbelievable, in virtue of their complexity, or other peculiarities.¹⁰⁸ Even more problematically, epistemic agents can have differing capabilities regarding the complexity of propositions they can represent mentally; as a result, it seems that the norm cannot both be universal, nor anything close thereto, and adhere to Kant’s law.

While I shall attempt to stake a moderate position between these two extreme views, the position for which I shall argue is closer to the unconditional formulation. I submit that our epistemic duty is, for every proposition about which one can form a belief, one ought to believe the proposition if it is true. Obviously, the pivotal term in the formulation is ‘can’; there are myriad senses in which it can be used, so we require some clarification along these lines if we are to come to a complete understanding of the present iteration of the truth norm.

Since I have endorsed doxastic involuntarism, there is one interpretation of ‘can’ that makes the term roughly identical with ‘will’ or ‘shall’. The principle of alternate possibilities is

¹⁰⁷ One might also wonder if there is anything of epistemic value in believing such a trivial proposition—I address this objection in section 6.4.

¹⁰⁸ For instance, if some metaphysical claim about the non-existence of the self is true, the very nature of belief may make this fact unbelievable.

false in the epistemic sphere, so there is only one doxastic profile that we can and will have. This line of reasoning harkens back to the argument from doxastic voluntarism. With hope, I have already shown both that we need not use the term ‘can’ in this sense, and that doing so is detrimental to the whole project of normative epistemology. It follows that we should adopt a broader understanding of ‘can’.

The opposing, maximally broad sense of ‘can’ brings us into a more complicated metaphysical discussion. It is unclear that my epistemic faculties are a part of my *quidditas*. As such, while I, in the actual world, might be incapable of believing the truth of some contingent, compound proposition involving trillions of other propositions, it may be nonetheless true that, in some possible world, there is something that shares in my essence, yet has this capability. Unless some range of epistemic capabilities is an essential part of what it is to be me, then I can believe any true proposition.¹⁰⁹ If possible, I would like to avoid having to answer this question here. We can observe at the very least that there is a sense of normativity with which we are concerned that slips through this very broad sense of ‘can.’ There is little value in evaluating me in terms of what I would believe if I had the mental processing power and storage space of an omniscient being. Our interpretation of ‘can’ should be narrowed in order to make the evaluations of the truth norm more applicable, though not to utmost degree, since that interpretation is also unfruitful for our purposes.

I have already hinted at my preferred strategy for finding a mean between these two extreme interpretations in section 2.4. That which we can believe is a function of our epistemic faculties: the categories we possess, our perceptual systems, our memory, etc. If our faculties are such that they can represent the world truly, then our epistemic obligation is to do so.

¹⁰⁹ We may also need to add my spatio-temporal condition to the ‘unless’ clause of this sentence.

In adopting this strategy, we are in a position to note that the truth norm is misleadingly simple in its formulation. While our obligation is to believe the truth, we are not always in a position to know which truths other agents, and indeed, even ourselves, are capable of believing. As such, the norm, despite its universal formulation, is not universal in the obligations it generates for each agent. This will make applying the norm difficult, both to ourselves and to others. It is difficult to know in one's own case, and even more so in cases of evaluating another, what truths one could have believed had one taken a different undergraduate major, etc. These difficulties, however, are strictly problems of application, and do not, by themselves, give us reason to think that the determinations of the truth norm thus interpreted are incorrect. Furthermore, this move to agent-relative norms means that the legitimacy of the truth norm is not universal to belief if one can think of believing agents that possess otherwise radically different psychological constitutions, but rather partly a result of human psychology.

Evaluations in terms of psychological possibility admit of two viable scopes. In the first place, we might be interested in determining how one is faring with regard to one's epistemic duty at this very moment. If so, then we need to make the corresponding determination concerning his or her faculties: which truths is he or she capable of believing right now? Alternatively, we could be interested in long-term evaluations. These would include not only the ways in which an individual can represent the world mentally, but also the ways in which he or she could (literally) change her or his mind in order better to represent the truth.

Both scopes have advantages and disadvantages. The present-term evaluations are easier to apply, since there are far fewer variables involved in determining what a person can do in a given moment. However, these evaluations also forgive what appear to be cases of culpable ignorance. Someone who wished for convenience's sake to remain ignorant of some troubling

fact might intentionally avoid acquiring the concepts necessary for believing the truth of the matter. For instance, I might keep myself ignorant of the parts of my car so as to avoid the beliefs (and corresponding unpleasant and expensive activities) pertaining to proper maintenance thereof. In such cases, if the subject succeeds, she or he would not transgress his or her epistemic duty, despite the obvious deviousness of such a strategy. The long-term evaluations handle these cases better; it is obvious that I can learn about my car, and as a result could no longer remain non-culpable in such cases of ignorance. However, these evaluations may be too harsh in some cases. There are a great many things about which I could learn, undoubtedly more than I could manage in a single lifetime. If my epistemic duty is to obtain all the true beliefs I could, then, regardless of what I actually end up believing, there will be truths that I could have believed (if I had acted and learned differently) but do not.

Although we can see that both interpretations mentioned above are beset with difficulties, the problems with the evaluations of narrower scope are worse than those that plague those of wider scope. Essentially, the objection to the latter is one of over-demandingness. On this interpretation, the truth norm demands more than can seriously be required of one, and, therefore, should not be thusly interpreted. However, if we recall that there is no good reason to expect that epistemic duty fulfillment or transgression will correspond to our practices (or axiological determinations) concerning praise or blame, as was argued in chapter 4, then the weighty demands of the truth norm are not so objectionable. One may non-culpably fail to believe every truth that she or he might have done. As such, she or he ought not to be blamed for failing to do so; however, this does not change the fact that he or she would have been better, epistemically, had he or she done so.¹¹⁰

¹¹⁰ Some ethicists have claimed that determinations of over-demandingness do not reflect a problem with norms, but rather with us, if we find them too demanding— memorably, (Singer, 1972). Intuitively, this argument

It may seem that endorsing a universal version of the norm is too great a price to avoid cases of culpable ignorance. I return to this concern in the next section, but I would here like to provide some supplemental reasons for endorsing the truth norm in an unconditional (or close thereto) form. On the assumption of doxastic involuntarism, we can make two broad observations about belief. In the first place, the only consideration that matters to us in forming a belief is whether the proposition considered by the belief is true.¹¹¹ Second, we are opinionated: if something seems true to us, then we believe that it is true. We find ourselves adopting beliefs regarding propositions that have no plausible value other than their truth.¹¹² The first observation motivates the demand that we believe only if a proposition is true, and the second motivates that we believe if it is true. Both of these characteristics are open to manipulation: we can intentionally avoid evidence with the hope of preventing the belief suggested by that evidence; we can also focus disproportionately on some evidence in the hopes of gaining a belief, and in so doing exclude other evidence, the total balance of which would suggest our targeted belief is false. Nonetheless, the fact that such goals require manipulation of our belief-forming processes suggests that the general rule concerning how we believe fits with the phenomena here mentioned. If this is a psychological fact about how we believe, and epistemic justification is a matter of believing well, then we ought to adopt an unconditional formulation of the truth norm.

Having covered the main clarificatory issues involved in articulating a norm of truth, we are left with the following formulation: *for every proposition about which one can form a belief,*

could be applied to epistemology as well, but the task may be more difficult given the possible interpretations of ‘can’ discussed above.

¹¹¹ Hieronymi (2005) and Lynch (2009) address the phenomenon in further detail.

¹¹² That is, the beliefs have no plausible alternative value. I am not sure if propositions can be the bearers of any value.

*one ought to believe the proposition if and only if it is true.*¹¹³ For brevity's sake, I omit recapitulating the two discussions concerning the scope of 'can' and of the evaluation in general, respectively. While, with hope, I have motivated a truth norm generally, and the precise formulation I give specifically, there remains a bevy of objections to such a rule; I now turn to presenting and addressing the most prominent thereof.

6.4 Objections to the Truth Norm

Objections to a truth norm of belief tend to follow two broad strategies. In the first place, some object to the spirit of the rule; these tend to aim at showing that some epistemic value deserves preferential consideration to truth, or that truth is not so valuable as to merit a rule aimed solely at securing it. Second, there are objections that object to the letter of the truth norm; these objections can accept the centrality of truth, but still object to implausible or untenable consequences of following particular versions of the truth norm.¹¹⁴ I shall begin with a review of some objections that adopt the first strategy.

I have already provided the means to responding to one of the more prominent objections to the spirit of the truth norm. Specifically, the objection is aimed at a deontological endorsement of the norm. It goes as follows: (1) according to the DCEJ, epistemic justification is a function of the extent to which one fulfills one's epistemic obligations. (2) The truth norm stipulates that we are epistemically obliged to believe the truth. (3) Therefore, one is justified only if one believes truly. (4) However, there are false, justified beliefs. It follows that the DCEJ cannot coherently endorse the truth norm.

¹¹³ I assume that if one can form a belief about a proposition, then one can also form a belief about that proposition's negation. Again assuming that contradictory beliefs are not possible, this formulation has the means to make error deontologically forbidden.

¹¹⁴ Of course, from this it should not be inferred that authors objecting to the letter of the truth norm thereby accept its spirit.

Obviously, the best route of response for my position is to object to (4). Paradigmatic cases of false, unjustified beliefs tend to lean on internalist intuitions concerning the nature of justification. One such case, discussed in chapter 4, is one in which someone believes in accordance with correct evidential and inferential principles, but is the victim of the systematic deception of a Cartesian demon. Since the victim is not at fault for her or his deception, his or her beliefs are justified, but false. However, because I have shown that the DCEJ need not endorse internalism, there is no reason to accept the verdict of this thought-experiment.

Not all purported instances of justified, false beliefs appeal to internalism; Goldman's process reliabilism would deem justified a false belief that was the result of an otherwise reliable cognitive process. Once we move to considering externalist theories, however, there is less reason to insist on the possibility of such cases. Since reliabilism emphasizes the importance of a non-accidental connection with the truth, it might coherently require perfect reliability for justification. To put this point more generally, *qua* externalist, there is no reason to think that, in principle, false, justified beliefs must be possible. Once one denies any necessary correlation between access and justification, one accepts the possibility that an agent might be blameless, and still have unjustified beliefs; the truth norm is only an extreme instance of such a possibility. One might think that false, justified beliefs are possible on the merits of additional argumentation, but the point ought to be the conclusion of a discussion, not asserted at the beginning thereof. If my arguments in favor of a truth norm succeed, internalism is false, and externalism should deny the possibility of false, unjustified beliefs.

One challenging objection to the letter of the truth norm is the existence of "blindspot" propositions—propositions the truth of which depend on the error of a given believer. For instance, it may be true both that it is raining and that I do not believe that it is raining. By

believing the two conjuncts, the proposition becomes false, and thus ought not to be believed. By not believing the conjuncts, they remain true, and thus the proposition ought to be believed.

(Bykvist and Hattiangadi 2007: 280) To some extent, putting the truth norm in terms of psychological possibility can disarm this objection. The strength of the objection depends on the fact that blindspot propositions cannot be believed, but I have already deemed that the scope of the norm should only apply to propositions that can be believed. Nonetheless, replying through these means leaves something to be desired; it is clear that one has made an error, and there seems no reason to think that one should be exempt from one's epistemic duties specifically because one erred.¹¹⁵

The following diagnosis leads to a more satisfactory treatment of blindspot propositions. I ought simply to believe that it is raining. In so doing the blindspot proposition will be falsified, and therefore not a barrier to my meeting my obligations regarding the truth. Obviously, the interesting aspect of these propositions is that they take place after I have formed the errant belief. Even so, if any blindspot proposition is true of me, then it will also be the case that there is a truth that I could believe, yet do not.¹¹⁶ The only way for me not to transgress my duty is to believe the non-self-referential truth. Since only believing that it is raining will result in my not transgressing my duty, I ought so to believe.¹¹⁷

¹¹⁵ i.e., intuitively, the determination is that the person should simply have checked if it was raining, or at least not believed anything on the matter if she or he was unable so to ascertain.

¹¹⁶ In this case "it is raining."

¹¹⁷ One of the authors against whom Bykvist and Hattiangadi's objection is explicitly aimed, Wedgwood (2002), gives a similar response. However, since he argues for the truth norm in general, rather than from a deontic perspective, his response is put in modal terms rather than in terms of duty-fulfillment. (2013: 136-138)

McHugh (2012) develops additional problematic cases for the truth norm. He asks one to consider a proposition that, if believed, will prove true and if disbelieved, will prove false.¹¹⁸ Since either outcome will result in adherence to epistemic duty, it seems that we ought both to believe and not to believe, but this is absurd. (13) There is a much more forgiving interpretation of such cases. According to the truth norm, we ought to believe or to disbelieve, but not both. Only one of the outcomes can be believed, and neither will result in a transgression of epistemic duty, so these cases, far from being problematic, are a free pass of sorts.¹¹⁹

Both of these sorts of problem cases operate by looking at an instant in one's epistemic life. However, since, in identifying the correct scope of the obligations generated by the truth norm, I have already provided independent motivation for thinking of epistemic evaluation as operating over a period of time, we have a natural means of looking past some of the apparent rigidity of the truth norm as I have formulated it. Since duty will be concerned with how one stands with regard to one's obligations over a relevant period of time, appealing to problems that only hold together for an instant will not be an effective strategy in arguing against the truth norm. Blindspot propositions are problematic because one has already failed to believe a truth, while self-fulfilling beliefs appeal to the undetermined truth-value of certain propositions. However, by looking a little ahead or behind, temporally, one can adopt a perspective from which these cases are of no particular concern.

McHugh also argues against the apparent implications generated by the truth norm regarding the value of withholding belief. (15-16) There are two ways of drawing out these

¹¹⁸ His example is about a belief concerning the result of a tennis match one is playing, noting that the psychological effect of believing that one will win or lose might have a decisive effect on the outcome of the match. (13-14)

¹¹⁹ The truth norm would still dictate that one ought to have some belief on the matter, rather than, say, withholding belief altogether.

implications as an objection to the truth norm. In the first place, one might argue that, at least in some instances, to withhold belief is the correct doxastic attitude. This objection, however, rather clearly depends on internalist intuitions that have already been called into question. The only apparent reason for preferring withholding to believing truly seems to be that withholding would secure a better fit to a subject's evidence. With hope, enough has been said along these lines already.

Alternatively, one could challenge that the truth norm incorrectly determines that withholding and erring are of equal undesirability; this seems to be McHugh's qualm. (15) Since the truth norm determines that our obligation is merely to believe the truth, and both withholding belief and erring fail to satisfy this obligation, the truth norm reprimands both states equally, but surely, it is argued, withholding is not as bad as erring with regard to our epistemic obligations.

The basic observation behind this objection, that erring and withholding have the same deontic status, is correct, and must be accepted if the truth norm is to be endorsed; however, its evaluation can be somewhat more subtle. In most cases, beliefs bear important inferential relations to other beliefs; from my belief that it is Saturday, I can infer that it is the weekend. Given these relations, errant beliefs pose a greater risk of leading to yet more error. From withholding belief, I can make no inferences, and therefore, at the very least, do not increase my susceptibility to fail to meet my epistemic obligations elsewhere. Thus, while the value is not directly relevant to epistemic justification, we can see that there is some reason to prefer withholding to erring in our attempts to fulfill our epistemic obligation.

Even if the truth norm can handle the most pressing objections to its motivation and articulation, there is more to be said in a defense thereof. In particular, though I have argued for the centrality of true belief to normative epistemology, there is a variety of alternative sources of

epistemic value. Although I am already committed to saying that accommodating these values is not essential to the viability of the truth norm, if I can show the truth norm can account for these alternative values as well, then the thesis of this chapter will prove more congenial to the general literature in epistemology.

6.5 Derivative Norms

While the truth norm has some radical implications when endorsed as the most fundamental norm in epistemology, it is more amenable to many of the traditional views than it might seem. The goal of this section is to consider some of the prominent alternative fundamental epistemic norms, and to consider the extent to which a view that adopts the truth norm can accommodate the former norms as derived from the latter.¹²⁰ Thus, for each of the following subsections, the methodology is largely the same. First, I identify the view and its central claims. Accomplishing this task in any great detail would exceed our present interests, so the presentations at this stage will be admitted over-simplifications, with hope, to no great detriment in coming to understand the theories. Since the focus of this chapter is normative epistemology, rather than meta-epistemology, the norms considered below should all be thought of as taking place within a deontological conception of epistemic justification. Having already endorsed a deontological metaepistemological position, the extent to which the truth norm can accommodate norms that presuppose different metaepistemological theories is of little interest. Second, I attempt to give a simple formulation of the norm taken by that theory to be fundamental. Third, I consider the relation between the truth norm and the norm identified by the previous step. Of course, some norms are derived more readily or more completely from the truth norm than others, but the general claim of this section is that the truth norm accommodates

¹²⁰ Here I understand the term ‘derived norm’ as follows: Norm A can be derived from norm B iff being obliged by norm B suffices for being obliged by norm A.

at least a variety of the motivations to rival theories, if not directly the norm endorsed by rival positions, to a degree insufficient to make its endorsement a radical departure from traditional normative epistemology.

Any additional rules that we identify as being entailed by the truth norm do not introduce new epistemic obligations. Because derived norms are necessary consequences of the norm from which they are derived¹²¹, they can only speak to a narrower set of conditions, or appeal to intuitions by putting their value in terms of something else of apparent importance.¹²² The obligations indicated by derived norms need to be followed as surely as do those of the truth norm, but they introduce no new information, nor do they generate any possibility of conflicting obligations. To cite a more familiar instance of such a process from the ethical sphere, if we accept that we ought not to treat someone as a mere means, then it follows that we ought not to enslave that person. Adhering to the former suffices for the latter, and the two can never generate conflicting obligations, but it might be important, depending on the context, that we recognize that such an implication holds. In such a case, we could have reason to articulate the derivative norm, even if it is a necessary consequence of the more general rule. Because of the traditional significance of some of the norms considered below, it is a worthwhile endeavor to show if, and if so, to what extent, they can be derived from the truth norm.

6.5.1: Foundationalism

Given that the two historical champions of the DCEJ are also foundationalists, one might expect the central commitments of foundationalism to mesh well with any deontological view. This is not the case, however, and I shall argue that, of the major views on the nature of justification, foundationalism is the least compatible with a truth-centric deontological theory.

¹²¹ And not also sufficient, since they are not the same norm

¹²² e.g., reliability or coherence

According to the foundationalist, justification is a function of the extent to which our beliefs are founded on some special class of basic beliefs. The basic beliefs are special because, they do not stand in the same need of justification as other beliefs. They are, as Chisholm notes, an epistemic analog of an unmoved mover. (1966: 33, n16) Such beliefs may acquire their special status for a variety of reasons, the particulars of which we need not concern ourselves here. What must concern us, however, is that there is no particular value to any foundational interrelations of beliefs when considered from the end of truth. True beliefs could be intuited, inferred from false beliefs, or from true beliefs using poor reasoning. Any true belief will meet the necessary criterion for justification identified by the truth norm, while any foundationalist view will say of at least some true beliefs that they are unjustified.¹²³

The two theories are slightly more compatible at an idealized extreme. Consider the ideal epistemic agent according to the truth norm— one who believes all and only true propositions. It may well be the case that one's doxastic profile must reflect a certain foundational structure in order to reach such a state. If this were true, then those who are justified according to the truth norm would also be justified according to foundationalism.

Nonetheless, this accord is largely superficial, and the psychological fact posited to bring the two views into harmony is far from established. The core motivation of foundationalism seems essentially internalist, and the truth norm is essentially externalist. This is not to say that foundationalism is wholly unconcerned with truth; presumably, the class of basic beliefs will be chosen in virtue of its fitness to serve as the base for a sound epistemic structure, one of the requirements of which is surely a wide array of true beliefs. However, by reducing the requirements for justification to merely a matter of the correct correlations between beliefs (as

¹²³ The only way of circumventing this difficulty would be to equate true beliefs with basic beliefs, but on such a view, no beliefs would be justified, because no true belief would have any foundations.

opposed to the content, particularly the truth, of basic beliefs), foundationalism distances itself from the truth norm. This move seems motivated by the sort of reasoning we saw endorsed by blame-centered, internalist versions of the DCEJ earlier. If the arguments in chapters three and four have succeeded, then foundationalism is left as an unappealing alternative to the truth norm. In other words, if the endorsement of true beliefs is an important part of normative epistemology, and truth is relevant to a belief's normative status (i.e. externalism is accepted), then there is no reason to reduce justification merely to a matter of relations between beliefs. A good epistemic structure as evaluated by the truth norm may reflect a foundational structure that would be endorsed by foundationalism in some cases, but there is little reason to suggest these two traits will necessarily coincide.

6.5.2: Coherentism

If foundationalism has us imagine our epistemic lives as an architectural structure, then the corresponding image of coherentism is that of a web.¹²⁴ Justification is again understood largely in terms of the co-relations of beliefs, but on this view, no particular beliefs hold a privileged position in the conferring of justification. Instead, large, coherent sets of belief are preferred to small, incoherent sets of beliefs.¹²⁵ The central norm regarding justification can be construed in terms of sets of beliefs or individual beliefs. Concerning sets of beliefs, the prescription would be to hold the largest possible maximally coherent set of beliefs. Concerning individual beliefs, the determination would be that beliefs are justified if and only if (or to the degree that) they are coherent with the rest of the subject's beliefs.

¹²⁴ Or a raft, as Sosa (1980) suggests

¹²⁵ 'Coherence' is understood to be a relation somewhat stronger than consistency, such that the truth of members of a coherent set supports the truth of the other members. Prominent works in this tradition include Lewis (1946), Bonjour (1985) and Lehrer (1990).

The extent to which this norm can be derived from the truth norm will settle along the lines of one of the more significant debates concerning coherentism, namely, whether or not coherence is truth conducive. The origin of the debate comes from Klein and Warfield (1994)'s answer in the negative.¹²⁶ Obviously, if coherence is not truth conducive, then there will be little reason for a system that accepts the truth norm as primary to recognize a derivative value in coherence. Rather than attempting to determine the winner of this debate, it will be simpler to make our response conditional. It is a mark in favor of coherentism if it can show that coherence is truth conducive.¹²⁷ Since coherentism is an alternative to the truth norm, the more generous move to my opposition will be to assume that coherence *is* truth conducive; if it is not, then the stakes of the subsequent argument are lowered significantly. If coherentism is implausible in its own right, then the ability of the truth norm to derive it is unimportant.

If coherence is truth conducive, then, given that coherence is defined in part by reference to truth, truth should also be conducive to coherence. To the extent that one believes all (and only) truths, one's set of beliefs will be coherent. Therefore, adhering to the truth norm will suffice for adhering to the norm of coherence as well. Of course, one might also say that adhering to the norm of coherence (perfectly) will result in adhering to the truth norm, since the set of all truths is also the most coherent set. The best reason for thinking of coherence as a norm derived from truth, rather than vice-versa, is that the truth norm enjoys a better non-epistemic grounding than coherence. The importance of truth can be explained in terms of the arguments in the preceding sections, while I submit, the importance of coherence can only be derived from the epistemic value of truth.

¹²⁶ One of the notable responses, for this project in particular, is Cross (1999).

¹²⁷ See, for instance, Bonjour (1985, 7-8)

6.5.3: Evidentialism

We have already discussed evidentialism in some detail in the second and third chapters, so here we need only recall that the central norm thereof is that one ought to believe in proportion to one's evidence, or that one is justified to the extent that one believes in proportion to one's evidence. Our two exemplars of evidentialism, Feldman and Steup, are also both internalists. Therefore, the greatest challenge that one can expect the truth norm to face in attempting to provide a derivative version of an evidentialist norm will be in accommodating the internalist intuitions that are connected with the motivations for evidentialism.

Depending on how one defines 'evidence,' the truth norm can fare quite well in deriving an evidentialist norm. If one employs a functional definition of the term, such that 'evidence' is just that which is used as rational support for a given belief (or, more importantly, that there is nothing that is evidence unless it actually plays an evidential role with regard to some belief), then the truth norm can accommodate evidentialism. With hope, this is an uncontroversial stipulation, as the opposing claim would require that there could be evidence that is not evidence *for anything*.

As I have presented it, the truth norm requires that we believe all and only (psychologically believable) truths. In satisfying this requirement, we also satisfy the requirement of evidentialism. Since all of our beliefs would be true and all of our evidence would be the result of accurate beliefs, all of our beliefs would fit our evidence. This is perhaps illustrated more clearly through a *reductio*. In order for the truth norm to fail to derive the evidentialist norm, it would have to be possible to satisfy the former while failing to satisfy the latter. To fail to satisfy the latter would be to have a discrepancy between one's actual doxastic profile and the one recommended by one's evidence. Assuming one satisfies the truth norm,

one's doxastic profile will be the set of all believable truths. For one's evidence to recommend a departure from this set would require that one have false or at least incomplete evidence. Having given a functional definition of evidence, it is not possible that one can believe the set of all truths while having false evidence, since false evidence would lead to false belief, or at least require erring regarding what is suggested by one's evidence. Denying that one can have incomplete evidence in such a situation requires appeal to a psychological fact: namely, that if one believes p , q , and that p is evidence for q , then p is evidence for q for the subject (i.e. the subject will use p as evidence for q). Since, by hypothesis, one believes all truths, including those regarding evidential relations, one will have perfectly complete evidence. By reductio, satisfying the truth norm suffices for satisfying the norm of evidentialism.

Although this argument also appeals to a psychological fact, as did the argument from the truth norm to the foundationalist norm, the claim about when one uses evidence seems less controversial, especially having already assumed the truth of doxastic involuntarism. To recognize something as evidence, yet not use it as such would amount to a strong disregard for the truth, against which the phenomenon of doxastic involuntarism speaks strongly.

6.5.4: Reliabilism

Despite their agreement on the importance of truth to epistemic justification, the central norm of reliabilism cannot be derived from the truth norm. Adjusting the kernel of the theory to a deontological formulation that makes beliefs the objects of its epistemic evaluations, we are left with the following norm as representative of reliabilism: a belief is justified only if (or to the extent that) it is the result of a reliable belief forming process. We can hypothesize a subject that has all and only true beliefs, but does not make use of any reliable belief forming processes. In such a case, the subject would meet the requirements of the truth norm, but would fail to meet

those of the above reliabilist norm. So, being obliged by the truth norm is not sufficient for being obliged by the reliabilist norm.

Even so, one might wonder whether whatever unspecified processes our hypothetical subject did adopt could fail to be reliable, since the known outcomes of those processes are perfectly true. Settling this matter would likely involve grappling with the difficult ‘generality problem’, a concern regarding the delineation of scope for the processes involved in the reliabilist analysis of justification. Instead, the conceptual distance between the two theories can be made clear by observing that the truth norm lacks the means for evaluating processes in any terms other than their direct, actual products: the beliefs they produce. So, unless evaluating belief-forming processes and beliefs is fundamentally the same activity, the possible divergence of the two norms is ineliminable. Any minimally generous interpretation of reliabilism would yield the result that there is a difference in these two types of evaluation, since, if there were no difference, reliabilists would likely have expressed their norm in terms of beliefs, and not processes.

However, in considering the basic motivations of reliabilism, it is not clear that the emphasis on processes should be essential to the theory. I shall take Goldman’s (1986) argument as my model, since his is the most prominent version of the theory. Let us begin with his observation that justification is best thought of in terms of adherence to a rule-based framework. (60) This much is consistent with the DCEJ, since according to it, justification is a function of adherence to one’s epistemic duties, which can presumably be expressed as rules. Goldman also commits to forming his normative epistemological system in terms of beliefs (14), so we can assume the correct rules of our system (J-rules, as he calls them), essentially to involve beliefs. Later, he describes two forms of J-rules, direct and indirect: “A direct form of J-rules would

expressly permit certain beliefs... [a]n indirect form of J-rules would license the execution of certain belief-forming processes or operations; the rules would not specify the belief output of these licensed processes.”¹²⁸ From indirect rules, we are permitted to use a cognitive process or operation, which means justification is a matter of permitted state transitions. (77) Finally, Goldman refines the notion of state transitions to processes, the latter including a causal element absent in the broader notion of state transitions. (85) We can note that, in these terms, the difference between the truth norm and the schematic reliabilist norm is that the former is direct, and the latter is indirect. What we lack from Goldman is a reason to think that the set of correct J-rules should include only indirect rules. The same motivations that prompt the requirement for reliable belief forming processes in an indirect J-rule support the requirement for true beliefs in a direct rule. So, while the central norm of reliabilism is not derivable from the truth norm, absent further argument in favor of indirect rules, reliabilists should at least recognize the theoretical motivations of their own norm as supporting the truth norm as well.¹²⁹

6.5.5: Proper Functionalism

I already briefly considered the prospects of a deontological version of proper functionalism in chapter 5. Recall that the basic norm of this view is that beliefs are justified (or have positive epistemic status, to adopt Plantinga’s more neutral name for the concept) if and only if, or to the extent that they are the result of the proper function of one’s epistemic faculties. Thus, a deontological version of this norm would deem that it is one’s epistemic duty to hold

¹²⁸ At this point, Goldman has already argued against rules that using the operators of obligation and prohibition. In order fully to unite the two views under consideration in this subsection, I would need to respond to this argument as well; however, for the time being, I shall only focus on the significance of the division between direct and indirect J-rules.

¹²⁹ This is not to make the very bold claim that all reliabilists should, on pain of inconsistency, adopt the categorical version of the truth norm that I have advanced earlier in the chapter. Instead, my suggestion is that the plausibility of such a rule is greatly increased when one starts theorizing from some set of truth-centric assumptions about normative epistemic value.

beliefs that are the result of the proper function of one's epistemic faculties. The best case for deriving the norm of proper function from the truth norm would be if the proper function of our epistemic faculties resulted in true beliefs. In such a case, satisfying the truth norm would be a necessary condition of having fulfilled the norm of proper function, but it would not be a sufficient condition; true beliefs that result from epistemic malfunctions are justified according to the truth norm, but unjustified according to the norm of proper function. Therefore, the norm of proper function cannot be derived from the truth norm. Since proper functionalism only evaluates the beliefs one has, it cannot derive the positive demands of the truth norm.

In one sense, this incompatibility should come as little surprise, since the concept of proper function has no analytic connection to truth. Coherence entails consistency, reliabilism requires a favorable ratio of true beliefs to false beliefs while evidentialism and foundationalism require that we make logically valid (i.e., truth-preserving) use of our evidence or basic beliefs, respectively. A deontological version of the norm of proper function can only explain what makes a particular function proper in terms of adherence to epistemic duty.¹³⁰ But if one's epistemic duty is just to have beliefs that are the result of proper function, then this definition of justification is circular. Proper functionalism very clearly is meant to derive its epistemic significance from the design plan or designer that fixes which functions are proper, but this is not a move available to the DCEJ.

Conclusion to chapter 6

Having finished the previous chapter with an endorsement of a deontological conception of epistemic justification, this chapter's aim was to provide a substantive epistemic norm so as to provide another portion of a complete theory of epistemic justification. The theory now on offer

¹³⁰ If it were to appeal to something else, then adherence to duty would not be the basic concept in terms of which justification is defined, so the theory would not be properly deontological.

is one that is deontological, and that sees our fundamental epistemic duty as believing truly. Merely adopting the truth norm will not suffice for a complete theory of epistemic justification, because even through giving a non-epistemic grounding for the truth norm, we lack the conceptual resources to preclude the legitimacy of other epistemic norms derived from other sources of epistemic value. The importance of belief as the basic unit of epistemic agency suggests the truth norm might hold some privileged position amongst our possible epistemic duties, but the stronger, negative existential claim, that there are no other duties, is too nebulous a question to advance a confident answer here.

The past two chapters have discussed the meta and normative levels of epistemology, but little has been said concerning the applied level. Having adopted a truth norm, the usefulness of developing any such theory is minimized. The goal of an applied epistemology would be to tell one what to believe in a given situation. However, since what we ought to believe is just what is true, one must have access to the truth in order to complete this goal. If one already has access to this truth, then one will already believe it, because of doxastic involuntarism. As such, the most to which applied epistemology could amount would be testimony— something of the form: “believe P because it is true.” Since it seems implicit in any act of testimony that one ought to believe the testified proposition, the form above would reduce to “P is true.”

If applied epistemology must remain a vacuous level in a theory of justification, the conclusions of the previous two chapters approach a complete account of epistemic justification. However, a significant question remains for normative epistemology, broadly construed: what is the relation between our epistemic obligations and, among others, our moral obligations? For instance, one might adopt the position that epistemic obligations are always trumped by moral obligations. Without understanding this relation, we cannot claim a complete understanding of

epistemic obligations, since we do not yet know if they hold unconditionally (in an *ultima facie* sense). The next chapter shall explore the standing of epistemic norms in the larger normative landscape.

CHAPTER 7

THE SPECIES OF NORMATIVITY

If the previous chapters have succeeded, we have a wide array of epistemic obligations. If a great many other philosophers are to be believed, we have yet more obligations from other sources, including, but perhaps not limited to, moral obligations.¹³¹ What, if anything, do our epistemic obligations mean for our moral obligations, and vice-versa? I shall answer this question, among others, in this chapter. The answer that I shall defend is rather stark: ‘nothing.’ In addition to an incommensurability between the spheres of moral and epistemic normativity, the prospects for an all-things-considered, general normativity are bleak. In order to support this position, I examine two concepts: an ethics of belief and an epistemology of action. In both sections, I argue that a good ethical and epistemological theory, respectively, will need to incorporate these concepts. Next, I submit that neither can accommodate the other. With this shown, the only remaining option for bringing the various normativities into concert with each other is to appeal to some unifying, higher-order normativity. In the third section, I defend the position that any such higher-order normative theory will fail to accommodate *any* of the lower-ordered systems. Accordingly, the chapter concludes by observing that it is possible that we could have obligations according to various species of normativity the fulfillments of which are mutually incompatible, and that such conflicts do not admit of any satisfying resolution.

¹³¹ I shall use the categories ‘moral’ and ‘ethical’ interchangeably in this chapter.

7.1 The Epistemology of Action

This section is named for the analogous partner to the far more famous ‘ethics of belief.’ To my knowledge, there is no systematic attempt to provide a normative epistemological theory of action,¹³² though, as discussed in the next section, there are notable attempts to provide a normative moral theory of beliefs. An epistemology of action describes the epistemically normative status of actions, such that some actions might be described as (epistemically) justified or unjustified.

There are a few objections to the possibility of such a theory that merit immediate consideration. In the first place, I have already committed myself to carrying out epistemic evaluations in terms of belief (sect. 6.1). Actions are not beliefs, so they are unsuited to bearing ascriptions of epistemic justification or the lack thereof. I think this objection is correct in that it views belief as the bottom line of epistemic value, but it takes too rigid a view of the matter.

Consider the following example: I am unsure of the official language of Luxembourg, and I am in a position to ascertain the fact through a Wikipedia search.¹³³ I contend that I ought, epistemically, to do so. To the extent that it is in my power to fulfill my obligation to the truth, I ought to do whatever I can to fulfill my obligation. However, the example merits some unpacking. In the first place, there is no necessary connection between the action and the epistemic state. I might attempt the search, only to find that my internet connection has been disabled. In this case, the action has been performed, but my normative-epistemic status has not changed, since I am in the same position as before with regard to my obligations. However, this only shows that *some* actions will fail to change my normative-epistemic status; one can still

¹³² There are uses of the phrase ‘epistemology of action’ extant in scholarly literature, but these uses are in the descriptive, rather than normative, sense of ‘epistemology.’

¹³³ In fact, Luxembourg has three official languages: French, German and Luxembourgish.

maintain that one ought to perform those actions that will improve one's standing and avoid those that will harm it (or those that will fail to raise it). Of course, we are not always in a position to know what actions will result in such changes (e.g., we don't always know if our internet connection is working), and one might be concerned that we could be subject to an overwhelming number of action-regarding epistemic norms without our ever being aware of them. This concern, however, while widely held, ultimately relies on some internalist principle of justification, which I have already argued that we should reject.

We can also note that epistemic obligations to act seem to be permissive, rather than categorical, like obligations to believe. It seems a psychological fact that, for any belief, there is more than one way one could come to have the belief: I could ask a civically well-informed Luxembourger, succeed with a lucky guess, or be the recipient of divine inspiration, all of which would put me into the requisite state. Thus, our epistemic obligations regarding actions might be better described as disjunctive, such that we ought to perform at least one of the possible actions available to us that would result in our holding a given true belief. These obligations, however, still suffice for the viability of an epistemology of action.

One might also have recurrent concerns over doxastic voluntarism in endorsing such obligations. If we have epistemic obligations regarding actions, then it seems at least some epistemic 'oughts' entail voluntary 'cans,' contrary to my characterization of the entailment relation in chapter two. However, it was only through avoiding the entailment of such voluntary 'cans' that I suggested one might be able to disarm the argument from doxastic voluntarism. To resolve this concern, we can return to a minor point mentioned in chapter two. For at least some beliefs, we do have voluntary indirect control (flipping the light switch to change our belief about whether the lights are on, etc.). We noted that this control would be insufficient to make all

the demands of the DCEJ fulfillable, but there is no difficulty in ascribing obligations that only range over actions that are epistemically efficacious. Suppose these two statements are true:

“June 1st, 2017 CE falls on a Wednesday” and “June 1st, 2017 BCE saw unseasonably cool temperatures.” I can perform an action to make myself believe the former, but not the latter.

Therefore, the only belief with regard to which I am obliged to perform an action is the former.

We can admit of instances that entail indirect doxastic voluntarism where appropriate without having to make the whole normative system entail doxastic voluntarism in any strong sense.

It is difficult to provide anything but a very general description of what an epistemology of action would involve. This is largely because, in endorsing an extreme form of externalism, I have effectively eliminated the ability for a theory of justification to provide guidance, either with regard to what we ought to believe¹³⁴ or, more applicably to the present topic, what epistemically efficacious actions we ought to take. Thus, the prescriptions regarding actions will be very idealized as well: one ought to perform all actions that will effect true beliefs in oneself. The most striking aspect of this stark formulation is that it admits of no exceptions with regard to ethical or prudential concerns; even if the action is ethically or prudentially bad, we ought still to do so. We might wish to clarify the previous sentence by noting that the ‘ought’ it employs is used strictly in the epistemic sense. However, so long as epistemology is accepted as a legitimate species of normativity, and therefore a field the determinations of which can give rise to true ‘ought’ claims, there is some sense of the word in which the sentence is true. From an epistemic perspective, on the other hand, this implication should not be surprising. However one wants to settle the nature of moral value, it is unlikely to have any necessary connection to true belief.

¹³⁴ That is, any specific recommendations— the theory can tell us that we ought to believe the truth, *qua* truth, but not what propositions are true, and thus what ought to be believed.

Because I have described epistemic obligations to action as permissive, rather than categorical, some cases of conflicts between epistemic and moral obligations could be resolved if one of the disjuncts of one's epistemic obligation were morally permissible, while perhaps some others were not. Even so, it seems clear that there could be cases in which there is only one action that will satisfy one's epistemic obligations, which happens to be morally impermissible as well. Suppose Eddington was running late in his expedition to photograph the solar eclipse of 1919, which helped to corroborate the general theory of relativity. The only way he could have made it to an adequate site to take the photographs in time was to take a shortcut over a tract of private land, and he did not have time to ask for permission to cross. Epistemically, he ought to have crossed the land so as to help bring about the breakthrough in physics resulting from the shift to the paradigm of relativistic gravitation. Morally, he ought not to have trespassed. It is not the case that both of these obligations could have been respected.¹³⁵

In starting with a discussion of the epistemic obligations to action one might have, I have taken on the more controversial topic already. As previously noted, the idea of an ethics of belief is comparably well developed. Nonetheless, in order to establish the broader thesis of the chapter, we must also see about moral obligations to believe, and if these obligations can conflict with our epistemic obligations. If we could derive all moral norms from epistemic norms, then it might still be possible for there to be a unified theory of normativity; if neither species can accommodate the other, then this possibility is much harder to defend.

¹³⁵ This example is not worded in terms of my primary interests in epistemic justification: true belief, but the epistemic significance of general relativity theory should be enough to motivate most peoples' normative-epistemic intuitions.

7.2 The Ethics of Belief

This section adopts a structure similar to that of the previous. In the first place, I consider what an ethics of belief is, and what it might require. Since this project is not focused on the topic of moral justification, I will not appeal to any substantive moral norms, so the description here will be very general. Because the idea of an ethics of belief is comparatively well accepted, I shall not focus on objections to the very possibility thereof. Instead, I shall consider Clifford's famous account of moral obligations concerning our beliefs, accepting it as a good general principle for a moral theory to adopt. I then argue that any system of moral norms will (possibly) come into conflict with our epistemic obligations. I conclude by considering a view according to which all epistemic norms are a subset of moral norms, and argue against the persuasiveness of this position.

Applying the analogous description from the previous section should give us a clear understanding of what an ethics of belief would look like. Leaving aside the question of what special property (or properties) bestows moral justification, we can note that the essential claim of an ethics of belief is that some beliefs have moral justification while others are morally unjustified, or that one might be morally unjustified in the absence of some requisite belief.

The analogy also suggests that moral obligations to believe need ultimately to be understood in terms of the actions they motivate (or fail to motivate).¹³⁶ This in turn suggests that these obligations are also permissive: as long as the requisite action is performed (or the forbidden action avoided), it matters not exactly what belief one holds on the matter.

Conceptualizing how belief gives rise to action is rather more complex than the case of how

¹³⁶ This suggestion is not supported as well as the analogous claim in the case of epistemology, since I have not given an argument to the effect that actions are the proper object of moral evaluation. However, I think any moral theory must deal with action at least at some level; e.g. some actions are indicative of a vicious character, the failure to perform some actions are indicative of willing a maxim that cannot consistently be willed as a universal law, etc.

actions can give rise to belief.¹³⁷ It may be the case only one set of beliefs can give rise to a given action A, namely, the practical syllogism concluding in the belief ‘I ought to do A.’ In any case, the beliefs subject to moral evaluation will be action-regarding.

Certainly, the most famous instance of a theory concerning the ethics of belief is that of W.K. Clifford. His position is articulated memorably in the dictum “it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence.” (2001: 18) This dictum alone might be read as an epistemic principle, but Clifford’s reasoning that leads to this conclusion, that credulity risks turning society into a “den of thieves” (16) and that a credulous person “is father to the liar and the cheat”, (17) suggests that this obligation is a moral one. Clifford takes a wide view of the actions that can be effected by our beliefs; since the dictum regards any beliefs, it need not be the case that only beliefs that *cause* us to act in a particular way are subject to moral evaluation. Since believing on insufficient evidence gives rise to credulity, and credulity to the general collapse of society, all of our beliefs are subject to moral norms.

Without adopting a normative ethical theory, I cannot provide a well-reasoned endorsement of Clifford’s views here. Nonetheless, his famous example of the negligent ship-owner agrees with most people’s moral intuitions. This is not to say that our intuitions about the example prove the moral general dictum mentioned above, but rather that Clifford’s position appeals to the common recognition that some beliefs can be morally significant, and that some moral assessments are incomplete if they fail to include an evaluation of the agent’s beliefs.

Curiously, it is not as easy to contrive examples that show that our moral obligations regarding belief can require us to transgress our epistemic obligations. Some have suggested that

¹³⁷ This is not to suggest that the latter is not complex, but merely that most, if not all, of the complexities had by the latter are also had by the former.

we have moral obligations to believe favorably concerning the character of loved-ones, at least in the absence of compelling evidence to the contrary. Assuming such an obligation, a conflict between the moral and the epistemic would arise if it were true (though not evident) that a loved-one was of a vicious character. However, that we have any such obligations is contentious. Some of Clifford's examples suggest that we can have a moral obligation to believe falsely; however, these only occur when one has misleading evidence, so there is a suggestion that the moral and epistemic are not really at odds in such a case. If one did have better evidence, then there would be no conflict between the moral and the epistemic. Even if we cannot contrive an example of a moral requirement to believe something that is epistemically forbidden, it is worth noting that the example from the previous section, that of Eddington trespassing to take his photos of the eclipse, works just as well when considered from the reverse perspective; after all, what a normative theory tells us we ought not to do is just as significant a topic as what it tells us we ought to do.¹³⁸

Nonetheless, I think one can construct more convincing analogous cases by taking a more practical view of belief, particularly in considering how it can motivate us to action. Suppose that my neighbor's house is on fire, but that I care little for any of the residents other than the family dog. Suppose further that the dog is out of the house, but one of the resident humans is trapped within, and that it is in my power to save her or him. If I believe the truth of the situation, I will do nothing, and thereby fail to fulfill my moral obligation to save the resident. If I believe falsely, that the dog is in the house, I will attempt to rescue it and, seeing as how I had already done half the job, save the human. If my motivational structure is as described, then I suggest it

¹³⁸ "Works as well" in the sense that we can generate conflicting obligations; not, necessarily, that this is a case of a particularly difficult conflict on which to pick a side

is my moral obligation to have the false belief.¹³⁹ Again, the obligation is disjunctive: if I believed falsely that a friend had been trapped inside during a botched attempt at burglary, I might be motivated to take the same action, but, crucially, no true belief will suffice to motivate me to perform my moral obligation,¹⁴⁰ so I have a moral obligation to have a false belief instead.

The project of deriving all epistemic norms from moral norms seems like a more feasible task than the analogous project of the previous section. One advantage it holds is that, because ethics has received significantly more attention than normative epistemology, there are fewer widely accepted epistemic norms that need to be derived. The examples discussed in this chapter should suffice to show that a strong truth norm of belief cannot be derived from ethical norms, since many, if not all, of the most plausible ethical norms can be put into competition with true beliefs. Perhaps, though, one could fare better in deriving other epistemic norms. Although it is not clear that he rejects the possibility of epistemic norms, Clifford's ethics of belief might provide a framework for deriving epistemic norms from moral norms. In particular, he seems to endorse a restrictive evidentialist norm (it is wrong always, etc. to believe on insufficient evidence) from a broadly utilitarian norm according to which credulity is to be avoided. The universality of the evidentialist norm means it is rather effective in guiding evaluations of beliefs.¹⁴¹ It is therefore true that we can derive a sophisticated system for the evaluation of beliefs solely from moral norms.

¹³⁹ Given how I have construed moral and epistemic obligations, it may be the case that if one has a *moral* obligation to believe, then that obligation entails doxastic voluntarism. Thankfully, since I am interested only in defending a view of epistemic justification, I need not explore that potentially problematic entailment.

¹⁴⁰ It would be best if I had a better motivational structure, such that I wanted to save the person in the house as well. In this case I assume that is not the sort of thing I could change in the window of time available for me to perform the morally required action.

¹⁴¹ Assuming, of course, we can work out the details of what makes for 'sufficient' evidence.

Nonetheless, the success of any such system is compatible with the existence of other systems of evaluation derived from other species of normativity. In the previous chapters, I have given positive arguments for why we should think our beliefs are subject to a norm derived from purely epistemic concerns. Unless the advocate of the ethical-as-foundational position can undermine all the norms derived from purely epistemic concerns, it does not matter how sophisticated or expansive an ethics of belief he or she might put forth. If anything, the more applications an ethics of belief has, the more conflicts we are likely to encounter between our epistemic and moral obligations. Even if I am wrong about the truth norm of belief, and we are obliged by an epistemic norm the requirements of which coincide exactly with the norm generated by Clifford's ethics of belief, we still cannot say that all epistemic norms are derived from ethical norms. This outcome would eliminate the possibility of conflict between obligations, but if both norms were motivated by considerations about the acquisition of two independent types of *value*, then the absence of conflict would be merely the result of a happy coincidence, rather than the success of the ethical-as-foundational project.

Finally, one might also wish to claim that "moral norms trump non-moral norms." One might recognize the existence of, and potential conflict between, different species of norms, but claim that one enjoys a privileged status. I think this claim must appeal to one of two implicit justifications. First, moral norms might enjoy a privileged status as a result of some ordering determined by a general normativity; that possibility is challenged in the next section. Alternatively, one might look to some intrinsic quality of moral norms to establish their privileged status. A platitude concerning morality is that, even if it may be the imprudent thing to do, one ought to meet one's moral obligations. If the principle illustrated by this platitude raises moral norms above prudential ones, then perhaps the same can be said in comparing the former

to epistemic norms. Rather than consider this application, I want to raise a brief consideration against the platitude. The word ‘ought’ is ambiguous, as it might refer to what we ought to do according to different species of normativity. Thus, it might be true that we ought to favor moral norms to prudential ones, but also that we ought to privilege prudential norms to moral ones. This apparent contradiction can be resolved by making clear the sense in which one uses ‘ought.’ Taking ought_m as an abbreviation for ought in the moral sense, and ought_e for the corresponding epistemic concept, we can say that we ought_e privilege our epistemic duties while we ought_m privilege our moral ones. Since we are still without any means to balance the two different senses of the term, this seems a prudent point at which to consider the prospects of a general normativity.

Not all conflicts are worth dwelling on, however. A norm of kindness derived from moral considerations may conflict with some norm that aims at the maximization of suffering, but that conflict is neither philosophically interesting, nor psychologically distressing. Not having laid the groundwork for any theory of moral justification, I am not in a position to give an argument to why we should care about moral norms, but it seems a safe presumption that most people have at least some interest in being moral; no one treats suffering as an unconditional good.¹⁴² Epistemic norms are more like moral norms in this respect, since, as I argued in chapter 6, our psychology betrays a standing interest in adhering to the truth norm. If a conflict stems from two sets of norms that are derived from considerations about what make us what we are (*qua* human, or rational agent, etc.), then we have good reason to be concerned about such conflicts, and to hope that there is a satisfying solution thereto.

¹⁴² Some people may be undisturbed by causing suffering in the pursuit of some other aims, like personal amusement, but presumably do not value the suffering in itself.

7.3 General Normativity

What I have in mind in using the term ‘general normativity’ is partially, though not entirely, in keeping with the classical usage of genus and species. A general normative theory should account for that which all species of normativity have in common, but it should have more than just explanatory power. In particular, we are here seeking a means to resolve conflicts between species of norms. If we succeed in finding such a theory, contrary to the assertion at the end of the previous section, there would be an *ultima facie* sense of the word ‘ought’, the correct application of which would suffice for our present desideratum. Given the task of the present section, we would do well to reflect on the criteria according to which a general normativity could be said successfully to have accommodated a species thereof.

We should think of a complete preservation of the specific norm as neither necessary nor sufficient for its accommodation in a general theory. It is not necessary, since, on such a demand, the existence of inter-special conflicts would make a general theory of normative incoherent. It is not sufficient, since the theory is to be general, and hence not identical with the species. What we need instead is some degree of commensurability of the species. The general theory should put the species into common terms, in such a way that we are, by virtue of the general theory, in a position to make comparative evaluations. Although a merely descriptive commonality will not enable us to make such evaluations, considering the general properties of norms seems like the best place to start in considering the prospects for the general theory.

Norms aim at the promotion or securing of some value. I have identified truth as the fundamental epistemic value and accordingly endorsed a norm that mandates true belief. The two most prominent loci of ethical value are utility and good will. Utilitarianism and Kantian ethics endorse norms that, if followed, would promote or secure these values, respectively. If this is an

accurate characterization of ethical and epistemic norms, and if the language of species and genus is to be taken seriously, this suggests that any general normativity would adhere to the same structure: first, identify the appropriate value; second, formulate and endorse norms that, if followed, would suffice to secure or promote that value.

Identifying the value sought by a general normativity is surely the more difficult of these two tasks, and giving any kind of substantive hypothesis on the topic would be rather presumptuous. On the current conception, this value will account for what is valuable in the True and the (moral and prudential) Good. Rather than put forward any candidate for this role, I wish only to work with the observation that, whatever it is, this value will not be the value identified by any of the species of normativity.¹⁴³ Let us call this value GV for ease of reference. Now consider the specific value of truth. I have argued that truth has epistemic value because beliefs succeed only if they are true, and beliefs are the fundamental means of epistemic evaluation. If the truth norm is the central element of normative epistemology, then GV will have to account for the value of adhering thereto. But if GV is some value other than truth, we are left without motivation to follow the truth norm, since the only value following it will secure is truth. Thus, truth must be an element of GV. If this is true, and the two are not identical, then the value of truth must be like an element of the broader set of GV.

If we could articulate a general value that accounts for the values targeted by all the (correct) specific norms, then we would have a single reason to justify every true ‘ought’ claim. However, this conception of GV does not solve the problem that motivated its introduction in the first place, namely, that there are (at least two sets of) norms that have legitimate independent grounding, in the adherence to which we have inherent interest. It is not enough that all the

¹⁴³ For, if the value is to be general, it cannot be specific at the same time.

values are accommodated; rather, we require some sort of ordering, or hierarchy of the norms if we are to use GV to resolve the significant conflicts we face.¹⁴⁴ However, if GV is merely a set of independent values, then there could be no rational basis for preferring one value to another. Only if we have an additional element, the valuing of which suffices to provide some resolution to the hypothesized conflicts is GV useful. If there is such an additional element, however, then GV is really *only* that additional element, and we have returned to the option rejected in the previous paragraph.

This problem is exacerbated by, but not dependent on, the categorical formulation of the truth norm given in chapter 6. If we ought to believe all and only truths, and every belief is either true or untrue, then all instances of belief have a normative standing with regard to the truth norm. So, for any instance of a GV endorsing the pursuit of utility through some belief at the cost of truth of that belief, GV will motivate the transgression of a legitimate, binding norm. If we consider a weaker epistemic norm, say, that we ought to have a high ratio of true beliefs to false ones, which might also be motivated by the value of true belief, instances of pursuing utility in preference to truth could be consistent with this. In such a case, the epistemic norm is respected, thus, so too is the value targeted by said norm. The content of the norms to which we are subject determines the frequency and prospective solubility of conflicts, but the basic meaning of two values being divergent is the fundamental source of the possibility of conflict.

Note that there may also be conflicts within specific normative theories; one recalls Aquinas' doctrine of double effect.¹⁴⁵ Thus, we may have a *prima facie* obligation not to kill, but

¹⁴⁴ Though it would be unfair to the position I am presently considering to suggest this ordering had to be categorical (such that, for instance, moral norms always trump epistemic norms); rather, it would simply have to provide us with some reason to prefer some of our obligations to others in cases of conflict. Note also that this ordering cannot be permissive (i.e. 'just pick one') since that prescription would put us in no better a position than if there were no general normativity at all.

¹⁴⁵ Discussed in the *Summa Theologica, Secunda Secundae Partis*, Q64

that obligation can be over-ridden by familiar concerns about people being tied up to railroad tracks. In the case of morality, however, there needs to be some common element, in virtue of which all the various (moral) norms to which we are subject are all moral norms (and not norms from some other species). This common element is the plausible source for our ability to resolve intra-moral conflicts.¹⁴⁶ Since the various specific norms are motivated by specific considerations, we cannot expect the same finality in our attempts to dissolve inter-normative conflicts.

We may consider the argument in simpler terms. Any general theory of value will appeal either to a mere aggregation of specific values, or to some value that does not include, but suffices to explain the values targeted by the species of normativity.¹⁴⁷ If the former strategy is adopted, we will lack a systematic way to resolve inter-normative conflicts. If the latter strategy is adopted, then none of the specific norms can be motivated, since all the specific norms hold in virtue of considerations regarding specific values. A general value could suffice to generate *more* norms, but it would only produce a unified, conflict-free normative theory if there were also no binding specific norms. I have argued that there is at least one such norm: the truth norm of belief.

¹⁴⁶ One might also think that there could be no moral conflicts, in which case the picture is decidedly clearer (from the theoretical standpoint, if not the practical).

¹⁴⁷ Again, the middle ground, which treats general normativity as an aggregation of the specific values while bringing in some independent reason to resolve inter-normative conflicts is really just the second horn of the present dilemma, since the independent reason should suffice to explain why we should prefer the obligations generated by one norm in favor of the other. If the independent reason did not suffice to explain the value of the specific norms, then it would be unclear why it should be treated as having the authority to settle the dispute. For example, I were to appeal to the principle “God says that in cases of conflict between moral and epistemic obligations, if failing to adhere to the moral norm would result only a venial sin, then respect to the epistemic obligation; if failing to adhere to the moral norm would result in a mortal sin, then respect the moral obligation”, I could resolve inter-normative conflicts by an ordering that follows from some divine-command theoretic conception of moral and epistemic normativity. If epistemic or moral normativity was derived from different conceptions of value (as I have argued is the case), then we would have no reason to try to apply this principle.

7.3.1: Eudaimonia and General Value

The concept of eudaimonia, as originally developed in Aristotle's *Nichomachean Ethics*, grounds a prominent group of theories that purport to accomplish exactly what I have deemed impossible in the earlier sections of this paper, namely, articulating a common ground for moral and epistemic normativity. Although it is not typically considered as such, I think that eudaimonist theory offers nothing more than a well-articulated type of general value, hence the present section's being included in the larger treatment of general normativity. To support this claim, I must accomplish three things in this section. First, we need a clear articulation of the conceptual commitments of eudaimonism. Second, I must establish that the value offered by eudaimonia is only a particular attempt at providing a general value, in the sense described in the previous section. Third, I must show that the arguments from the previous section apply to the particular account of general value in question.

I shall focus on Aristotle's conception of eudaimonia, if only because of his foundational role in the development thereof. Other eudaimonist theories exist, but presumably do not stray too far from the concept's landmark articulation. The term is most literally translated as 'good spirits', but is more popularly translated as 'happiness' or 'flourishing.' More technically, it is the condition that is necessary and sufficient for having a good life. The eudaimon life is characterized by the proper function (i.e., the virtue) of our various psychological faculties, and it is the particular faculties that should provide the basis for the various species of normativity. (NE: 1098a8-17) The significance of each of the species is presumably to be derived from their role in the eudaimon life. Although the analogy is not perfect, the two broad categories of normativity we have discussed to this point, the ethical and the epistemological, correspond

roughly to the virtues of practical wisdom (phronesis) and wisdom (sophia), respectively. (NE: Bk. VI, ch. 7)¹⁴⁸

Clearly, this strategy is an instance of taking the second horn to the dilemma presented in the previous section. The value of adhering to one's normative and epistemic obligations is to be found in their serving as a necessary condition for eudaimonia; for instance, if the eudaimon life were impossible for one, then one would have no reason to be virtuous. Furthermore, if we accept Aristotle's premises, then eudaimonia provides an acceptable account of a general value.¹⁴⁹ Nonetheless, I have, I hope, already provided sufficient argumentation to support the view that we ought not to accept the present account of our epistemic and ethical obligations.¹⁵⁰ I need not deny Aristotle's claims about the nature of virtue or eudaimonia in order to support this conclusion. Rather, we need simply to observe that the value enjoyed by eudaimonia might exist, but stem (in part) from the fulfillment of epistemic norms other than the truth norm. Even if this is so, we are not given a reason to think we are not also bound by the truth norm, as I have argued we are. In other words, even if we can successfully derive a set of epistemic norms solely through considerations about a general value like eudaimonia, we are not in a position to deny the legitimacy of other normative systems promoting specific values, nor can we excuse ourselves of these more specific obligations for the sake of avoiding normative conflicts. To preclude the possibility of conflicts between norms, the eudaimonist must claim that eudaimonia

¹⁴⁸ I say 'roughly' because there seems to be an epistemic element to practical wisdom, even if it is wisdom regarding how one ought to act.

¹⁴⁹ This is facilitated by the suggestion that wisdom concerns eternal things, and practical wisdom concerns temporal things. On such a view, we could not be epistemically obligated to act, nor ethically obligated to believe in the manner described in section 7.1

¹⁵⁰ It may be somewhat unnatural to deem the recommendations of a eudaimonist theory 'obligations.' However, they hold to roughly the same conditional structure as a deontological obligation, namely "you must do/believe thus-and-so if you want to be morally/ethically good" only the eudaimonist instead targets eudaimonia.

is the only value that can give rise to a normative system. If the previous chapter has succeeded, this cannot be done.¹⁵¹

7.4 Objections and Replies

One seeking a happier resolution of these potential conflicts might appeal to a provision made in the fifth chapter, concerning the respective merits of deontological and virtue epistemology, respectively. There, I concluded that if deontology enjoys a greater degree of legitimacy, it stems from its narrow focus on truth, rather than the broader concerns of virtue theory. At this point, however, it has become apparent that the narrow focus of a deontological theory that endorses the truth norm comes with the side effect of allowing for irresolvable conflicts between norms that obligate us. Thus, a virtue theorist might now claim that what was earlier taken to be an advantage of the deontological theory is now seen to be a problematic tenet. As such, the argument continues, the DCEJ has no theoretical advantages to virtue theory, so the latter ought to be preferred to the former.

While this argument might give us reason to *hope* that virtue theory is true and the DCEJ is false, it is an illegitimate inference to accept or reject a theory on the grounds of its conforming well or poorly to our hopes regarding the truth. If the arguments in favor of the DCEJ and the truth norm are sound, then, like it or not, there are likely irresolvable conflicts between the normative demands with which we are faced.

A different eudaimonist objection finds support in the preeminent desirability of the eudaimon life. There are surely a great many norms that could make a claim on human action, not all about which we should care. Some, though coherently formulated, will lack any sort of force in their demands; consider for instance, the norm that requires us to have only an even

¹⁵¹ It is perhaps imprudent to phrase the antecedent thus, as it depends on a very controversial position; rather, if there is any legitimate, non-eudaimonist normative epistemological theory, then the consequent holds.

number of beliefs. Our continuing efforts to meet the demands of the truth norm would often put us at odds with this norm, but it is obvious that we should not care about our failure to oblige to the even-belief norm. If the truth norm is like the even-belief norm, then, even if there are legitimate cases of conflicts between attaining truth and eudaimonia, we can resolve the conflict by noting that only one of the norms should interest us. However, I have argued that the truth norm has a special status due to the nature of human psychology: the fact that we all believe in accordance with how things seem to us indicates that we have an interest in adhering to at least some form of a truth norm of belief. Because we have a demonstrated interest in adhering to the truth norm, as well as to attaining eudaimonia, we should be concerned about cases of conflicts between these two goals.

A final concern about the possibility of normative conflicts takes the project back to its point of departure: the ‘ought implies can’ principle. If we have two obligations, the satisfactions of which are formally incompatible, it follows that we ought to do each. If we ought to do each, then we ought to both, and if we ought to do both, then we can do both, contrary to hypothesis. To put the argument formally, where P and Q are two legitimate obligations the fulfillment of which are incompatible:

- 1) $O:P$ [assumption]
- 2) $O:Q$ [assumption]
- 3) $\sim\Diamond(P \ \& \ Q)$ [assumption]
- 4) $O: (P \ \& \ Q)$ [from 1 and 2]
- 5) $\Diamond(P \ \& \ Q)$ [from 4, Kant’s law]
- 6) $\Diamond(P \ \& \ Q) \ \& \ \sim\Diamond(P \ \& \ Q)$ [from 3 and 5, RAA]

Thus, one interpretation of such examples would be to say that contrary obligations stand opposed to Kant's law, so it cannot be the case that both are legitimate obligations.

However, even in the absence of a well-developed semantics for deontic logic, we can see that the inference from the conjunction of two obligations to a conjunctive obligation (step 4 above) is invalid. As is precisely the case by hypothesis, the fulfillment of some obligations might preclude the fulfillment of other obligations, so even if it is possible that we fulfill our obligations P and Q, it nonetheless might be impossible that we fulfill both. We can also present the response formally:

- 1) $O:P$ [assumption]
- 2) $O:Q$ [assumption]
- 3) $\sim\Diamond(P \ \& \ Q)$ [assumption]
- 4) $\sim O:(P\&Q)$ [from 3, contrapositive of Kant's law]¹⁵²

As such, there remain no convincing objections to the mere possibility of conflicting obligations from the various species of normativity, nor any satisfying means of resolving such conflicts if one accepts their possibility.

Conclusion:

The conclusions of all of the preceding chapters admit of clear and concise summaries: Contrary to the established view, the DCEJ's sole essential commitment is the claim that epistemic justification is a function of adherence to epistemic duties. One such duty is a requirement of true belief in proportion to the psychological constitution of the believer. Instances of conflict between this duty and other species of obligation do not admit of theoretically rigorous resolution.

¹⁵² Evidently then, even without the details of the semantics for our provisional deontic system, it will have to be different from those of the more popular modal logic systems, since the first proof, substituting the 'O:' operator for the 'necessary' operator, is a theorem therein.

The most apparent avenue for continued research regarding the DCEJ pertains to the conclusion of the sixth chapter. A complete deontological theory of epistemic justification should specify not just that we have some duties, but should aim for a complete list thereof. The grounding for the truth norm suggests that our epistemic faculties (e.g., our capacity to believe, etc.) should inform the obligations by which we are bound. However, one should also recall that the commitment to the truth norm is independent of the DCEJ proper, and further research may indicate that the latter functions as a better hypothesis in the absence of the former.

A complete list of epistemic duties, coupled with the analysis of epistemic justification at the core of the DCEJ, would provide a foundation for extensive expansion into other topics in epistemology. Traditional epistemology has treated justification and knowledge as importantly related, so one could reasonably hope for progress in research on the latter, having settled the former, even if such progress would still fall short of a complete analysis of knowledge. A deontological account of intra-agential justification could possibly be extrapolated to a set of normative epistemological rules concerning testimony and collective agency.

More realistically, however, continuing the project of the DCEJ is most likely to involve further research at the level of metaepistemology. As noted in chapter 5, it is methodologically difficult to construct arguments that carry much finality in that sphere of inquiry, and my efforts in that chapter are unlikely to prove conclusive. It is my modest hope that, even if the arguments in favor of the DCEJ are rejected, the division between the different levels of inquiry in epistemology is adopted; even if we are still asking questions, asking the right questions shall carry us a great distance toward our shared goal.

REFERENCES

- Ahlstrom-Vij, K. & Dunn, J. (2014). A Defence of Epistemic Consequentialism. *Philosophical Quarterly* 64 (257):541-551.
- Ahlstrom-Vij, Kristoffer (2013). In Defense of Veritistic Value Monism. *Pacific Philosophical Quarterly* 94 (1):19-40.
- Alexande, Larry–Moore (2007). Michael. Deontological Ethics. In Thaddeus Metz (ed.), *Stanford Encyclopedia of Philosophy*.
- Alston, William P. (1985). Concepts of Epistemic Justification. *The Monist* 68 (1):57-89.
- Andow, James (2016). Do non-philosophers think epistemic consequentialism is counterintuitive? *Synthese*:1-13.
- Bergmann, Michael (2006). *Justification Without Awareness: A Defense of Epistemic Externalism*. Oxford University Press.
- Berker, Selim (2013). The Rejection of Epistemic Consequentialism. *Philosophical Issues* 23 (1):363-387.
- Boghossian, Paul A. (2003). The normativity of content. *Philosophical Issues* 13 (1):31-45.
- BonJour, Laurence (1985). *The Structure of Empirical Knowledge*. Harvard University Press.
- Bonjour, Laurence (1980). Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy* 5 (1):53-73.
- BonJour, Laurence (1985). *The Structure of Empirical Knowledge*. Harvard University Press.
- Bykvist, Krister & Hattiangadi, Anandi (2007). Does thought imply ought? *Analysis* 67 (296):277–285.
- Chan, Timothy (ed.) (2013). *The Aim of Belief*. Oxford University Press.
- Chignell, Andrew (2016). The ethics of belief. *Stanford Encyclopedia of Philosophy*.
- Chisholm, Roderick M. (1966). *Theory of Knowledge*. Englewood Cliffs, N.J., Prentice-Hall.
- Clifford, W. K. (1999). "The Ethics of Belief". In *The Ethics of Belief and Other Essays*. Amherst, New York: Prometheus Books. pp. 70-97.

- Côté-Bouchard, Charles (2016). Can the aim of belief ground epistemic normativity? *Philosophical Studies* 173 (12):3181-3198.
- Cross, Charles B. (1999). Coherence and truth conducive justification. *Analysis* 59 (263):186–193.
- Cross, Charles B. (1995). Probability, evidence, and the coherence of the whole truth. *Synthese* 103 (2):153 — 170.
- Feldman, Richard (2000). The ethics of belief. *Philosophy and Phenomenological Research* 60 (3):667-695.
- Firth, Roderick (1981). “Epistemic Merit, Intrinsic and Instrumental”. *Proceedings and Addresses of the American Philosophical Association* 55 (1):5-23.
- Frankfurt, Harry (1969). “Moral Responsibility and the Principle of Alternative Possibilities”. *Journal of Philosophy* 66:829--839.
- Frankfurt, Harry G. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy* 68 (1):5-20.
- Ginet, Carl (1975). *Knowledge, Perception, and Memory*. D. Reidel Pub. Co..
- Goldman, Alvin I. (2015). Reliabilism, veritism, and epistemic consequentialism. *Episteme* 12 (2):131-143.
- Goldman, Alvin I. (2002). *Pathways to Knowledge: Private and Public*. Oxford University Press.
- Goldman, Alvin I. (1986). *Epistemology and Cognition*. Harvard University Press.
- Hall, Richard J. & Johnson, Charles R. (1998). The Epistemic Duty to Seek More Evidence. *American Philosophical Quarterly* 35 (2):129 — 139.
- Hazlett, Allan (2013). *A Luxury of the Understanding: On the Value of True Belief*. Oxford University Press.
- Hieronymi, Pamela (2005). The Wrong Kind of Reason. *Journal of Philosophy* 102 (9):437 — 457.
- Kant, Immanuel (1996). *Practical Philosophy*. Cambridge University Press.
- Klein, Peter & Warfield, Ted A. (1994). What Price Coherence? *Analysis* 54 (3):129 — 132.
- Lehrer, Keith (1990). *Theory of Knowledge*. Routledge.
- Lewis, C. I. (1946). *An Analysis of Knowledge and Valuation*. Open Court.

- Littlejohn, Clayton (2012). *Justification and the Truth-Connection*. Cambridge University Press.
- Lynch, Michael (2009). The Truth of Values and the Values of Truth'. In Pritchard, Haddock & Millar (eds.), *Epistemic Value*. Oxford: Oxford University Press.
- McConnell, Terrance (1988). Ross on Duty and Ignorance. *History of Philosophy Quarterly* 5 (1):79 — 95.
- Percival, P. R. (2002). Epistemic Consequentialism: Philip Percival. *Supplement to the Proceedings of the Aristotelian Society* 76 (1):121-151.
- Plantinga, Alvin (1993). *Warrant: The Current Debate*. Oxford University Press.
- Plantinga, Alvin (1988). Positive epistemic status and proper function. *Philosophical Perspectives* 2:1-50.
- Quine, W. V. (1969). Epistemology Naturalized. In *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Ross, W. D. (1930). *The Right and the Good*. Clarendon Press.
- Sosa, Ernest (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- Stace, W. T. (1952). *Religion and the Modern Mind*. Greenwood Press.
- Steglich-Petersen, Asbjørn (2006). No Norm needed: On the aim of belief. *Philosophical Quarterly* 56 (225):499–516.
- Steup, Matthias (ed.) (2001). *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.
- Steup, Matthias (2000). Doxastic Voluntarism and Epistemic Deontology. *Acta Analytica* 15 (1):25-56.
- Velleman, David (2000). *The Possibility of Practical Reason*. Oxford University Press.
- Wedgwood, Ralph (2002). The aim of belief. *Philosophical Perspectives* 36 (s16):267-97.
- Whiting, Daniel (2013). Truth: the Aim and Norm of Belief. *Teorema: International Journal of Philosophy* 32 (3):121-136.
- Williams, B. (1970). Deciding to believe. In Bernard Williams (ed.), *Problems of the Self*. Cambridge University Press. pp. 136--51.
- Zagzebski, Linda (2003). The Search for the Source of Epistemic Good. *Metaphilosophy* 34 (1-2):12-28.