

In Defence of the Epistemological Objection to Divine Command Theory

By John Danaher, NUI Galway

Forthcoming in *Sophia*

Abstract:

Divine Command Theories (DCTs) comes in several different forms but at their core all of these theories claims that certain moral statuses (most typically the status of being *obligatory*) exist in virtue of the fact that God has commanded them to exist. Several authors argue that this core version of the DCT is vulnerable to an epistemological objection. According to this objection, DCT is deficient because certain groups of moral agents lack epistemic access to God's commands. But there is confusion as to the precise nature and significance of this objection, and critiques of its key premises. In this article I try to clear up this confusion and address these critiques. I do so in three ways. First, I offer a simplified general version of the objection. Second, I address the leading criticisms of the premises of this objection, focusing in particular on the role of moral risk/uncertainty in our understanding of God's commands. And third, I outline four possible interpretations of the argument, each with a differing degree of significance for the proponent of the DCT.

1. Introduction

This article critiques Divine Command Theories (DCTs) of metaethics on the grounds that they involve an intimate connection between moral epistemology and moral ontology. More precisely, it argues that all DCTs incorporate an epistemic condition into their account of moral ontology and because of this they are vulnerable to an objection, *viz.* an important class of moral agents fail to satisfy this epistemic condition. This casts into doubt the metaethical adequacy of DCTs, and has significant implications for anyone wishing to rely upon them.

This objection to DCT is not new (see: Morrision 2009; Peoples 2011; and Wielenberg 2014). It is referred to as the 'Epistemological Objection' (Morrision 2009; Peoples 2011), though there is some confusion as to whether this name is appropriate given the argument's implications (Wielenberg 2014). In this article,

I hope to clear up some of the confusion associated with this objection, and provide a strengthened defence of it. I do so in four parts. First, I offer a simple, general version of the objection. This consists of two premises and a conclusion. Second, I defend the first premise of the argument, highlighting the important connections between moral ontology and epistemology in the DCT. Third, I defend the second premise, illustrating the problem with the primary objection to it from proponents of DCT and showing how recent insights into the phenomenon of moral risk/uncertainty might illuminate key features of this premise. Fourth, and finally, I consider the implications of the argument. Does it offer a strong objection to DCT? Or can proponents of DCT simply shrug their shoulders and say ‘So what?’

2. What is the Epistemological Objection?

To understand the epistemological objection, we must first understand DCT. Here, I take DCT to be a label that can attach to a family of metaethical theories, each of which is concerned with accounting for, explaining, or grounding the existence of one or more species of moral fact by reference to God’s commands (Murphy 2012).

Moral facts come in two main flavours. There are *value*-based facts, which concern whether a particular states of affairs is *good* or *bad* or *neutral*. And there are *deontic* facts, which concern whether or not particular actions are *obligatory*, *forbidden*, *permissible*, *supererogatory* (and so on).¹ Although it is possible for a DCT to be structured in such a way that it accounts for both value-based and deontic facts, this is not now the norm. The main reason for this lies in the classic Euthyphro dilemma, which challenges proponents of DCT to account for the necessity of certain moral values or duties (e.g. the duty to refrain from torturing innocent children for fun). If DCT is true, and if God is free to command anything, then it seems like the duty to refrain from torturing children could be contingent.

¹ Deontic logic is multivalent. According to some theories there are more than the four deontic statuses mentioned in the text, hence the ‘and so on’.

That does not sit well with many people's foundational moral intuitions.² Many DCTists resolve this dilemma by appealing to God's nature. They argue that God's nature necessarily exemplifies or instantiates certain properties like lovingkindness and that these properties in turn ensure that God would never command something terrible like the torture of innocent children. This approach is most closely associated with the work of Robert Adams (1999), though it has been defended by others (e.g. Alston 1991; Quinn 2006).

The upshot is a more complex understanding of the possible relationships between God and moral facts. God's commands are now thought to ground a limited set of moral facts, most typically the fact that some actions are obligatory, while other aspects of God's nature are thought to account for other moral facts. Wielenberg (2014) offers a useful categorisation of these theories. He argues that there is a general category of *theological stateist* theories which hold that moral facts are dependent for their existence on one or more of God's states (e.g. values might be said to be dependent on God's essential nature); within this general category there is a more discrete category of *theological voluntarist* theories, which hold that moral facts are dependent for their existence on one or more of God's voluntary acts (e.g. his intending or willing that X be so); and finally, within the category of voluntarist theories, there is the even more circumscribed category of DCTs, which focus specifically on God's commands.

Taking this onboard, in this article I assume that all DCTs are relatively narrow in scope. They focus solely on the relationship between God's commands and moral facts. And, in light of the responses to the Euthyphro dilemma, they only really focus on one type of moral fact, namely *obligations*. They try to argue that moral obligations are grounded in God's commands (Adams 1999; Quinn 2006). The essence of the DCT is thus that without a divine command there would be no moral obligation. Adams (1999) is most explicit about this commitment, holding that a divine *command* is essential because without it there would be no discernible difference between a supererogatory act (one that is

² Though whether we should trust those intuitions is a separate question. See Joyce 2002 for an analysis and critique of the Euthyphro dilemma.

above and beyond the call of moral duty) and an obligatory one. I will return to this feature of Adams's theory below, but even if some DCTists reject this particular claim concerning the necessity of the divine command for distinguishing between supererogations and obligations, I believe they will accept the dependency-relationship between obligations and divine commands.

It is this dependency-relationship that gives rise to the epistemological objection. Put most simply, this objection holds that certain classes of moral agent (specifically, *reasonable non-believers*) lack epistemic access to God's commands. In other words, they do not know or cannot reasonably be expected to know, what God has commanded. The result is that, if the DCT is true, then for this class of moral agents, moral obligations no longer exist. It is, however, wrong to suppose that reasonable non-believers have no moral obligations. Consequently, something must be rotten at the heart of the DCT. That's the essence of the epistemological objection.

There is some confusion as to whether this objection is strictly epistemological (Peoples 2011) or ontological in nature (Wielenberg 2014). There is also some confusion as to how serious the objection really is (Peoples 2011). I hope to clear up this confusion in what follows. I will argue that the objection is *both* epistemological and ontological in nature, and that it can be quite serious. This is because there is an important connection between epistemology and ontology in the DCT (and, indeed, in most metaethical theories); and the significance of this depends largely on what type of theist you are. To facilitate making these arguments, I will work with the following simple and generalised version of the objection. This version formalises some of the preceding discussion and is relatively abstract and non-committal in its scope. This contrasts it with previous, narrower versions (Morrison 2009):

(1) DCTs, either explicitly or implicitly, include an epistemic condition in their account of moral obligations, *viz.* you must either know or successfully receive communication (implying knowledge) of divine commands in order for you to be morally bound.

(2) There are such things as reasonable non-believers (*i.e.* non-believers who do not violate any epistemic duties in their non-belief) and for these reasonable non-believers (RNBs), satisfaction of the epistemic condition of DCT is not possible.

(3) Therefore, on DCT, there are no moral obligations for reasonable non-believers.

In the remainder of the article I clarify and defend the two premises of this argument, and assess its overall implications.

3. Why is an epistemic condition essential?

The first premise of the argument states that all DCTs incorporate an epistemic condition into their account of moral ontology. It is important to understand what this means. It is not simply that it would be a good or desirable thing for people to know what God has commanded. It is that without knowledge of the command, moral obligations fail to exist. Suppose God has commanded you to give 10% of your income to charity. Premise (1) is saying that under DCT this would give rise to an obligation only if you have knowledge of the command. Premise (1) is thus pointing to an intimate connection between moral epistemology and moral ontology in the case of the DCT.

What can be said in favour of premise (1)? For one thing, proponents of DCT already seem to accept it. This is certainly the case for Robert Adams, perhaps the most famous contemporary defender of a modified DCT. His theory is explicitly *social* and *communicative* in nature. He says that having an obligation is something that only arises in a social context, *i.e.* in the context of an interpersonal relationship between two or more individuals. And that consequently being placed under an obligation 'essentially involves

communicative acts' (Adams 1999, 262).³ One person must explicitly tell another person what is required of them in order for the other to be bound. In the case of moral obligations, the only person with the right kind of authority and nature to issue such directives is God. Hence, being under a moral obligation requires that one is in a communicative relationship with God. This implies that one must have epistemic access to the relevant communicative acts, *i.e.* God's commands. It is very difficult to see how this account of moral obligations could be sustained without incorporating an epistemic condition. The result, I believe, is that anyone who favours an Adamsian version of DCT should accept premise (1).

But the Adamsian version is not the only game in town. Peoples (2011) argues that at least some versions of DCT do not incorporate an explicit epistemic condition. An example is supposedly the theory defended by Quinn (2006). The key claim in this theory is that ontology and epistemology are *separate* and *separable*. How something comes to be (an ontological matter) is distinct from how we come to know of its existence (an epistemological matter). As Quinn himself says:

"Our theory asserts that divine commands are conditions causally necessary and sufficient for moral obligations and prohibitions to be in force. It makes no claims at all about how we might come to know just what God has commanded... After all, it is a philosophical truism that the causal order and the order of learning need not be the same."

(Quinn 2006, 44-45)

This seems emphatic, but two interpretative points need to be made. First, it is not clear from this passage whether Quinn completely rejects an epistemic condition. Elsewhere in the same section of text he talks about how someone could first come to know that something is forbidden and then come to know that it was commanded by God. This suggests that he thinks knowledge of an obligation is essential, but knowledge of the source is not. If so, then he does not

³ Note: this quote was originally sourced through Morrison 2009

really reject the epistemic condition. He merely offers a modified version of it. His objection then collapses into the ‘content-not-source’ objection which I discuss in the next section. Second, there is some dispute in the literature as to the correct classification of Quinn’s moral theory. Although Quinn referred to it as a command theory (Peoples 2011), others such as Morrison (2009) have suggested that it really belongs to the more general non-command-based voluntarist family.

Even if we grant that the Quinn-type theory belongs to the more general category of non-command-based voluntarist theories it probably makes no difference because something like the epistemic condition is going to be essential to all voluntarist theories. The main reason for my confidence in this regard is that most metaethicists now concede that a plausible account of moral ontology almost always requires the satisfaction of an epistemic condition. If DCTists want to deny this connection, so be it: their theory will be weaker as a result. I defend this view in two ways: by appealing to common methodologies of argument in the field of metaethics, and by analogy to an ongoing dispute within that field between moral realists and antirealists.

Let’s look at methodology first. As Joyce (2002; 2006) notes, most contemporary metaethicists approach the task of finding an appropriate grounding for moral facts in a similar manner (e.g. Joyce 2002 and 2006; Beyleveld 1991; Smith 1994; Enoch 2011; Adams 1999). First, they identify a set of moral platitudes, *i.e.* widely accepted statements about the nature of moral requirements. Examples would include things like ‘moral obligations are *other-regarding*’, ‘morality is *impartial*’, ‘moral obligations are *conatively independent*’, and ‘moral obligations are *universal* in scope (i.e. they apply to all similarly situated moral agents)’ and so on.⁴ These lists of moral platitudes can be more or less exhaustive (compare Smith 1994 and Enoch 2011) and there can be disputes about which ones should be included. Nevertheless, once the platitudes have been identified, the metaethicist’s job is to work out a grounding that best accounts for these platitudes. This is no easy task, with some arguing that no

⁴ This particular list of moral platitudes is taken, with some modifications, from Beyleveld 1991.

plausible grounding can be found (Joyce 2002 and 2006) and others arguing that the best we can hope for is a theory that accounts for a majority of them and so wins out because it has the most ‘plausibility points’ (Enoch 2011).

Who wins the war is not hugely important in the present context. What is important is that the DCT will be much less likely to win the war if it excludes an epistemic condition. The reason for this is that most widely-accepted lists of the key moral platitudes – particularly those relating to moral obligations – include conditions that either explicitly or implicitly require epistemic access to the demands of morality. There are two illustrations of this. One is that it is widely-agreed that moral obligations are action-guiding and motivationally salient. That is to say, they have some sort of effect on how we choose to behave (even if they do not completely overwhelm our other internal reasons for action). It is not clear how moral obligations could have these properties if the people to whom they are directed lack epistemic access to their content. Similarly, there is the widely-accepted Kantian maxim that ‘ought implies can’. This suggests that in order for something to count as a moral obligation we must have the ability to follow it. This ‘ability’ consists in the power to control one’s actions in the way specified by the obligation and also, crucially, the ability to *know* what the obligation demands.⁵ So an epistemic condition looks like it would be an important part of any plausible list of moral platitudes, and hence something that any sound metaethical grounding for moral obligations would account for.

This importance is reinforced by the fact that the absence of a plausible account of moral epistemology is at the heart of one of the most contentious contemporary debates in metaethics: the debate between realists and antirealists. Recent years have witnessed a resurgence in defences of non-natural moral realism (e.g. Shafer-Landau 2003; Parfit 2011; Enoch 2011; Wielenberg 2014). According to this theory, moral facts, including both values

⁵ For a defence of the claim that ‘ought implies can’ also implies knowledge, see Howard-Snyder 1997. In that paper, Howard-Snyder specifically critiques objective consequentialism, arguing that people cannot follow the core imperative of objective consequentialism because they don’t always know which action produces the best consequences. This view has been criticized on the grounds that Howard-Snyder confuses know-how with know-that in her argument (Andrić 2016). But even defenders of this criticism concede that know-which (i.e. knowledge of which actions are required by a given moral principle) is essential to the ‘ought implies can’ maxim (Andrić 2017). I am indebted to an anonymous reviewer for pushing me on this point.

and obligations, are akin to Platonic mathematical facts. They are abstract, mind-independent and metaphysically necessary features of reality. They do not depend for their existence on other more metaphysically basic facts (except to the extent that some moral facts depend on other moral facts). Moral antirealists have an opposing view: they think that moral facts (particularly obligations) are constructed out of practical attitudes shared by most human agents (Street, 2006; Southwood 2010). Antirealists sometimes argue that the non-natural realist position is problematic because it fails to provide a plausible account of moral epistemology (Street 2006; Joyce 2006). Realists generally agree (Enoch 2011) that we must *have* epistemic access to moral truth; that we must know, innately or through serious reflection what is morally required of us. But how does this come to be the case? On the realist view, it seems to be a mere coincidence that our moral beliefs and attitudes correspond with the mind-independent moral reality. But antirealists argue that this is troubling in light of the fact that our cognitive faculties are the result of a messy and highly contingent evolutionary history (Joyce 2006; Street 2006). Realists have some answers to this. They try to argue that there is a correlation between pleasure/pain and the moral properties of goodness/badness and since evolutionary processes are sensitive to the former, there is a reason to believe that our evolved cognitive faculties would pick up on these deeper moral truths (Enoch 2011; Skarskaune 2011; Wielenberg 2014). Who comes out top in this debate is unimportant; the important point is that realists generally agree that they owe the antirealists some account of moral epistemology, i.e. that unless their metaethical theory incorporates an epistemic condition (one that explains how we can come to know what is moral) it is deficient in some respect (Enoch 2010). My contention is that the proponent of DCT is in a similar position to the non-natural realist:⁶ they can ignore moral epistemology if they wish, but this means that their attempt to provide a grounding for morality ends up being much less plausible.

⁶ Not least because DCT is itself a species of non-natural moral realism, albeit slightly different from the version described in the text.

To sum up, I think that premise (1) of the epistemological objection is sound. Any proponent of the DCT that adopts the Adamsian social and communicative version of DCT would seem bound to accept that knowledge of God's communicative acts (i.e. commands) is an essential precondition for the existence of an obligation. And any proponent of the DCT (or voluntarism more broadly) that wishes to hang tough and insist upon the separation of moral epistemology and ontology faces a problem: their theory becomes much less plausible if they fail to incorporate an epistemic condition because epistemic conditions are included among the standard lists of moral platitudes that any satisfactory metaethical theory ought to account for.

4. Defending the second premise: Why RNBs lack the requisite knowledge

Premise (2) is the real sticking point of the epistemological objection. That premise has two parts to it. The first part claims that there are such things as *reasonable non-believers* (i.e. non-believers who violate no epistemic duties in their nonbelief); the second part claims that these non-believers lack epistemic access to God's commands (i.e. that they cannot know what God has obligated them to do).

I have relatively little to say in defence of the first part. There are three reasons for this. The first is that I think this element of the argument may be the one toward which proponents of DCT should direct most of their energy. I return to this point in the final section. The second is that, notwithstanding what I have just said, I think that others have already proffered admirable defences of the existence of such non-believers (e.g. Schellenberg 2007),⁷ and I have nothing new to add to such defences. Certainly, I like to think of myself as a reasonable non-believer, someone who was once happy to believe, but who lost their belief in the course of an honest search for the truth. I also think the existence of numerous people, both past and present, for whom belief in God was never a live option (for demographic or historical reasons) is good support for the existence

⁷ Schellenberg now prefers the term *nonresistant nonbelief* for reasons we will not get into. We stick with the term *reasonable nonbelief* on the grounds that it is still being used in the literature on the epistemological objection to DCT (e.g. Morrison 2009; Peoples 2011; and Wielenberg 2014).

of reasonable nonbelievers. The third reason I have little to say in defence of their existence is that I think the argument may not be limited to them anyway. As will become clear in the discussion below, I also think that there are particular classes of believers who are vulnerable to the same problem. Nevertheless, I keep the scope limited to reasonable non-believers at the outset because I think the best case can be made in relation to them.

This brings us to the second part of premise (2). This is perhaps the most crucial element of the argument and the one upon which most attention has focused. The *prima facie* defence of it is straightforward: reasonable non-believers do not believe in God. Consequently, they cannot interpret anything in the world as being a command emanating from God. To take an obvious example, they do not believe that Bible is the divinely inspired word of God. Thus, they cannot believe that the ten commandments are really divine commands. Consequently, they do not *know* that these are moral obligations (if indeed they are). This is just an example. According to premise (2), any candidate moral obligation under DCT will fall foul of the same basic problem.

Theists have developed responses to this *prima facie* defence. These defences typically start out by arguing for an enriched understanding of the possible forms of divine communication. God, we are told, need not simply communicate to us via a revealed text, he can also communicate to us via conscience, moral intuition, natural law (all understood as forms of general revelation) and special revelation.⁸ Once we have this enriched understanding of the forms of communication, we can start to see how it is possible for a reasonable non-believer to acquire the requisite moral knowledge. Two

⁸ Several authors make this point about the diverse forms of divine communication. Peoples (2011) summarises the contributions to the debate. See also Evans (2013), pp 37-45 on the different ways in which God's commands may be promulgated. As one of the anonymous reviewers to this paper pointed out, missing from the list provided in the text are Kant's suggested methods for using reason to arrive at moral knowledge in the *Groundwork on the Metaphysics of Morals* (2012/1785 at 4: 401-402 and 4:421-424) and *Critique of Practical Reason* (2015/1788 at 5:25). The Kantian method would seem to deliver moral obligations as dictates/commands of reason. But it is noteworthy in this regard that Kant did not think that the moral law had an author. He thought God had a role to play in morality, but that it was a regulative/practical role, not a constitutive/grounding role. He argued that we needed to postulate God in order to make it practical to attain the highest good (2015/1788 at 5: 113-132), but that the content of the moral law itself was grounded solely in reason. This Kantian view of God's role in the moral order is different from the one being targeted by the epistemological objection. On this interpretation of Kant's argument, see Byrne (2011).

arguments in particular can be used to support this conclusion. I call them the 'content-not-source' argument and the 'probability-possibility' argument. I deal with them both in turn.

4.1 – The Content not Source Argument

The 'content not source' argument comes from the work of Robert Adams, as modified and expanded by Glenn Peoples, and from the work of C. Stephen Evans. It is the leading response in the current literature. I'll consider both the Adams/Peoples' version and the Evans version separately since, although they are quite similar in their overall thrust, there are some subtle and important differences between them.

Let's start with the Adams/Peoples' version. As mentioned earlier, Adams is one of the proponents of DCT who concedes the link between ontology and epistemology. He thinks that the DCTist owes some account of how it is that people come to know what is commanded of them. To do this, Adams offers an account of what it would take for God to successfully issue a moral command. This results in something I shall call 'Adams's Communicative Standard'. It has three elements to it:

Adams's Communicative Standard: 'In my opinion, a satisfactory account of these matters will have three main points. (1) A divine command will always involve a *sign*, as we may call it, that is intentionally caused by God. (2) In causing the sign God must intend to issue a command, and what is commanded is what God intends to command thereby. (3) The sign must be such that the intended audience could understand it as conveying the intended command.' (Adams 1999, 265)

The first two conditions are unproblematic. We can readily concede that God, if He exists, would have the power to intentionally cause the creation of signs that convey his commands. The crucial condition is the third. It is this one that allegedly causes problems for the epistemological objection. For when you look at it, this condition is claiming that access to God's commands does not

specifically require knowledge of the *source* of the command. It only requires knowledge of what is being commanded. This creates a problem for the epistemological objection because when you factor in the different possible forms of divine communication (i.e. the different possible 'signs') then you can imagine cases in which the reasonable non-believer knows what is morally required of them thanks to a divinely communicated sign, even though they lack knowledge of the sign's origins. Peoples illustrates the idea:

"Consider for example the possibility that God conveys the "sign" to people regarding some act (let's pick murder) via a proper function of the human conscience. Nobody needs to know what conscience is, how we got one, or that God uses it to ensure that we have some true beliefs in order for them to know, via conscience, that murder is wrong..."

(Peoples 2011)

In short, reasonable non-believers can know the content of moral commands without knowing their source.

There are at least two difficulties with this version of the argument. The first has to do with how Adams's communicative standard ought to be understood. To work, the third condition in this standard must state that knowledge of content is all that is required. But Adams's original wording suggests that knowledge of *form* is needed as well. After all, he says that you must understand the sign as 'conveying the intended command', suggesting that the agent to whom the command is directed must know that they are being commanded to do something. This is problematic because it is highly unlikely that reasonable nonbelievers have epistemic access to both the *form and the content* of the divine commands. Such a nonbeliever might have the phenomenological feeling that murder is wrong and that they ought not to do it, but may nevertheless have no sense that they are being *commanded* to refrain from it. Wielenberg puts the point nicely when he says that '[t]here is a difference between *commanding* someone to perform a certain act A and *causing*

someone to believe that he is morally obligated to perform act A.' (Wielenberg 2014, 79). All that is happening in Peoples's imagined case is that the reasonable nonbeliever is being caused to that murder is very bad; not that they are being caused to believe that they are being commanded not to murder. This suggests that insisting that the moral agent must have knowledge of the form and content is not going to be sufficient to rebut the epistemological objection.

It is probably for this reason that Peoples insists on a pure 'content not source' objection, going so far as to explicitly state that the reasonable nonbeliever 'does not even need to know that [what is being communicated to him] is a command, provided the command can be conveyed to him' (Peoples 2011). This seems to be a more sensible but here we run into a second major difficulty. Once you insist that knowledge of content alone is sufficient for the satisfaction of the epistemic condition you seem to undermine the criterion that the DCTist uses to distinguish genuine moral obligations from other obligation-like claims on our moral agency..The essence of DCT is that an obligation is *moral* only if it is commanded by a *being with the right kinds of properties* that can create truly moral obligations. But if a particular moral agent does not know that their conscience-delivered moral beliefs derive from such a being, it is very difficult to see how, on DCT, moral obligations are supposed to exist for that agent. Communication of the source would seem to be essential to successful communication of the command. The theist may insist that the obligations still exist in some objective, agent-independent sense, but then they are moving away from a communication-based theory of moral obligation, and ending up with a more general theological stateist view.⁹

We see the problem in the case of the modified version of the DCT defended by Adams. One part of the motivation for this modified DCT was to come up with a satisfactory response to the Euthyphro dilemma. This was achieved by insisting that God has the right kind of essence or nature to issue moral directives. Another part of the motivation for Adams's modified DCT was the need to distinguish between obligatory and supererogatory acts. Our

⁹ The problems with this move are discussed briefly in the final section.

conscience may often tell us what is morally good and what is morally bad; indeed, it may often go so far as to tell us that something is really really good or really really bad. But it's not clear that conscience alone is sufficient for knowledge of obligations, certainly for many hotly disputed moral issues. Am I obligated give large chunks of my money to charities in the developing world, just as I would be obligated to rescue a drowning child, or is that merely supererogatory? Adams's point is that we need a being with the right mix of properties and the right kind of authority to help us tell the difference. If we don't know that the pangs of conscience we experience in relation to charitable giving are command-like signs from such a being, how are we to know what we are obligated to do?

Allied to this, and more importantly, there is the problem that (moral) command-like signs are cheap and easy-to-fake. Humans can create books claiming to contain God's commands but which are actually made up to satisfy their own selfish political needs; preachers can claim to channel God's commands in their sermons for their own nefarious purposes; more outlandishly (but consistent with certain religious worldviews) there could be demons or devils interfering with our moral consciences and telling us that certain things are right or wrong even when they are not. All of these cases could result in command-like signs being communicated to moral agents, but none of them result in the creation of moral obligations. Why not? Because in each case we would know or have reasonable doubt about the authority of their sources. Knowledge of the source is thus crucial if we are to distinguish moral commands from fakes and for those commands to be successfully communicated to us. In the end, then, it is very difficult to see how a proponent of the Adamsian-style DCT could insist that knowledge of content by itself is sufficient for the satisfaction of the epistemic condition. For them, it is the source of the commands that makes all the difference

Is there any way to avoid this problem? C. Stephen Evans tries valiantly to do so. In its initial steps, his argument is the same as that provided by Adams/Peoples. He agrees with Adams and Peoples that knowledge of the

content of God's commands is essential if moral obligations are to exist (Evans 2013, 37). Indeed, he admits that 'it is hard to see how an unknowable moral obligation could be an obligation at all' (Evans 2013, 38). He goes on to claim that conscience is a key source of moral knowledge and that conscience is open to all, believer and non-believer alike (2013, 41). Conscience, according to him, gives us an immediate intuitive knowledge of the rightness or wrongness of our actions, although he admits that the knowledge it provides may be 'highly fallible' (2013, 41). Where Evans differs from Adams/People is in offering a more precise diagnosis of where it is that the proponent of the epistemological objection goes wrong and why it is that knowledge of the source of commands is not necessary for moral knowledge.¹⁰ As he puts it, the proponent of DCT must:

“defend[] the claim that God’s commands can generate obligations even for those who do not recognize those commands as coming from God. To do this two distinctions must be made. First one must distinguish between a recognition of a moral obligation and the recognition of a moral obligation as a divine command. This distinction in turn requires that one distinguish a recognition of a moral obligation from an explanation of the existence of a moral obligation. The defender of DCT... must hold that it is possible for a reasonable non-believer to recognize a moral obligation without realizing that the obligation is in fact a divine command. This ignorance...may prevent the non-believer from being able to give an adequate explanation of the existence of the moral obligation, but there are many cases in which a person may reasonably accept the existence of something, but lack an adequate explanation of the existence of something.

(Evans 2013, 112-113)

There is a lot going on in this quoted passage, but the essence of it is this:

Evans's Epistemic Condition: On DCT, a RNB can have knowledge of a moral obligation M provided that they can recognise that M is a moral obligation;

¹⁰ I am indebted to an anonymous reviewer for encouraging me to engage with this aspect of Evans's argument. Minor textual note: Evans refers to the epistemological objection as the 'promulgation objection'. The terminological difference is unimportant here.

they do not need to have an adequate explanation of why M is a moral obligation in order to have the requisite knowledge.

His argument then is that reasonable non-believer can *recognise* the existence of certain moral obligations (a) because their consciences will reveal to them that particular actions are right/wrong and (b) because they are familiar with other kinds of obligations from their social lives (like Adams, Evans develops his DCT by analogy with the obligations arising in social life). They do not need a deeper explanation of why it is that the particular obligation is a moral obligation.

Evans uses a lengthy analogy to defend this point. I discuss that analogy in the next section because it raises additional problems. We don't need it to understand the flaw in Evans's argument. The flaw stems from Evans's attempt to distinguish the explanation of M from knowledge of M in the context of DCT. The attempted separation is understandable: DCT is, after all, usually proffered as an explanatory grounding for moral obligations that we already believe ourselves to have. It states that the reason why we have those moral obligations is because God commanded them to be moral requirements. But the proponent of the epistemological objection is highlighting is that the attempted separation of explanation and recognition is not possible in the case of DCT. There is an intimate connection between knowledge of moral obligations and the kind of explanation of the existence of moral obligations that DCT purports to provide.

It is easiest to make this point by way of analogy. Consider the fact that many people prefer to eat organic food. But organic food is not simply any old food, it is food that is produced by a particular causal mechanism, namely: an organic farming method (i.e. one that is free from chemical fertilisers, pesticides and so on). Consequently, for people who prefer organic food, it is essential that they know that the food they consume is produced by organic farming methods. In many countries, official certifications and food labels are used to provide consumers with knowledge of organic farming methods. Provided the consumer has faith in the food labelling and certification system, they can be reasonably confident that any food they purchase with the relevant labels and certifications

is indeed organic in nature. If such a consumer wanders down a food aisle in a supermarket and sees food on the shelves, they might be able to recognise that food *qua* food by its general appearance, but they cannot recognise it *qua* organic food, until they see the relevant label certifying the process by which it is produced. In this scenario, knowledge of the process that produced the food *explains* why the food counts as organic food, but knowledge of that process is essential for the consumer if they are to recognise the food *qua* organic food. So in this case, the ability to recognise X *qua* organic food is not separable or distinguishable from an explanation of why X counts as organic food.¹¹

The analogy with the case of DCT and moral obligations should be straightforward, but let me spell it out. The essential claim of DCT is that a moral obligation is not simply any old obligation (or obligation-like sign/command) – we have lots of those in society as it is; it is an obligation that is produced by the right causal mechanism (i.e. by a command issued by God). Consequently, on DCT, for people who want to know their *moral* obligations, it is not enough that they recognise M as an obligation-like sign/command, they must know that it was produced by the right causal mechanism. Without that knowledge they may be able to recognise M *qua* obligation but they will not be able to recognise M *qua* moral obligation. Now, it is true that the divine causal origin is also what explains why M is a moral obligation, but that is exactly the point that the proponent of the epistemological objection is making: explanation of moral obligations is not distinguishable from recognition of moral obligations, given the way in which DCT purports to explain the existence of moral obligations. That's the essence of the epistemological objection: that given the kind of metaethical grounding that DCT provides for moral obligations it is not possible to separate explanation from knowledge. Consequently, the content-not-source objection does not work.

¹¹ There is of course one crucial difference between the cases, but this merely strengthens the epistemological objection to DCT. Organic food will continue to be organic food even if it is mislabeled. But on DCT, a moral obligation will not continue to be a moral obligation unless it is successfully communicated and, according to what is being argued here, it is not successfully communicated until knowledge of its source is provided.

4.2 – The Probability-Possibility Argument

Is there another way to defend DCT from the epistemological objection? Perhaps. One thing that the preceding analysis assumes is that the RNB completely (and reasonably) denies the existence of God and so they are epistemically closed to the possibility of moral obligations emanating from God. But, of course, that might not be a fair representation of their phenomenology. The RNB is likely to be (or have been) open to the possibility of there being a God, and thus to harbour some potential doubt as to the source of their moral beliefs. Is there any way in which this doubt can be leveraged into a defence of the DCT? Evans (2013) hints at this in his writings, as do others.¹² As far as I am aware, however, no one has offered an extended defence and analysis of this suggestion. I thus have two goals in this section. The first is to present this argument – which I call the ‘probability-possibility’ argument – in what I believe to be its most persuasive form; the second is to critique it.

We can understand the gist of the argument by considering an extended thought experiment from Evans. This is presented as part of his defence of the ‘content-not-source’ objection but its wording and structure it is not fully analogous with a pure version of that objection. This is because, as you will see in the quoted passage below, it focuses on a scenario that combines uncertainty as to the source of a command with knowledge of the content of a command, and not on a scenario that just involves knowledge of a command. The thought experiment is worth quoting in full:

“Suppose I am hiking in a remote region on the border between Iraq and Iran. I become lost and I am not sure exactly what country I am in. I suddenly see a sign, which (translated) reads as follows: “You Must Not Leave This Path”. As I walk further, I see loudspeakers, and from them I hear further instructions: “Leaving the path is strictly forbidden.” In such a situation it would be reasonable for me to form

¹² One of the reviewers to this paper suggested that when considering the alleged status of a moral obligation under DCT “[t]he RNB can still examine the commands and properties of the alleged God of DCT, and can even specifically consider the claim that God has employed conscience or natural law to communicate his moral will, as well as verbal commands.” This suggests a scenario in which the RNB is at least open to God as a possible source. Is being open to this possibility enough to generate knowledge of an obligation? That’s the question pursued in the text.

a belief that I have an obligation to stay on the path, even if I do not know the source of the commands. For all I know¹³ the commands may come from the government of Iraq or the government of Iran, or perhaps from some regional arm of government, or even from a private landowner whose property I am on. In such a situation I might reasonably believe that the commands communicated to me create obligations for me, even if I do not know for sure who gave the commands.”

(Evans 2013, 113)

There are three essential features to this thought experiment: (i) there is a sign communicating a command; (ii) there is uncertainty as to the true source of that sign; and (iii) there is the *probability/possibility* that it came from someone with the right kind of authority. The conjunction of these three features is thought to provide knowledge that is sufficient to generate an obligation to obey. The idea then is that the situation faced by the reasonable non-believer is similar. Hence there might be knowledge sufficient to generate an obligation to obey in their case too.

There are two ways in which to interpret this argument. One, which I shall call the *pragmatic enrichment interpretation*, seems to be implied by Wielenberg in his response to Evans’s argument (Wielenberg 2014). Pragmatic enrichment is the well-known linguistic phenomenon whereby the meaning of an utterance is enriched by the pragmatic context in which it was uttered.¹⁴ So, for example, I might say to you ‘Close all the doors’. In saying this I am issuing a command, but what does it require you to do? I surely cannot mean for you to close all the doors in the world. That would be extreme, but that interpretation is consistent with the literal meaning of the utterance. As it happens, however, I say this to you while we are sitting in a room with two open doors. This is the pragmatic context in which the command was issued. Given that pragmatic context, you rightly interpret the command to mean that you should close the

¹³ This phrasing seems crucial. It suggests that, within the thought experiment, there is a live epistemic possibility that the commands come from some particular source.

¹⁴ Linguistic philosophers distinguish between *semantics*, which has to do with the general conventional meanings, and *pragmatics*, which has to do with the token-specific properties of an utterance.

two open doors. The meaning has thus been enriched by the context. Pragmatic enrichment of this sort is a routine feature of ordinary conversations and communications. Speakers and listeners share a set of assumptions and understandings about how the world works. That set of assumptions always influences what we take to a reasonable or plausible interpretation of that utterance. The idea would be that the receiver of any command-like sign makes certain background assumptions about the reasonable way in which to interpret those commands, and those assumptions might include the possibility of them emanating from God.

There is a problem with this interpretation of the argument. Evans's original thought experiment only works exploits shared background knowledge about the pragmatic context in which the obligation-imposing sign is produced. We all know that the Iran-Iraq border is subject to political-legal control (i.e. that there are authorities, be they public or private, who control the land in that region). We also know, given the contested nature of that landscape, that these political-legal authorities may have good reason for telling people to stay on the path. These shared assumptions make it plausible for us to interpret the sign in an obligation-imposing way. The problem is that these shared assumptions about pragmatic context evaporate in the case of the reasonable nonbeliever who is trying to interpret allegedly obligation-imposing signs. The reasonable nonbeliever does not think that his or her conscience is the product of a divine will, or that the local preacher has a hotline to God, or that the Bible is a divinely inspired text. They do not share those assumptions with the believers. Consequently, it is difficult to see how they could understand those signs as commands coming from a potentially legitimate authority (Wielenberg 2014, 79).

There is, however, another way in which to understand the argument. Wielenberg's claim that the reasonable nonbeliever won't interpret the sign as being a command only really works in the case of moral conscience or intuition. In the case of signs communicated via a holy text or through the medium of a prophet/preacher, the command-like form of the sign may be obvious enough. In those cases, the pragmatic enrichment interpretation misses the crucial role that

moral uncertainty seems to play in Evans's thought experiment. We can pay due homage to that role by considering the argument in light of the burgeoning literature on the topic of moral uncertainty (Lockhart 2000; Sepielli 2008; Guererro 2007; Moller 2011; Weatherson 2014; Harman 2015).

We need some background first. The current debate about moral uncertainty was kickstarted by a number of philosophers (Lockhardt 2000; Guererro 2007; Moller 2011) claiming that uncertainty as to the moral status of an act should, on at least some occasions, influence our decision-making with respect to that act. It has long been accepted that uncertainty with respect to our factual beliefs should influence our moral decision-making; the moral uncertainty merely claimed that the same should be true with respect to our moral beliefs. Analogies and stories can illustrate their point. Consider the following two examples (adapted from Weatherson 2014):

Poisoning: Alice is caring for her sick friend Bob. She has a bottle containing what she believes to be useful but non-essential medication. But she is only 90% certain of this. She knows that there is a 10% risk that the bottle contains a lethal poison that would instantly kill Bob.

Vegetarianism: Rotimi is invited to a meal at her friend's house. Steak is on the menu, but there is also a vegetarian option. She is 90%¹⁵ certain that eating meat is morally permissible, but she accepts that there is a 10% risk that eating meat is gravely morally wrong (nearly as bad as killing and eating an innocent human being).

The first case involves factual uncertainty, *i.e.* uncertainty with respect to the contents of the bottle. There is no uncertainty with respect to the morality of lethally poisoning one's friend. Most people accept that the factual uncertainty in this case should influence the outcome: given that the medication is useful but non-essential, there is no way Alice should run the risk of lethally poisoning Bob.

¹⁵ Obviously these figures are misleadingly precise. No one could accurately estimate the truth of their moral beliefs like this. In reality, we would have to estimate subjective probability ranges. The precise figures are used for illustrative purposes only.

Moral uncertainty argues that similar reasoning applies to Rotimi's case. But in that case the uncertainty is directly related to moral beliefs, not factual ones.¹⁶ Rotimi is not sure about the ethical status of eating meat. It could be permissible but it could also be gravely morally wrong. In that context, the uncertainty will argue that she ought not to run the moral risk of doing something that could be gravely morally wrong. She should take the less risky option.

The Rotimi example highlights an argumentative structure that is common throughout the literature on moral uncertainty, something we can call the *risk asymmetry argument* (Weatherson 2014; Moller 2011). The gist of the argument is that when an agent is faced with two or more options, one of which has a non-negligible probability of being gravely morally wrong, the other of which is pretty certain to be morally permissible, they ought not take the potentially gravely morally wrong option. If this sounds somewhat similar to the argumentative structure underlying Pascal's Wager, then be reassured: it is effectively a moral version of it.

Evans's thought experiment can be reinterpreted in light of the concepts and argumentative structures used in the debate about moral uncertainty. The thought experiment focuses on a scenario in which (a) there is some uncertainty regarding the moral status of a sign (i.e. whether it is obligation-imposing) and (b) where uncertainties with respect to the moral status of the sign are sufficiently asymmetrical to generate a moral obligation to follow the sign. In the case of the sign on the Iran-Iraq border, the uncertainty stems from uncertainty with respect to the origin or source of the sign. One might argue that this isn't really moral uncertainty at all since the truth regarding the origin or source of the sign is factual in nature, not moral. But that isn't quite right. Although the distinction between factual and moral beliefs is fuzzy, the whole point of the Iran-Iraq thought experiment is that the source of a command determines its normative status. If it comes from a legitimate authority, it is obligation-imposing; if it does not come from a legitimate authority, it is not. The concept of

¹⁶ Of course, some people will argue that moral beliefs are factual since moral propositions are capable of being true or false. I don't dispute this and although I think the distinction between the moral and the non-moral is fuzzy, I think the distinction is clear enough for this discussion to make sense.

legitimate authority is itself highly moralised, so it is difficult to see this as anything other than a type of moral uncertainty.¹⁷ This applies *a fortiori* to cases in which the uncertainty with respect to the origin concerns whether or not God is the source of the sign since God functions as the legitimate authority for *moral* obligations.

The thought experiment also structures the risk asymmetries in such a way that the obligation-imposing interpretation of the sign is the most compelling. It is, after all, based in a disputed region between two historically belligerent nations. In that context, it seems highly likely that the sign comes from someone with the right kind of authority, and that one would be running a grave moral/legal risk by not following the command.¹⁸ But this is where the thought experiment unravels if it attempts to provide broader support for DCT. The thought experiment seems to involve a scenario in which we are 90% sure that the command comes from a legitimate authority. It is easy enough to see how that degree of confidence might suffice to generate belief in the obligation. The tougher case (and the one that is presumably more analogous with the case of the reasonable nonbeliever) would be where the probabilities are stacked the other way around, i.e. we are 90% sure the signs do not come from the right source, but think there is a 10% chance that they do. Would that be sufficient to generate belief in an obligation? It might be, though it would presumably depend on the potential moral status of disobeying the command. If walking off the path confers no obvious moral benefits, then we might agree that the non-negligible risk that the sign is obligation-imposing is sufficient to generate belief in that obligation. But if there are moral benefits from disobedience, the situation might be rather different. The risk asymmetry may work the other way.

¹⁷ It is certainly no different than the kind of uncertainty regarding whether a foetus is a *person* or a non-human animal an entity with the right not to be killed for our consumption. Uncertainty with respect to those kinds of beliefs is accepted as an example of moral uncertainty in the existing literature (Moller 2011). That said, I readily acknowledge that some critics of moral uncertainty think there are subtle distinctions between moral and non-moral facts that may make a crucial difference in this debate (e.g. Harman 2015 and Weatherston 2014). Since I ultimately appeal to the views of these critics in the response to Evans, I don't believe that my view is in tension with theirs.

¹⁸ The thought experiment also layers prudential risk on top of legal/moral risk, with the risk asymmetries working in the same way. This makes it doubly compelling.

Consider how the reasoning would apply to the case of the reasonable non-believer. The idea, presumably, would be that the nonbeliever is also confronted with a sign (e.g. a pang of conscience, or an utterance from a holy book or a preacher etc); and that they are uncertain as to the true source of that sign. They are reasonably confident (say 90%) that it does not come from a being with legitimate authority to create commands (i.e. God) but they accept that there is a non-negligible probability that it does (say 10%). The claim then would probability of the sign coming from the right source is sufficient for the reasonable nonbeliever to know that they are under an obligation and that this in turn is sufficient to satisfy the epistemic condition of DCT.¹⁹

There are four problems with this as an attempt to ward off the epistemological objection. The first is theological in nature. I think it is questionable whether this sort of uncertainty-based view of communication is compatible with the properties of God as traditionally conceived. One would presume that God has the power to communicate clearly to us that he is the source of a particular sign; just as our local legislature has the power to do the same. The fact that he does not, for at least some people, and that he consequently has to leverage uncertainty in the manner just outlined seems surprising. The theist owes us some justification/explanation for why God has to adopt such an imperfect system of communication. Theists will no doubt oblige by offering explanations that are broadly similar to those provided in response to the problem of evil or the problem of divine hiddenness (e.g. Evans 2013, 114-115). But there is one problem with all such attempts to explain away the uncertainty. Remember what is at stake if the epistemic condition is not satisfied? Moral obligations cease to exist for a certain segment of the population. The theist needs to explain why God is entitled to run the risk of excluding people from the realm of moral duty. For instance, a typical theodical response to the problem of evil is to argue that God cannot intervene to prevent evil because it would compromise our ability to develop and acquire true moral agency, where this agency consists in the ability to recognise right or wrong, freely

¹⁹ This interpretation of the argument subsumes the preceding pragmatic enrichment interpretation because what is now being alleged is that uncertainty with respect to the pragmatic context is sufficient for successful communication of the command.

decide one's actions, and take moral responsibility for the outcomes. A theist might be tempted to offer a similar response in this instance, arguing that God doesn't communicate his will to us perfectly in order to give us the freedom to develop our own moral agency. But this theodicy is not going to work in the present context. If the divine communication fails, it isn't simply the case that the non-believer is left to develop their own moral agency without the guiding hand of God; it is that moral obligations cease to exist for this agent and hence there is no moral agency for them to develop. This would actually defeat the rationale for the alleged theodicy.

The second reason I think this is not a plausible defence of the DCT is that the impact of moral uncertainty on our obligations is itself hotly disputed. Although several authors think that moral uncertainties and risk asymmetries can combine to alter the content of our obligations (Moller 2011; Lockhart 2000; Guerrero 2007), there are others who have challenged this belief (Weatherson 2014; Harman 2015). Weatherson for instance argues that one problem with taking moral uncertainties onboard in one's practical reasoning is that it results in a type of moral fetishism: one acts for the moral good *whatever ever that may be*, not for the reasons that typically motivate moral action (e.g. because an act does harm). In essence, the uncertaintist acts morally *de dicto* rather than *de re*. Harman agrees with this *de dicto/de re* distinction and adds the observation that the uncertaintist position rests on the mistaken belief that false moral beliefs can exculpate. We cannot adjudicate on these arguments here, but if there is indeed something suspicious about the uncertaintist's position it casts into doubt this rendering of the argument.

The third reason for doubting this defence of the DCT is that if you accept that non-negligible probabilities can alter the status of a 'moral' sign, you start down a slippery slope. There are plenty of other non-negligible probabilities that could lead to a very different interpretation of the allegedly-divine signs. For a non-believer, the probability of God's existence may be roughly on a par with the existence of other supernatural beings with similar powers (e.g. Evil God, the devil, demons, angels etc.). When confronted with a putatively divine sign (such

as a pang of conscience) they may recognise the non-negligible risk that it emanates from God, alongside a similar non-negligible risk that it emanates from another supernatural being like the devil or a demon. But, of course, if the sign comes from one of these other sources, it wouldn't be obligation-imposing. Quite the contrary: it might be the devil trying to trick the person into doing something genuinely malicious or evil. In that case, they would be running a grave moral risk by following the sign. The risk asymmetries may thus counsel against viewing the sign as obligation-imposing. In short, risk-based analyses of this sort can cut both ways. I suspect it will most often cut against the theistic interpretation in the case of the typical reasonable non-believer. In fact it is even worse than that because the same problem applies to a significant group of *believers*. Many believers think that the devil's existence is highly probable. For them – more so than the reasonable non-believer – the risk of misinterpreting a sign is far higher. They may find themselves paralysed in cases where the source of the sign is uncertain. The same may be true for those who embrace the skeptical theist position. Many have argued that skeptical theism undermines our epistemic confidence in allegedly divine signs to a similarly paralysing effect (Maitzen 2007; Wielenberg 2010; Law 2014). This suggests, incidentally, that the epistemological objection could have a far broader scope than originally assumed. It could be a problem for believer and nonbeliever alike.

Finally, I think the response fails because for many reasonable non-believers the probability of God's existence may be truly negligible. Most defenders of the view that moral uncertainty can alter our understanding of our moral obligations accept that this only holds true when the risks in question are non-negligible (e.g. Moller 2011). So if the probability of God's existence is negligible, the argument will not work.

To briefly sum up, the probability/possibility argument has some interesting features, and is hinted at in the literature, but even after developing it in what I believe to be a charitable form, I find it unsatisfactory as a response to the epistemological objection.

5. Conclusion: Assessing the implications

The preceding discussion points to one conclusion: the epistemological objection – in the form presented – is plausible. Divine Command Theories really do incorporate an epistemic condition into their account of moral ontology; that condition fails to be satisfied for a sizeable class of nonbelievers; and the result is that moral obligations fail to exist for that class of nonbelievers. At this point we return to the question: so what? Is this something that should bother the proponent of DCT? Or can they simply shrug off the objection? There are several possibilities.

If one is of a *universalist* persuasion when it comes to God and the moral law (i.e. believes that everyone is subject to the same moral law), then one should be disturbed by the argument. It is suggesting that a significant proportion of the population are outside the realm of God's moral law, able to act with impunity and without moral consequence. At the very least, this would seem to be a significant practical problem; and more likely a deep theological problem.

If one is persuaded by the notion that metaethical theories battle it out by earning more plausibility points, then one should also be disturbed by the argument. It suggests that the DCT misses one important mark because it has no way to account for the moral obligations of reasonable non-believers. This might mean that the believer should rest their hopes on an alternative theistic metaethical theory. This strikes me as being the most compelling interpretation of the argument since epistemic conditions would seem to be an important component in any moral theory.

That said, a *hard line* interpretation is also available. The theist could insist on the pure separation of epistemology and ontology. This would probably require a retreat to a more general theological stateist position, and not a command-based view based on successful communication. In that case, the theist could argue that their account of moral ontology is internally consistent (even if

it fails to account for moral knowledge),²⁰ and then simply note that the epistemological objection merely confirms their worst suspicions about an atheistic worldview: It turns out that even reasonable atheists really are beyond the moral pale.

This may, in turn, point the way back to the neglected part of the argument: the existence of reasonable nonbelievers. Although I think the existence of such nonbelievers is highly likely, I admit that I have a bias. The best hope for the proponent of the DCT might be to challenge the reality of reasonable nonbelievers.

Bibliography

Andrić, V. (2017). Objective Consequentialism and the Rationales of ‘ “Ought” Implies “Can” ’. *Ratio* 30(1): 72-87

Andrić, V. (2016). Is Objective Consequentialism Compatible with the Principle that “Ought” Implies “Can”? *Philosophia* 44: 63-77.

Adams, R (1999) *Finite and Infinite Goods*. Oxford: OUP

Baggett, D and Walls, J (2011) *Good God: The Theistic Foundations of Morality*. Oxford: OUP

Beyleveld, D (1991). *The Dialectical Necessity of Morality*. Chicago, IL: University of Chicago Press

Byrne, P (2011). Kant and the Moral Argument. In Jordan, J. (ed) *Key Thinkers in Philosophy of Religion*. London: Continuum.

Danaher, J. (2014). Skeptical theism and divine permission: A reply to Anderson. *International Journal for Philosophy of Religion*, 75(2): 101–118.

Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: OUP.

Evans, C. Stephen (2013) *God and Moral Obligation*. Oxford: OUP

Guerrero, A. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies* 136: 59-97

²⁰ To be clear, I do not accept here that DCT is otherwise internally coherent in its account of moral ontology. There are other objections one can make. I merely grant this possibility *arguendo*.

- Harman, E. (2015). The Irrelevance of Moral Uncertainty. In Shafer-Landau (ed) *Oxford Studies in Metaethics Volume 10*. Oxford: OUP
- Howard-Snyder, F. (1997). The Rejection of Objective Consequentialism. *Utilitas* 9(2): 241-248
- Joyce, R. (2002) *The Myth of Morality*. Cambridge: CUP
- Joyce, R. (2006) *The Evolution of Morality*. Cambridge, MA: MIT Press
- Kant, I. (2012). *Groundwork on the Metaphysics of Morals* (originally published 1785). Translated by Mary Gregor and Jens Timmerman. Cambridge: Cambridge University Press.
- Kant, I. (2015). *Critique of Practical Reason* (originally published 1788). Translated by Mary Gregor. Cambridge: Cambridge University Press.
- Lockhart, T. (2000). *Moral Uncertainty and Its Consequences*. Oxford: OUP.
- Maitzen, S. (2007). Skeptical Theism and God's Commands. *Sophia* 46: 237-243.
- Moller, D. (2011). Abortion and Moral Risk. *Philosophy* 86: 425-443
- Morrison, W (2009). The Moral Obligations of Reasonable Non-Believers: A Special Problem for Divine Command Metaethics. *International Journal for Philosophy of Religion* 65(1): 1 – 10.
- Murphy, M. (2012). Theological Voluntarism. *Stanford Encyclopedia of Philosophy*. Available at <http://plato.stanford.edu/entries/voluntarism-theological/>
- Peoples, G. (2011). The Epistemological Objection to Divine Command Ethics: Morrison on Reasonable Unbelievers and Moral Obligation. *Philosophia Christi* 13(2): 389-401
- Quinn, P. (2006). Divine Command Ethics: A Causal Theory. In Miller, C (ed) *Essays in the Philosophy of Religion*. Oxford: OUP.
- Sehon, S. (2010). The problem of evil: Skeptical theism leads to moral paralysis. *International Journal for the Philosophy of Religion*, 67: 67–80
- Sepielli, A. (2008). What to Do When You Don't Know What to Do. In Shafer-Landau (ed) *Oxford Studies in Metaethics*. Oxford: OUP.
- Smith, M (1994). *The Moral Problem*. Oxford: Wiley-Blackwell.
- Street, S. (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies* 109-166.

Weatherson, B (2014). Running Risks Morally. *Philosophical Studies* 167:141-163

Wielenberg, E. (2014). *Robust Ethics: The Metaphysics and Epistemology of Godless Normative Realism*. Oxford: OUP

Wielenberg, E. (2010). Sceptical theism and divine lies. *Religious Studies*, 46: 509-523.