

Theistic ethics and the Euthyphro dilemma

Richard Joyce

Penultimate draft of the article appearing in *Journal of Religious Ethics* 30 (2002) 49-75.

ABSTRACT:

It is widely believed that the Divine Command Theory is untenable due to the Euthyphro Dilemma. This article first examines the Platonic dialogue of that name, and shows that Socrates's reasoning is faulty. Second, the dilemma in the form in which many contemporary philosophers accept it is examined in detail, and this reasoning is also shown to be deficient. This is not to say, however, that the Divine Command Theory is true—merely that one popular argument for rejecting it is unsound. Finally, some brief thoughts are presented concerning where the real problems lie for the theory.

The label “Divine Command Theory” does not pick out any particular metaethical thesis, but rather a cluster of similar views. The version that I will begin with is the thesis that it is necessary that something is morally good if and only if God commands (wills, loves, approves of) that thing. I am content for present to leave open the relation that God bears to the thing in question, thus allowing that the name “Divine Command Theory” is possibly a misnomer. Nor does it matter terribly to this paper whether what is under analysis is moral *goodness*, as opposed to an action's being morally *required*.¹ Something about which I have made a decision is the presentation of the theory as a necessary biconditional (an “if and only if” claim), but we will encounter alternative readings below.

Philosophical discussion of the Divine Command Theory (DCT) has been rather peripheral to mainstream theoretical ethics in recent years. One reason for this is doubtlessly a widespread desire for a secular ethical theory, but another reason is that it is generally assumed that the DCT is subject to a fatal logical objection known as “the Euthyphro Dilemma”. In this paper I have nothing to say concerning the first reason for avoiding the DCT, but a great deal to say about the second reason. I will argue that the Euthyphro Dilemma represents *no threat* to the DCT.

The dilemma gets its name from Plato's dialogue of that title: therein Socrates apparently demolishes the position of his interlocutor, Euthyphro, who asserted something like a DCT. But the Euthyphro Dilemma, as it is now widely understood, bears little resemblance to anything presented by Socrates. My first task is to suggest just how Socrates's argument does proceed, and I will show that the reasoning is faulty. I will then outline the dilemma as it appears to be understood by many contemporary philosophers, and show that this reasoning is also unconvincing. The conclusion is that the DCT is destroyed neither by Socrates, nor by some argument bearing faint resemblance to the Socratic one. This is not to say that I think the DCT a true theory, for I do not. In the final section of this paper I will indicate where I think the troubles lie for the theory. However, my main objective in this article is to argue that the Euthyphro Dilemma is not the problem.

¹ Adams (1973, 1979, 1981, 1999) and Quinn (1978, 2000) have both preferred the deontological theory. My preference for “moral goodness” should not be taken as evidence that I think that they're mistaken. When I come to discuss Adams's views directly in the final section, I will shift terminology accordingly, but even there it doesn't really matter to my arguments.

1. What is Socrates's *Euthyphro* argument?

Euthyphro asserts that “holiness is what all the gods love, and unholiness is what they hate” (9E1-3). Socrates quickly presents him with a choice: “Do the gods love holiness because it is holy, or is it holy because they love it?” That Socrates presents Euthyphro with a disjunctive choice is not to say that he presents him with a dilemma. An argument by dilemma would proceed as follows: “Assume the first disjunct. This leads to unacceptable consequence X. Now assume the second disjunct. This leads to unacceptable consequence Y. Therefore whatever presupposition led us to the disjunction must be false.” The so-called “Euthyphro Dilemma” is often presented in just such a way (as we will see later), but it is not how Socrates proceeds. Rather, Socrates quickly gets Euthyphro to assent to one disjunct, then reasons from there. After a few leading questions, he is satisfied that he has overthrown Euthyphro's initial claim, and the other disjunct receives no further consideration.

It is not my intention to embark on textual exegesis of Plato's dialogue; rather, I will satisfy myself by appealing to an analysis of the text which I find the most charitable interpretation of Socrates's logic.² However, I'm afraid that the following is unlikely to be perspicuous read in isolation from the Platonic text (which is rather too long to quote here). As will become apparent, it's a fairly complicated argument, but a glance the text reveals that it must be.

The argument begins with Socrates getting Euthyphro's agreement to some preliminary statements (10A4-10C12):

- i) Something is a carried thing because it is carried; it is not the case that it is carried because it is a carried thing

- ii) Something is a led thing because someone leads it; it is not the case that someone leads it because it is a led thing

- iii) Something is a loved thing (is beloved) because someone loves it; it is not the case that someone loves it because it is beloved

Socrates is clearly trying to establish a general claim. The argument proper begins when he gets Euthyphro's assent to the pattern with the predicates “... is loved by the gods” and “... is holy”. Let me symbolise the sentence “Something is loved by the gods because it is holy” as:

1. L because H

In line with the pattern established by the earlier instances, Socrates gains acceptance of 1 and the denial of its converse:

² For capable analyses of the text, see Brown (1964) and Cohen (1971).

2. Not: (H because L)

And then he does the same with the proposition “Something is pleasing to the gods because it is loved by them” and its converse:

3. P because L

4. Not: (L because P)

Having won acquiescence on these four premises, Socrates thinks he’s almost home, declaring proudly that being pleasing to the gods cannot be the same property as being holy (10E2). Euthyphro is understandably confused about how such a conclusion follows, and in response Socrates presents two conditionals (10E14-11A4):

5. If (L because H) then (L because P)

6. If (P because L) then (H because L)

Quite how these conditionals are supposed to resolve matters is unclear. But we can progress if we treat each as the consequent of a conditional, the antecedent of which is Euthyphro’s claim. Let us revise them accordingly:

5*. If (H = P), then if (L because H) then (L because P)

6*. If (H = P), then if (P because L) then (H because L)

This pair of premises appears to be powered by a substitution principle (indeed, there seems no other way of making sense of them): the consequent of 5* involves replacing H with P, the consequent of 6* involves replacing P with H—both operations apparently permitted under the assumption that $H = P$. With these premises so understood, straightforward logic does the rest. The conjunction of premises 1 and 4 gives:

7. (L because H) and not (L because P)

while the conjunction of 2 and 3 gives:

8. (P because L) and not (H because L)

7 and 8 are the negations of the consequents of 5* and 6*, respectively. *Modus tollens* finishes things: the antecedents of 5* and 6* must be false. Thus we can conclude:

9. $H \neq P$

In other words, being holy is not the same thing as being pleasing to the gods.

I do not claim with any assurance that this is exactly how Socrates (Plato) took the argument to run, but I consider it to be a generous reading of some very dense passages, at least leaving him with a valid argument. In any case, my criticisms of Socrates do not depend on the above being precisely the argument. They focus on a substitution principle that is invoked, and on the initial assumption that Socrates coerces from Euthyphro. Even if one objected to the above interpretation, I do not think that it could be denied that Socrates's reasoning requires these two moves.

2. What Euthyphro should have said

Let us start with the intriguing claims (i)-(iii). Suppose John loves Mary (and pretend, for simplicity, that nobody else does). Mary is in a state of being loved; she is beloved. If we were to ask why John loves her, the answer "Because she is beloved" seems hopeless. Rather, we expect an answer along the lines of "Because of her intelligence, good looks, and sense of humor." In other words, we expect to hear about John's *reasons*. Compare the question "Why is she beloved?" Here the answer "Because John loves her" seems acceptable, but only if we understand the "why" as asking "in virtue of what?" In this case, the "why" does not ask for any agent's reasons.

That "why" questions allow of different interpretations is a thought that goes back at least to Aristotle. In his *Posterior Analytics* (2: ch. 11), he claims that the question "Why is the lantern shining?" may be answered "Because the surface of the lantern allows for the tiny light particles to pass through," or "To save us from stumbling." One answer is appropriate if we understand the question in "in virtue of what" terms; the other is appropriate if we are asked for our reasons for lighting the lantern.

Euthyphro should have pointed out that the relations he agreed to in (i)-(iii) are quite different from the relation he is claiming between being holy and being loved by the gods, which it is tempting to read as an assertion of identity between properties: the property of being holy *is* the property of being loved by the gods. For the time being, let us give in to this temptation, though we will back off from it later. If we try to construct correct analogues to (i)-(iii) using the predicates of Euthyphro's identity claim, then instead of 1 and 2 we would get something along the lines of:

- iv) Something is a holy thing because it is holy; it is not the case that it is holy because it is a holy thing

- v) Something is beloved by the gods because the gods love it; it is not the case that the gods love it because it is beloved by the gods

(v) seems as acceptable as (i)-(iii), but (iv) seems weird. The problem is that the adjectives "beloved," "carried," "led" are passive participles, all having obvious corresponding transitive verbs ("to love," "to carry," "to lead")—they indicate that some action has been performed to or upon the object in question. "Loved by the gods" is also like this, but "holy," along with a vast number of other adjectives that one can easily think of, is not. Therefore (i)-(iii) establish only a limited pattern, and

in agreeing to them Euthyphro in no way committed himself to 1 and 2. Indeed, the fact that Socrates attempts to “mix” predicates in 1 and 2—forcing a conclusion combining *being holy* and *being loved by the gods*, rather than, say, *being carried* and *being a carried thing*, or *being pushed* and *being a pushed thing*—should have aroused immediate suspicion.

Given Euthyphro’s initial claim, is he forced to admit *any* “because” statement? Suppose he responded thus: “I have made an identity claim, similar to ‘Jocasta is Oedipus’s mother.’ Am I required to say that she is Jocasta *because* she is Oedipus’s mother, or that she is Oedipus’s mother *because* she is Jocasta? I claim neither; I only say that Jocasta *is* Oedipus’s mother.” This can be our first defense on Euthyphro’s behalf: an identity claim need imply no “because” relation. But Euthyphro is made to overlook this.

Let us now turn attention to the substitution principle that Socrates invokes in order to produce 5* and 6*. This seemed the only way to make sense of the argument, but it is, nevertheless, an illegitimate principle, and 5* and 6* are false. We have already found the seeds of why this is so in Aristotle. A “because” claim is an explanation, generally in response to a “why” question. But a “why” question may have a number of quite different answers, each suitable to a different context of inquiry. What is a good explanation for one inquirer may not be for another. Consider the following call for an explanation: “Why did Oedipus blind himself?” To someone ignorant of the situation, the answer “Because he discovered that he’d slept with Jocasta” is no good, whereas the answer “Because he discovered that he’d slept with his mother” will be considered an adequate explanation. That is not to say that the latter is *the* correct explanation. To someone aware that Jocasta is Oedipus’s mother, the answer “Because he discovered that he’d slept with Jocasta” might serve as just as adequate an explanation.

In short, explanation contexts are opaque, which is to say: one cannot substitute a term with a co-referential term and be guaranteed that the resulting sentence will have the same truth value as the original. The standard example of an opaque context is the ascription of a belief: “Lois Lane believes that Clark Kent is coming to dinner” may be true, while (having substituted co-referential terms) “Lois Lane believes that Superman is coming to dinner” may be false. Similarly, from the fact that Lex Luther’s being in prison is *explained* by Superman’s having caught him it does not follow that his being in prison is explained by Clark Kent’s having caught him.

But (Socrates might object) Euthyphro’s initial claim is more than just an assertion that two terms are co-referential, it is a claim of *identity of properties*. What’s the difference? This is difficult to answer, since we do not know—not without indulging in philosophical speculation—just what *kind* of object a property name like “holiness” refers to. It cannot be just the name of the set of all holy things in the actual world, on pain of allowing that *having a heart* is the same property as *having a kidney* (to use Quine’s example³). Perhaps it is the name of the set of all holy things across all possible worlds.⁴ But this is not a question that Euthyphro

³ See Quine (1959, 204) and (1981, chapters 12 and 20).

⁴ In this paper I often use the term “possible world” as a useful device to refer to ways that things could have been. There is a possible world (quite a “close” one) at which Al Gore won the 2000 election; there is a possible world (a more distant one) at which Julius Caesar did not conquer Britain;

really needs to answer. He can simply say “Well, *whatever* kind of thing properties turn out to be, a single property could have two names, and when I say that holiness is being loved by the gods, I’m pointing out that two names are co-referential.”

One might be tempted to think that it makes a difference whether the alleged co-referentiality of these two names is supposed to be *a priori* or *a posteriori*. Someone could allow that explanation contexts are opaque when it comes to *a posteriori* equivalents (“Superman” and “Clark Kent”, for example), but insist on transparency when it comes to the *a priori*. This, however, turns out to be a red herring. Let me purloin an example of Elliott Sober’s for my own purposes (Sober 1982).

Imagine a machine that sorts shapes on a conveyor belt. As the shapes pass underneath, it determines whether each one has three sides, and then, let’s say, drops the three-sided shapes into a box. Suppose we ask, of a particular shape, X, “Why did the machine drop X in the box?” The natural answer would be “Because X has three sides (that is, is a trilateral).” The answer “Because X has three angles (that is, is a triangle)” is not nearly so adequate an explanation. It may not even be a true explanation at all. This is even more obvious if we are asking for an explanation of *how* the machine selected X. The answer “By determining that X is a trilateral” is correct; “By determining that X is a triangle” is simply false.

The conclusion is that explanation contexts are opaque for the substitution even of *a priori* equivalent terms. I have not claimed, however, that “trilaterality” and “triangularity” refer to the same property, even though clearly they are satisfied by the same objects across all possible worlds. The example seems to show (as is Sober’s intention) that property names must refer to something other than sets. But, as I said, this is not Euthyphro’s problem, and nor is it ours. If the worst comes to the worst, he can always retreat from making a strong claim of property identity, and instead claim the seemingly more modest necessary biconditional:

Necessarily: for any x , x is holy if and only if x is loved by the gods

As we have seen, Euthyphro can even hold this to be *a priori* true if he wishes (though he is not forced to), and Socrates has no argument to show that it is false.

We can sum up Euthyphro’s defense quickly. Suppose he starts out claiming the necessary biconditional just given. He need not admit that any “because” relation follows. And even if “because” relations do seem reasonable, there is no reason why he must choose one over the other. Perhaps relative to one context saying “X because Y” is correct; relative to another it is acceptable to claim “Y because X.” (See Van Fraassen 1980, 130-134.) Yet Socrates forces him, through appeal to dubious analogies, to assent to “Something is loved by the gods because it is holy; it is not the case that it is holy because it is loved by the gods.” But even with this admission Euthyphro’s thesis remains perfectly defensible, for Socrates’s argument hinges on the permissibility of substituting co-referential terms within an opaque context, which is exactly what you’re not allowed to do with an opaque context.

3. The Euthyphro Dilemma

there is no possible world containing square circles. The term is not intended to have any metaphysical implications.

Surprisingly, the “Euthyphro Dilemma” gets referred to frequently without any of the above arguments being entered into. Either people do not bother with Plato’s text, or they find a superior argument suggested by it. (Charitably, we’ll assume the latter.) The argument is presented as a dilemma. Pseudo-Socrates argues roughly as follows:

“You say, my dear Euthyphro, that necessarily: for any x , x is holy if and only if x is loved by the gods. So either something is holy because the gods love it, or the gods love it because it is holy. But either way leads to unpalatable consequences, and so you must give up your initial claim.”

We can see straight off that this reasoning reproduces the same fault as the real Socrates’s arguments: biconditionals—even necessary, *a priori* ones—do not need to imply any “because” claims holding between their relata. But I’ll put this important point aside in order to look at the rest of the argument. What are the alleged “unpalatable consequences” that each horn of the dilemma delivers? Let us consider the horns in turn, and while we’re at it let’s up-date the proceedings so as to consider a singular god and the predicate “... is morally good”.

Suppose that we assert that for any x , God loves x because x is morally good. The reputed problem with this is that it makes moral goodness independent of God. If the moral goodness of X is the reason for God’s loving X , then that goodness must be “prior” to God’s loving it, and must therefore be understandable in terms that do not invoke God’s love. This seems to undermine any DCT.

Suppose instead that we assert that for any x , x is morally good because God loves x . There are reputed to be two problems with this horn. First, statements that one might expect a proponent of a DCT to be committed to—statements like “God is good” and “God’s actions are good”—are rendered tautologous and empty. The second problem is that it makes the extension of “... is morally good” modally vulnerable—which is to say, what is supposed to be necessarily true turns out to be only contingently true. Had God loved killing babies instead of caring for them, then killing babies would have been good. But (the objection goes) killing babies is necessarily evil.

I will call these three problems, respectively, the “Independence Problem,” the “Emptiness Problem,” and the “Modal Vulnerability Problem.” The purpose of the following section is to argue that the defender of a DCT need not be bothered by any of them. However, it should be remembered that blunting only *one* horn will suffice.

4. Blunting the horns of the dilemma

The Independence Problem:

The assumption is that for any x , God loves x because x is morally good. What makes this seem like it renders goodness independent of God’s love is an almost overwhelming temptation to read the “because” as referring to *God’s reasons* for loving x . If my reason for loving hot chocolate is its warm sweetness, then its warm

sweetness must be a property that it has prior to my loving it. I perceive the hot chocolate, detect its warm sweetness, and love it for this attribute. Warm sweetness is thus a property that could be understood independently of my responses.

But we have already seen that there are other senses of “because” to which we should pay attention. We might say that John is a bachelor because he is an unmarried man. This isn’t his *reason* for being a bachelor. Perhaps his reason for being a bachelor is that he likes to eat TV dinners and watch football all day long. When we ask “Why is John a bachelor?” it is normal to interpret this as asking after his reasons, but it is not a mandatory reading. The “why?” may denote “in virtue of what?”—the questioner may be unclear on the conditions for bachelorhood, and be asking for a clarification of the concept. Suppose we start out making a reasonable necessary biconditional claim: “For any x , x is a bachelor if and only if x is an unmarried man.” We might just as easily choose to put this as a “because” claim: “For any x , x is a bachelor because x is an unmarried man.” Nothing wrong with that; we simply have to be clear that we’re not talking about any person’s *reasons* for being a bachelor.

Holding this in mind, consider the claim that for any x , God loves x because x is morally good. The advocate of a DCT can simply maintain that this “because” denotes an “in virtue of” relation. As to the question of God’s *reasons* for loving something—who knows? It has sometimes been claimed that if God’s reasons for loving things is not their goodness, then His love is *arbitrary*—but this is wildly overstated.⁵ The defender of a DCT is committed merely to what one of God’s reasons is *not*, which is a long way from claiming that He has no good reasons whatsoever. (See Adams 1973, 331 ff, and Quinn 1978, 135 ff.) Perhaps (just to give a concrete example) God loves things that maximize happiness—His reasons for loving things are based on utilitarian calculations. Suppose there is some token action—let’s name it “ ϕ ”—that maximizes happiness. God comes along (so to speak), detects this feature, and accordingly comes to love ϕ . Since, under the assumption of the DCT, there is a necessary connection between God’s loving something and its moral goodness, we may well assert that it is loved by Him in virtue of its goodness. Thus we could make two “because” claims:

- i) ϕ has the property of being loved by God because ϕ is morally good
- ii) God loves ϕ because ϕ maximizes happiness

The difference in wording of the portion before the “because” is intended to signal that only the latter concerns God’s reasons. Now there is nothing to prevent us creating a third “because” claim:

- iii) ϕ is morally good because ϕ maximizes happiness

⁵ Ralph Cudworth wrote: “divers modern theologers do not only seriously, but zealously contend ..., that the arbitrary will and pleasure of God ... by its commands and prohibitions, is the first and only rule and measure [of moral value].” The result of this position, he argues, is that “all moral good and evil, just and unjust are mere arbitrary and factitious things, that are created wholly by will.” (Cudworth 1731/1995, 3: 529, 532.)

And this, it may be thought, is the nub of the Independence Problem—for now we have an understanding of moral goodness with no mention of God. But this thought is quite mistaken. For a start, the procedure required a substitution within an opaque context. This is not to say that (iii) is false, given (i) and (ii); only that it does not follow from them. Second, there is no reason to think of (ii) as a necessary truth. Even if we thought that God is a utilitarian, we would not need to think that He *necessarily* is. And even if we thought that ϕ actually maximizes happiness, there is no reason to think that it *necessarily* does. Thus (iii) would not be a necessary truth either. It would tell us something about the particular action ϕ , but it would not provide us with an understanding of *moral goodness*.

Compare the case of John the bachelor. We can make two claims, each involving a different kind of “because”:

- iv) John is a bachelor because he is an unmarried man
- v) John is a bachelor because he likes TV dinners every day

From these we could go on to say:

- vi) John is an unmarried man because he likes TV dinners every day

But there are two problems. First, since all three statements are kinds of explanation, we have illegitimately performed a substitution in getting to (vi) from (iv) and (v). Second, it would clearly be foolish to think that (vi) has given us some understanding of *being an unmarried man* that is independent of the concept *bachelor*. Having seen that this horn of the dilemma is not as menacing as it is made out to be, let us now turn to the other horn.

The Emptiness Problem:

Now the assumption is that for any x , x is morally good because God loves x . The advocate of a DCT is probably a little more comfortable with this “because” claim than the converse, since there is no temptation to read it as referring to God’s reasons. Rather, it reads naturally as an “in virtue of” claim—God’s love “makes something” good. It is worth noting that, for many, this “because” claim *is* the Divine Command Theory, and the necessary biconditional version that I began with is too weak—is, perhaps, too “symmetrical”—to capture the intended relation between God and goodness. (See, for example, Burch 1980; Clark 1982; Quinn 2000.) But how should we understand this dependence relation? If the “because” claim is taken to be a necessary truth, then it cannot be a causal relation. It must therefore be some kind of constitutive relation, at least implying that being loved by God is necessarily a sufficient condition for being morally good.

The first reputed problem with this view is that statements like “God is good” and “What God does is good” become tautologous, in the sense of being empty and uninformative. (See Shaftesbury 1711/1900, 1: 264; MacIntyre 1969, 33). One kind of response would be to accept this consequence. The DCT supporter might just cease to make such assertions. Alternatively, she might claim that when she says that

God is good she means “good” in some special sense, distinct from the moral goodness which, according to her, God’s love constitutes. This is a stable position, though it raises the question of what this special sense of goodness might be, if not moral. But let’s assume that (A) the theist does want to make these assertions, and (B) the goodness in question is the same as that which is mentioned in the formula for the DCT. Are the charges of emptiness well-founded?

So far I have lumped together different things that a DCT might claim: that moral goodness is constituted by God’s love, or His approval, or His commands, or His will. This blurring of distinctions hasn’t really mattered, but now it does. I will proceed by discussing six different possibilities. First we will imagine that God’s *approval* defines the good. (This, I think, is close enough to speaking of His *love* that the latter doesn’t require individual treatment.) Under this assumption we will see what sense can be made of the claim “God is good” and then what sense can be made of “What God does is good”. Then we will turn to the possibility that God’s *commands* define the good, and examine these two claims again. Finally we will do the same with the assumption that God’s *will* defines the good.

Assume that God’s approval defines the good. Then the sentence “God is good” would be the claim that God approves of Himself. Now we might suppose that God does approve of Himself, but it is not an empty tautology that it is so. The sentence “What God does is good” would mean that God approves of His own actions. Again, the theist probably thinks that God does approve of His own actions, but it is not a truism that He does. So neither sentence is trivially true under this assumption. Of course, questions about whether God approves of Himself and His actions might be considered theologically misguided—a misapplication of human traits onto a Divine Being. So be it: all this shows, however, is that the fan of the DCT will also hold statements like “God is good” to be similarly defective, in which case the problem under discussion will evaporate. (This point should be borne in mind during the following discussion, too.)

Assume that God’s commands define the good. Commands pertain to actions, so what would it mean, on this account, to claim that a *person* (or *being*) is good? An obvious answer is that the person acts (usually? always?) in accordance with those commands. So “God is good” would mean that God acts in accordance with His own commands. One might worry that God doesn’t direct commands at Himself at all, in which case the sentence would certainly not be trivially true for it would be *false*—and the advocate of DCT doesn’t want that. It is better for the proponent of DCT to conceive of God’s commands not as directed at particular agents, but of a general nature (“Don’t kill”, etc.), in which case we can speak of God acting in accordance with His commands. As before, we can assume that God *would* act in accordance with His commands, but it is not a trivial matter that He does, for a person can issue commands without following them. By the same reasoning, the sentence “What God does is good” is non-trivial. We can conceive of God failing to act in accordance with His commands, without this implying the falsity of the DCT.

Assume that God’s will defines the good. As before, presumably what it means for a *person* to be good, according to this theory, is that he or she acts in accordance with God’s will. Now usually it is useful to draw a distinction between what one wants (desires) and what one wills. I may desire to go to the Caribbean for a week’s

holiday, but if I know that I cannot then I will not *will* that I go there. An act of will is what is involved in the actual decision to go—it is what prompts the action of buying the tickets, etc. But if this is how we understand “will” then it makes little sense to speak of an agent acting according to another person’s will. We can (literally) will only ourselves. This being the case, when we speak of humans acting in accordance with God’s will, presumably this must mean “in accordance with God’s *desires*.”

The sentence “God is good” would thus mean that God acts in accordance with His desires. This may seem uninformative on the grounds that part of what we mean by “action” is “a desired behavior”.⁶ However, there are a variety of familiar phenomena that complicate the relation between an person’s desires and actions. An agent may suffer from weakness of will, taking an extra slice of chocolate cake, say, while judging that she really ought not. The action of taking extra cake wasn’t an *accident*, so she still desired its execution, but perhaps we should say that she did not all-things-considered desire it. This may be understood in terms of higher order states: one may desire the cake, but desire that one did not desire the cake. The higher order desire is often called one’s “better judgment”, or what one “values” as opposed to merely desires. (See, for example, Frankfurt 1971, and Lewis 1989.)

Suppose that Bill is torn by weakness of will concerning something that he wants Betty to do. He wants her to come to coffee with him, but feels guilty because he knows that she should study. His “better” judgment—that is, the desire that is endorsed by his higher order desires—is in favor of leaving her to her work. However, suffering from weakness of will he reaches for the phone. In such a situation, what would count as Betty acting *in accordance with Bill’s desires*? There is no answer without simply disambiguating: to come for coffee is to act in accordance with one of his basic desires, to stay home is to act in accordance with his higher order desires. Yet it does seem reasonable to say that if Betty wants to act in a way that respects Bill, in such a way that Bill will not have cause later to castigate himself, then she should act in accordance with his higher order desire.

Presumably God is not like Bill. Suffering from weakness of will is a characteristic of certain classical Greek gods, perhaps, but not the Judaeo-Christian one.⁷ Nevertheless, reflection on the phenomenon of weakness of will may lead adherents of the DCT to revise their theory. It is not really God’s *desires* that they mean to highlight, but His *values*. Of course, in actual terms it makes little difference—since theists may be confident that God’s desires and values are in perfect, eternal harmony—but it is best to be precise in formulating one’s theory.⁸ Furthermore, however confident theists may be about the actual God, it is not obvious that it is an *a priori* truism that God does not suffer from weakness of will. Suppose there is a possible world at which God is torn concerning what He wants humans to do. He wants us to ϕ , but He also wishes that He didn’t want us to ϕ . As

⁶ This is so, at least, according to a certain Humean view of human action, which it is fair to call “orthodoxy.” I am very much inclined to think the Humean correct, but since in this paper I appeal to the view only hypothetically, there is no need to defend it.

⁷ Or, if you prefer: “...but not the Judaic God and nor the Christian God.”

⁸ In the same way, one might propose: “Necessarily: any three-angled shape is a trilateral,” but it would be preferable to replace “trilateral” with “triangle.”

with Betty and Bill, there would simply be no single answer as to what actions count as acting in accordance with His desires. Since defenders of the DCT want their theory to apply to counterfactual situations as well as actual ones, then they will need to revise it in order to cope with this bizarre world. Without revision, if a person at this possible world ϕ ed, then this may be both morally good and morally bad at the same time (assuming, as seems reasonable, that “moral badness” is defined in terms of acting contrary to God’s desires). Again, the natural thing, I think, would be to revise the DCT in such a way that the good is associated with God’s “better” judgments—that is, with his *values*.

But once this revision is made it is easily seen that “God is good” and “What God does is good” are not empty tautologies. Yet again, theists may safely assume that God *does* act in accordance with His values (thus He and His actions are good), but it is no triviality that this should be so. We can imagine a possible world at which He fails to so act, while still maintaining this form of the DCT.

In this paper I haven’t wanted to favor one form of DCT over another, but recent paragraphs have described a slide from talk of God’s will, to His desires, and finally to His values (that is, His higher order desires). I have given my reasons for finding this development necessary. This is not to say, of course, that theists and advocates of the DCT should give up talk of our acting in accordance with God’s will and God’s desires—only that when we get down to a meticulous defense of the DCT then such ways of speaking need to be made precise. Once that precision is in place, then it can be seen that the Emptiness Problem can be deflected.

I admit that there may be some lingering discomfort in what I have said concerning the Emptiness Problem. It may be conceded that the sentence “God is good” is not a tautology on the grounds that it is not a tautology that God approves of Himself. However, when a theist claims that God is good, surely she is not merely saying that God approves of Himself! Granted. All this shows, however, is that the theist who supports the DCT is probably employing a special sense of appraisal in positively evaluating God—one that is distinct from that which might be used in claiming that, say, St. Francis was good. (See Clark 1982, 345; and 1987.) And if this is so, then the Emptiness Problem doesn’t arise.

The Modal Vulnerability Problem:

The assumption is, again, that for any x , x is morally good because God loves x . This horn of the dilemma is supposed to have a second prong: that if God is free, omnipotent, etc., then He might have chosen to love different things than He actually loves. Surely He could have loved murder and mayhem had He wanted to. But (apparently) we have the intuition that murder and mayhem are *necessarily* morally wrong, and therefore the DCT is in trouble.

An obvious (and popular⁹) solution is to deny the contingency of God’s loving. Perhaps certain details of His love are contingent (for example, He loves Francis for his virtue; but since Francis is autonomous there are possible worlds where he is

⁹ See, for example, Burch (1980), Wierenga (1983), Adams (1999), Quinn (2000).

vicious and God does not love him¹⁰), but the broad, general aspects of His love (in favor of virtue, against baby-killing, etc.) are necessary attributes.

One worry with this solution is that God's abilities seem constrained. Some theologians have sympathies with Luther's claim that "God is He for Whose will no cause or ground may be laid down as its rule and standard; for nothing is on a level with it or above it, but it is itself the rule for all things" (Luther 1525/1957, 209). Furthermore, had God no real choices before Him, wouldn't the believer's tendency to *praise* Him be undermined? How can something be praiseworthy when it cannot be other than it is? These questions lead into theological territory so dense that if we pursued them it is doubtful that our discussion would ever again see the light of day, so I propose that we leave the issue as merely "a possible avenue" for the defender of DCT.

In any case, it seems to me that the other avenue—that which rejects that murder and mayhem, etc., are necessarily morally wrong—is more promising. (This is a bullet that William of Ockham was unashamedly willing to bite.¹¹) We are supposed to have the intuition that certain actions are so bad, so morally appalling, that they are *necessarily* wrong. Certainly there is a knee-jerk response in favor of the thought, but careful consideration reveals that things are not straightforward.

Consider a really appalling crime. Suppose Jack killed Jill from no motive other than hatred and self-gain. Call this action " ϕ ". Let's say that ϕ occurred on September 1st, 1888, in London's East End, with a knife.

Suppose that there's a possible world—call it " W_1 "—at which the only way in which Jack can save the lives of hundreds of innocents is to kill Jill. Imagine that this is what he chooses to do (on September 1st, 1888, etc.). Perhaps in those circumstances his action was morally correct, or at least permissible. (Some will say "Yes", some will say "No".) Suppose that there's another possible world, W_2 , where Jack and Jill are friends walking down the street in the East End, shopping for cutlery. Jack, fooling around with a knife he has just bought, pretends to stab Jill but accidentally does exactly that. Now consider W_3 , at which God's voice booms down and commands Jack to kill Jill (with a knife, on September 1st, etc.), which he proceeds to do. Or consider the possible world, W_4 , where it is readily evident that anyone who is violently killed goes on to an eternal heavenly paradise, regardless of how they lived their lives. Jack so loves Jill that he decides to send her there (albeit rather suddenly), and she is grateful for it.

I raise these bizarre thoughts in order to address this question: Is Jack's killing of Jill at W_1 , W_2 , W_3 and W_4 *the same action* as his killing of her at the actual world? They all involved Jill's death brought about by Jack's actions at the same time and place with the same implement. An important issue is whether Jack's *intentions* are a necessary feature of the action when it comes to identifying that action in counterfactual situations. We simply do not know what the identity criteria for a

¹⁰ One might object that God loves everyone, even the sinner. But the advocate of DCT cannot be referring to *that* kind of love, on pain of allowing that everyone is morally good. Perhaps the follower of DCT will prefer to frame the theory in terms of God's *approval*: although God may continue to love the sinner, He surely does not approve.

¹¹ William of Ockham, 1494-96/1962: vol. 4 (*Super quattuor libros sententiarum*) II 19 O, III 12 AA; vol. 1 (*Opus nonaginta dierum*) c.95.

token action are. But without confidence in these criteria, we have no business being sure that some actions are *necessarily* wrong. Change aspects of circumstances, motivations, or consequences—albeit dramatically in some cases—and all but the most theoretically entrenched intuitions will shift. (For similar considerations, see Wierenga 1983, 393-396.)

The opponent of the DCT might be unmoved. “Take an actual morally bad action,” he or she might say, “and fix its consequences and intentions across possible worlds.” In other words, let us just *stipulate* that “ ϕ ” is going to name Jack’s killing of Jill *from no motive other than hatred and self-gain*. The Modal Vulnerability Problem remains, for there is (we are imagining) a possible world at which ϕ occurs and God loves ϕ —in other words, a world at which ϕ is morally good.

But just how troubling is this consequence? Start with an analogy. If I tell you that my brother might be a serial killer, you’d naturally be shocked and appalled. But suppose it turns out that I don’t mean that it’s particularly likely—in fact, I’m 100% certain that he’s nothing of the sort, and nor is he in the least tempted to become one—I just mean that, at some other possible world, he’s a serial killer. All I’m saying, in other words, is that being an upstanding citizen is not an *essential* property my brother. So understood, the claim that he might be a killer turns out to be pretty innocuous and uninteresting. The claim that God *might* love murder and mayhem is no more threatening than this. It doesn’t mean that one should start worrying; indeed, one can be perfectly confident that He *doesn’t* love anything of the sort, and nor is He going to change His mind. It’s just a claim about logical possibility.

The same point should be borne in mind when we consider the possibility of an action which we morally abhor being only contingently morally abhorrent. The wrongness of an act of brutality being only contingent need not undermine our 100% confidence that it *is* morally wrong. Nor does it mean that the circumstances under which it would cease to be bad are particularly likely to arise. The world at which God would change His opinion of this action might, after all, be a very distant and strange place. It is almost as if one’s admitting that the wrongness of an action is only contingent places one under suspicion, as if such a person is not so committed to the wrongness of the action as someone who insists that the action is *necessarily* wrong—as if a full-blooded moral judgment entailed a modal commitment. But this is absurd. My admission that my brother is only contingently a good citizen does not in any way detract from my confidence in him. My positive moral assessment of his character would not in any way be enhanced or reinforced were I (extravagantly) to insist that he is *essentially* good.

That we might have modal intuitions about morality at odds with the DCT is not something that need cause the supporter of that theory undue concern. Of course, any proposed analysis or explication of the troublesome concept *moral goodness* must fit with moral intuitions. If a theory had the consequence that genocide is morally permissible, then that would count powerfully against the theory. And many moral intuitions *are* of a modal nature—for example, in the “runaway trolley” thought experiments made famous by Philippa Foot, we ask ourselves what we think would

be morally permissible if such an unlikely situation were to arise.¹² However, in such thought experiments the domain of possibility is restricted. We do not, for instance, consult our moral intuitions about runaway trolleys and people stuck to tracks for worlds where violent death is known to lead to post-mortem paradise. We do not wonder about the permissibility of genocide at worlds where the laws of nature do not hold. It is open to the DCT adherent to make a similar claim. “Of course,” she might say, “it is logically possible that God would love murder, and so, *ex hypothesi*, it is logically possible that murder be morally good. But such a world would be one where fundamental attributes of the Being who created and sustains the universe are dramatically different.” (Note: by “fundamental” we do not mean “essential”.) For a serious theist, imagining a world at which God loves different things than He actually does, where He makes quite different choices, is tantamount to imagining a world with different laws of nature. Who knows what kind of moral intuitions we should have about such worlds? We probably shouldn’t have any.¹³

5. The real challenge for the Divine Command Theory

The previous sections have shown that the Divine Command Theorist can happily ignore the Euthyphro Dilemma (both Socrates’s original and the modern version). Doing so may be counted as a form of defense of the theory, but in fact, in the long run, I doubt that the theory is true. In this final section of the paper I will switch roles, and indicate where I believe the problems lie. I do not pretend, though, to put forward a watertight case on this occasion.

We can start by asking the advocate of a DCT what kind of truth he or she takes the theory to express. Is the biconditional supposed to be a conceptual claim about what English speakers *mean* when they say “morally good”, or is it meant to be an *a posteriori* truth—perhaps comparable to “Something is water if and only if it is H₂O”?

The first option looks highly implausible. We need to bear in mind that people who do not believe in God by and large employ moral terms as sensibly and adamantly as do theists. Now the DCT proponent *could* claim that theists and atheists (and agnostics) simply mean different things when they say “morally good”, but that would be a most unattractive line to take. According to such a view, if an atheist were to claim, say, “It is morally permissible to ignore those in need,” the theist could not sensibly respond “No, you’re mistaken—one is morally required to help.” The two would be speaking different sub-languages, and their words would not express a disagreement at all. This is a deeply counter-intuitive result. If we are to allow for the possibility of genuine moral disputes between atheists, agnostics, and

¹² Philippa Foot’s original article, plus a good portion of the industry that it spawned, are collected in Fischer and Ravizza (1992).

¹³ This should not be confused with my earlier claim that fans of the DCT will want their theory to apply even to a bizarre world at which God suffers from weakness of will. The problem with that scenario was that the “unrevised” version of the DCT led to *contradictions*: to a person’s action’s being morally good and morally bad at the same time. A theory that leads to contradictions is in trouble. By comparison, a theory that just doesn’t know what to say about a fantastic counterfactual situation may be revealing (perfectly reasonable) imprecision.

theists, then we must allow that they all *mean the same thing* when they speak of “moral goodness”.

Given this, it would be a brave DCT defender who claimed that an adamant atheist *means* to refer to God when he or she employs moral language. Imagine a whole community of adamant atheists, contentedly employing a full range of moral concepts in their evaluation of each other and their actions. Whether their moral language is misguided or sensible is not the point; all that matters is that they qualify as *having* a moral language. In order to grasp what these people *mean* by some moral term we would need to try and get at *why* and *in virtue of what* they apply the term to the things to which they apply it, and deny it to the things to which they deny it. We would need to consult their moral intuitions, practices, and intentions, and then attempt to systematize these data into a claim of the form “For any x , x is morally good [say] if and only if x is ...” What replaces the ellipsis might turn out to be quite simple—like “the option that maximizes utility” or “what a rational person would choose”—or may be disjunctive and messy. The important point, though, is that we, as theorists analysing their meaning, would be at the mercy of *their* intuitions and intentions.

Now we can see that the prospect of the DCT being a claim about the *meaning* of “moral goodness” is very far-fetched, for it would require that if we were to line up all the moral intuitions held by a group of people who have never dreamed of there being such things as gods, or have even soundly denied their existence—and noted that these intuitions are stated in purely secular terms—we might nonetheless discover that the best systematization of those intuitions introduces something they have never thought of, or perhaps even denied: God. It is, perhaps, not inconceivable that this could happen, but it would involve ascribing a massive amount of self-deception to the atheists. In fact, it might be interpreted as revealing that our “atheists” are not really atheists at all, since their belief in moral goodness would commit them to a belief in God. This I take to be an unacceptable result. I’m quite prepared to admit that either the atheists or the theists are massively mistaken about the nature of the world (indeed, one of them *must* be), but claims that all atheists or theists are *self-deluded*—that they do not believe what they take themselves to believe—have no place in serious argument.

If the DCT biconditional is not put forward as a truth about the concept *moral goodness*, then perhaps it is supposed to be a truth about the property for which the concept stands. A useful model here might be the truth “Water is H₂O”. This is not a claim about the *meaning* of “water”, otherwise we’d have to hold that ignorance of chemistry amounted to linguistic incompetence, and that speakers of more than a couple of hundred years ago meant something quite different by their term “water”. But that is clearly absurd; when Shakespeare wrote of “all the water in the rough rude sea”, he meant the same by “water” as we do. The truth of “Water is H₂O” is an *a posteriori* matter—we had to investigate the world in order to come upon it. Yet the orthodox view, ever since Saul Kripke’s *Naming and Necessity*, is that it is also a *necessary* truth (Kripke 1972). Kripke’s revelation of the necessary *a posteriori* opened up new vistas to modern moral realists. They could now hope to present a theory of the form “Something is morally good if and only if it is ...” that might be

necessarily true, yet without the pretence of its being a piece of conceptual analysis.¹⁴

One way of acquiring a necessary non-conceptual truth would be to start with the agreed-upon extension of some predicate (such as "...is water"), and then investigate the items in that extension, perhaps discovering some other previously unknown fact about them ("All the stuff we've been calling 'water', and only that stuff, has the molecular structure H₂O"). But this seems unlikely to work as a model for the DCT biconditional because of the widespread fragmentation and disagreement in our moral views. In the case of "water", everyone pretty much agreed on which things were made of water and which things were not, and so we could then embark on an empirical investigation of the stuff in question. But there is no analogous agreement over the extension of our moral predicates. Is euthanasia morally permissible? We disagree. Is George W. Bush a morally good man? We disagree. Should the Elgin Marbles be returned to Greece? We disagree. When the extension of the predicate is so widely contested, and when the disputes are as deeply entrenched as they appear to be, then the possibility recedes of discovering that all these items (and only these items) satisfy some other predicate as well. We simply don't have a stable group of things available for investigation. If the DCT biconditional could capture one person's use of the term "morally good", then it must fail to capture the use of anyone who disagrees with him over a great many moral matters.

But there is another way of attaining necessary non-conceptual truths. Rather than starting with the *extension* of the predicate "... is morally good", investigating the items therein, perhaps we should start with what the predicate *means*. This route is that taken by R.M. Adams (1973, 1979, 1981, 1999) in his defense of a deontological DCT. (I will change terms accordingly, discussing God's *commands* and moral *requirements*.) Starting with the *meaning* might sound like a return to the view already rejected, but I will explain why it isn't.

We've already seen that in order for there to be genuine moral disagreement—say, when one person thinks that returning the Elgin Marbles to Greece is morally required, and the other person thinks it is not—there must first be agreement over what the terms *mean*, otherwise they would not in fact be disagreeing at all. How might we find out what people mean? I have already sketched an answer. We would take their shared intuitions, the platitudes that they endorse about *moral requirements*, and we would gather them together and try to encapsulate them into as simple a statement as we could. By "platitude" here we mean something of which a person must be disposed to indicate acceptance in order to be granted minimal competence with the term in question. For example, whether or not you think that returning the Elgin Marbles is required, you think that morally required actions are desirable, you think that they are important, you think that they have an impact on human affairs, you think that virtuous persons are (*ceteris paribus*) inclined to perform them, and so on. (For a more substantial list of what I have in mind, see Smith 1994, 39-41). A person who used the term "moral requirement" but denied any of these platitudes might well be accused of not knowing what she was talking

¹⁴ For *a posteriori* moral realists, see, for example, Boyd (1988), Brink (1984), and Lycan (1986).

about. When two people argue over whether some action is morally required, they are arguing over whether the action satisfies these platitudes.

Earlier I argued that it is exceedingly unlikely that the best systemization of our moral platitudes—bearing in mind that “our” includes atheists, polytheists, agnostics—is going to produce the DCT biconditional. Perhaps it will prove impossible to come up with a systemization at all—perhaps all we’ll be left with is a long list of more or less vague platitudes about moral matters. But then we might ask ourselves “Is there some property in the world that satisfies, or comes reasonably close to satisfying, all these platitudes? If so, then that’s the property of *being morally required*.” Here’s where the proponent of the DCT could step in. “Yes,” she might claim, “the property of being commanded by God, if there is a God, comes close to satisfying all these platitudes—closer, at least, than any other property.” This would not be a claim about the *meaning* of the term “moral requirement”. The meaning would be captured in the platitudes—platitudes with which *ex hypothesi* atheists and theists alike will be satisfied. What the defender of the DCT is suggesting is that there is a kind of property that, if instantiated, would do a good job of satisfying these platitudes, and therefore if this property *is* instantiated we should consider it to *be* the property of *being moral required*. (In what follows, for brevity I will call the property that the adherent of the DCT is proposing—in this case, *being commanded by God*—“the DCT-property”.)

The atheist has several options in response. First, he could say “Yes, I now see that the DCT-property is the *only* property that will do a decent job of satisfying all our moral platitudes; but since I don’t believe in God, I now see that there is no such thing as moral requirements.” This atheist endorses the DCT, but at the price of embracing moral nihilism. Second, he could say “Yes, I now see that the DCT-property is the *best* property for satisfying all our moral platitudes; but since I don’t believe in God, I will try to come up with a secular property that also does a decent job. Granted, it may not do such a good job, but at least it will allow me to avoid moral nihilism while remaining an atheist.” Although this second option is a perfectly unobjectionable strategy, the theistic DCT advocate might be pleased with the concession, for if she can then establish to her own satisfaction that God exists (and issues commands), then she will have established to her own satisfaction that the DCT is true. In the dialectic of this debate, that counts as a victory for the supporter of the DCT, for she is not necessarily trying to persuade everyone that the theory is true, merely that it’s acceptable for her to hold it. Needless to say, the atheist will not grant the existence of God, but the important thing is that the terms of the debate will have shifted. The dispute *was* over the truth of the DCT, but now it is over the familiar question of God’s existence. The theistic proponent of the DCT could be well-satisfied with having collapsed two thorny problems into one.

It is the atheist’s third strategy that interests me. Here he denies that the DCT property *is* the best contender for the property that maximally satisfies our moral platitudes. If the arguments to this conclusion are sound, then the supporter of the DCT—the kind of DCT that we have been considering in the last couple of paragraphs—is defeated. In order to assess this route, we must turn to specific important moral platitudes and wonder whether the property *being commanded by*

God is likely to satisfy them. Here I don't pretend to develop an argument in detail, but highlight it as an avenue for further dispute.

Consider, first, the cluster of platitudes that express the *authority*, the *inescapability*, of moral prescriptions. If an action is morally forbidden, then you must not perform it even if you want to, even if you stand to gain from it, even if you can reliably evade the punishments that might as a consequence be visited upon you. Even if you *will* be punished as a result of your moral crimes, the "wrongness" of your actions hardly derives from the unpleasantness of the punishment. In other words, enshrined in our moral platitudes is the vague but firm idea that moral prescriptions are not the same as pieces of prudential advice. When we say to a Nazi "You ought not kill these innocent people", is it at all plausible that the strength of our moral "ought" is that the Nazi is harming *himself*? If that were so, what sense would there be in thinking that he deserves punishment—that is, more harm—for his crimes?

Our question, then, is whether the DCT-property is likely to provide persons with such *non-prudential practical reasons*. We've already seen that the fact that *God* is typically understood to be the distributor of great rewards and penalties will not provide Him with the appropriate kind of authority to underwrite the moral law. Nor does His attribute of having created us and the universe in any obvious way grant His commands authority over us. (Though we generally consider parents to have authority over their children, I submit that we do not think that this authority derives *per se* from the fact that the parents in a very real sense "created" their children.) One might instead advert to the characteristics of *omniscience* and *lovingness* that are generally attributed to *God*. Now it is true that if there is a person who loves me a great deal, and who knows everything there is to know about me (and the consequences of my actions), and who judges that I ought to ϕ , then ϕ ing is likely to be a good idea for me. But its being "a good idea" for me to act in certain ways hardly captures the *authority* that we demand from our moral code. We don't think that it would be "a good idea" for torturers to cease their despicable activities. The fact that someone is in a position to give me excellent advice does not thereby grant her the authority to command me.

That the DCT-property might not do such a good job of underwriting moral authority is brought out by the following reflection. Imagine someone saying "I acknowledge that it is morally wrong for me to ϕ , but I don't care about morality—I'm going to do it anyway." We don't interpret this simply as manifesting villainy; we don't merely think that the person *should* care about morality. Rather, the claim seems *contradictory*—it would appear that in acknowledging what he morally ought to do the speaker is acknowledging what he *all-things-considered* should do. For him then to announce that he intends not to do what he judges he *all-things-considered* ought to do doesn't just seem like a poor choice—it seems completely irrational. We might even think that he didn't properly understand the terms used in the sentence that he uttered. The point is put very well by W.D. Falk:

It seems paradoxical that moral conduct should require more than one kind of justification: that having first convinced someone that ... he was morally bound to do some act we should then be called upon to convince him as well that he had some sufficiently strong reason for doing this

same act. “You have made me realize that I ought, now convince me that I really need to” seems a spurious request, inviting the retort “if you really were convinced of the first, you would not seriously doubt the second.” (Falk 1948, 121-2)

But now compare this: “I acknowledge that God commands me not to ϕ , but I don’t care about God’s commands—I’m going to do it anyway.” This person, it seems to me at least, provokes a very different “feel”. Someone who ignores the God that he believes in we think of as misguided, as pitiable. But it hardly has the same air of *paradox* as the former case. Someone who acknowledges but ignores God we think of as a wretched and foolhardy figure; someone who claims to acknowledge but reject morality seems literally not to know what he’s saying.

Things may be worse yet for the fan of the DCT, for we must also heed moral platitudes other than those concerning authority. Consider, for example, the epistemological truisms concerning how we can come to know moral truths. For the DCT-property to be a good contender for being the referent of “moral requirement”, the means by which we discover what God commands must be the same as the means by which we discover moral requirements. I am not going to hazard suggestions about how we are supposed to come to know the mind of God, but I submit that it doesn’t look much like the ways by which we’re supposed to come by moral truths—for there, enshrined in our platitudes (remember: the platitudes shared by theists and atheists alike), are ideas like “Often you can work out what you morally ought to do by putting yourself in another’s shoes” or “People might seem morally good on the outside, but you never really know until you know their motives”. One might object that these are not *inconsistent* with what we think about our epistemic access to God—it’s *possible* that we could come to know what God commands by adopting another person’s point of view—but the point is that this hardly has the central and prioritised role in theistic epistemology that it does in moral epistemology.¹⁵

There are also platitudes concerning the nature of moral agency that must be considered. Here I will mention three. First, many philosophers have argued that sincere use of our moral discourse evinces a commitment to the following: if a person makes a moral judgment that some available action is morally required, she must have some (defeasible) motivation in favor of performing that action.¹⁶ But the DCT-property doesn’t on the face of it satisfy this apparent platitude. Of course, a great many people who believe that God commands them *do* feel motivated to act

¹⁵ Certain theists may object that there are longstanding traditions according to which how we come by moral truths is *exactly* like how we come to know the mind of God: reading the Bible, listening to church leaders, etc. However, these beliefs cannot be part of the *meaning* of our moral terms—they cannot feature in the platitudes that we are trying to gather—on pain of admitting that theists and atheists mean different things with their moral terms—and we have already noted the disagreeable consequences of that.

¹⁶ The person who springs to mind most readily as the defender of this view is Hume, who wrote: “Morals excite the passions, and produce or prevent actions” (Hume 1739/1978, 457). W.K. Frankena wrote that “it would seem paradoxical if one were to say ‘X is good’ or ‘Y is right’ but be absolutely indifferent to its being done or sought by himself or anyone else. If he were indifferent in this way, we would take him to mean that it is generally regarded as good or right, but that he did not so regard it himself” (Frankena 1973, 100). See also Hare 1952, and Smith 1994. (Smith’s version, it is true, is somewhat different, but not enough to affect my argument.)

accordingly, but it is hardly a necessary truth that they do. Second, if an action is impossible for a person, we think that it cannot possibly be the case that she is morally required to perform it. Yet there is nothing to prevent authoritative figures commanding impossible actions, and there seems nothing impossible in the idea of God commanding something impossible. (Perhaps there is some lesson He hopes to teach by doing so.) Third, we restrict the application of certain moral terms to agents who *freely* choose their actions: someone who brings about catastrophe by non-negligent accident evades the moral opprobrium that would have been levelled at him had he done it on purpose. Yet there is no reason to assume that God's commands or approval patterns respect this distinction. Again, one might believe (or hope) that they *do*, but it is hardly a logical truth that they do; and thus again we find a mismatch between our shared moral platitudes and the proposed DCT-property.

I could go on. But in response to this strategy of finding mismatches the proponent of the DCT has a fair response—namely, that she is claiming only that the DCT-property is the *best* satisfier of the platitudes, and this is quite consistent with there being *many* mismatches. The only way really to succeed in demonstrating that the DCT-property is not the best contender is to present an instantiated property that does better still. Either that, or demonstrate that God doesn't exist. Neither of these tasks (I hope you won't be too disappointed to hear!) will be attempted in the final two paragraphs. All I have really been trying to accomplish in this final section is (A) to show that the theistic supporter of the DCT still has arguments to win, and (B) to indicate some grounds for pessimism that she will succeed.

Let me close by recapping a rather complex dialectic. First, it should be emphasized that unless the theist can convince the atheist (or the agnostic) of the existence of God, then the latter have no good reason to endorse the Divine Command Theory. (That said, an atheist *could* endorse the theory, leading to moral nihilism.) But it is widely believed that even if the theist does have good grounds for believing in God, he *still* should not endorse the DCT because of the Euthyphro Problem. I have shown that this is false. The trouble with the DCT is nothing to do with Socrates's bamboozling, nor some modern "dilemma," each horn of which threatens to impale the theorist. The upshot might then be taken to be that the theist (who, by definition, takes himself to have good grounds for believing in God) can happily endorse the DCT after all. But this, I have argued, cannot be assumed without further debate. First the adherent of the DCT needs to make absolutely clear just what *kind* of truth the theory expresses. Latterly, I have been discussing what I consider to be the most promising line that might be taken in this respect: that pursued by Adams. I have argued that even with this clarification made, the theist still needs to earn the right, via argument, to the DCT. It must be demonstrated that the property of *being commanded by God* is the best contender for fulfilling our shared moral platitudes. I have indicated some reasons for thinking that the property will fall short of a good fit. This argument leaves the supporter of the Divine Command Theory with defensible ground (conditional on *theism* being defensible, of course), but with work to do.

REFERENCES

Adams, Robert Merrihew

- 1973 "A Modified Divine Command Theory of Ethical Wrongness." In *Religion and Morality*, edited by Gene Outka and John Reeder, Jr., 318-47. Garden City, NY: Anchor Books.
- 1979 "Divine Command Metaethics Modified Again." *Journal of Religious Ethics* 7: 66-79.
- 1981 "Divine Command Ethics as Necessary A posteriori." In *Divine Commands and Morality*, edited by Paul Helm, 109-119. Oxford: Oxford University Press.
- 1999 *Finite and Infinite Goods*. Oxford: Oxford University Press.

Boyd, Richard

- 1988 "How to Be a Moral Realist." In *Essays in Moral Realism*, edited by Geoffrey Sayre-McCord, 181-228. Ithaca: Cornell University Press.

Brink, David

- 1984 *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.

Brown, J.

- 1964 "The Logic of the *Euthyphro* 10A-11B." *Philosophical Quarterly* 14: 1-14.

Burch, Robert W.

- 1980 "Objective Values and the Divine Command Theory." *The New Scholasticism* 54: 279-304.

Clark, Stephen

- 1982 "God's Law and Morality." *Philosophical Quarterly* 32: 339-47.
- 1987 "Discussion: God's Law and Chandler." *Philosophical Quarterly* 37: 203-08

Cohen, S. Marc

- 1971 "Socrates on the Definition of Piety: *Euthyphro* 10A-11B." *Journal of the History of Philosophy* 9: 1-13.

Cudworth, Ralph

- 1995 "A Treatise Concerning Eternal and Immutable Morality." 1731. In *The True Intellectual System of the Universe*, volume 3, 519-646. Bristol: Thoemmes Press.

Falk, W.D.

- 1948 "'Ought' and Motivation." *Proceedings of the Aristotelian Society* 48: 492-510.

Fischer, J.M. and Ravizza, M., eds.

1992 *Ethics: Problems and Principles*. New York: Harcourt Brace Jovanovich College Publishers.

Frankena, W.K.

1973 *Ethics*. New Jersey: Prentice-Hall, Inc.

Frankfurt, Harry

1971 "Freedom of the Will and the Concept of a Person." *Journal of Philosophy* 68, 5-20.

Hare, Richard

1952 *The Language of Morals*. Oxford: Clarendon Press

Kripke, Saul

1972 *Naming and Necessity*. Cambridge: Harvard University Press.

Lewis, David

1989 "Dispositional Theories of Value." *Proceedings of the Aristotelian Society*, supplementary volume 63: 113-37.

Luther, Martin

1957 *The Bondage of the Will*. 1525. Translated by J.I. Packer and O.R. Johnston. London: James Clarke.

Lycan, William

1986 "Moral Facts and Moral Knowledge." *The Spindel Conference: The Southern Journal of Philosophy*, supplementary volume 24: 79-94.

MacIntyre, Alisdair (with Ricoeur, Paul)

1969 *The Religious Significance of Atheism*. New York: Columbia University Press.

Quine, Willard van Orman

1959 *Methods of Logic*. New York: Henry Holt and Co., Inc.

1981 *Theories and Things*. Cambridge: Harvard University Press.

Quinn, Philip

1978 *Divine Commands and Moral Requirements*. Oxford: Clarendon Press.

2000 "Divine Command Theory." In *The Blackwell Guide to Ethical Theory*, edited by Hugh LaFollette, 53-73. Cambridge, Mass: Blackwell.

Shaftesbury, Anthony Earl of

1900 *Characteristics of Men, Manners, Opinions, Times, etc.* 1711.
London: Grant Richards.

Smith, Michael

1994 *The Moral Problem.* Oxford: Blackwell.

Sober, Elliott

1982 “Why Logically Equivalent Sentences may Pick out Different
Properties.” *American Philosophical Quarterly* 19: 183-89.

Van Fraassen, Bas

1980 *The Scientific Image.* Oxford: Clarendon Press.

Wierenga, Edward

1983 “A Defensible Divine Command Theory.” *Noûs* 17: 387-407.

William of Ockham

1962 *Opera plurima.* 1494-96. Farnborough, England: Gregg Press.