Divine will/divine command moral theories and the problem of arbitrariness

THOMAS L. CARSON

Department of Philosophy, Loyola University Chicago, Chicago, Illinois 60626, USA
e-mail: tcarson@luc.edu

Abstract: A well-known objection to divine will/divine command moral theories is that they commit us to the view that God's will is arbitrary. I argue that several versions of divine will/divine command moral theories, including two of Robert Adams's versions of the DCT and my own divine preference theory, can be successfully defended against this objection. I argue that, even if God's preferences are somewhat arbitrary, *we* have reasons to conform our wills to them. It is not a fatal objection to divine will/divine command moral theories if they imply that God's will/God's commands is/are arbitrary, to some extent.

> Putting God in charge of morality is one way to solve the problem [the problem of whether 'the distinction between right and wrong' is 'real'], of course, but Plato made short work of it 2400 years ago. Does God have a good reason for designating certain acts as moral and others as immoral? If not – his dictates are divine whims – why should we take them seriously? Suppose that God commanded us to torture a child. Would that make it all right, or would some other standard give us reasons to resist? And if, on the other hand, God was forced by moral reasons to issue some dictates and not others – if a command to torture a child was never an option – then why not appeal to those reasons directly. (Pinker (2008), 56)

One well-known objection to divine will/divine command moral theories is that they commit us to the view that God's will is arbitrary, and the arbitrary will of God is not a plausible basis for morality. The gist of the argument can be summarized as follows: since divine will/divine command moral theories claim that moral standards are dependent on God's will/commands, they imply that God cannot be motivated by moral reasons. Among other things, divine will/divine command moral theories imply that God can't will/command something because it is right, since according to such theories nothing can be right independently of God's will/commands. Therefore, divine will/divine command moral theories imply that God's will/commands is/are arbitrary and that God has no reason to

445

will or command one thing rather than another. However, if God's will/commands is/are arbitrary there is no reason for us to take it/them to be a moral standard that is binding on us. This argument or something like it is often attributed to Plato and is widely thought to be a fatal objection to all such theories.[1]

I argue that this objection is ill-founded and that several versions of divine will/divine command moral theories can be successfully defended against it. In the first section of the article, I state the objection and explain its roots in Plato's *Euthyphro*; I also formulate several close analogues of Socrates' *Euthyphro* argument that apply to monotheistic divine will/divine command moral theories. Although the *Euthyphro* argument is ingenious and justly celebrated, neither Socrates nor Euthyphro gives any reason why we should accept Euthyphro's answer to the question: 'do the gods love what is pious because it is pious, or is the pious pious because the gods love it?' The enduring appeal of the *Euthyphro* argument is because many think that Euthyphro's answer to the question (the gods love what is pious because it is pious) is obviously correct and can be easily defended. Many, I dare say most, contemporary philosophers think that Euthyphro's answer to the question is obviously correct, since the other answer (what is pious is pious because the gods love it) makes the loves and hates of the gods arbitrary. However, I will contend that this widely held view is mistaken; the view that things are pious because the gods love them does not imply that the loves and hates of the gods are arbitrary.

I examine five recent divine will/divine command moral theories: Adams's three different formulations of the DCT from his papers published in the 1970s and his later book *Finite and Infinite Goods*, Linda Zagzebski's divine motivation theory, and my own divine preference theory. Adams's first two versions of his theory (from his two papers) can easily handle the objection. Since these theories hold that there are standards of good and bad that are independent of God's commands, Adams can say that God commands what he commands because of his loving concern for human beings and their good/welfare, and that, therefore, God has good reasons to command what he commands. But it is not so clear that Adams's most recent theory in *Finite and Infinite Goods*, a theory that combines a DCT of right and wrong and a divine resemblance theory of good and bad, allows him to give a satisfactory answer to the Euthyphro problem.

Zagzebski and I attempt to base all moral standards, including standards of good and bad, on God's motivations/preferences. Our theories don't permit us to give the quick and ready answer to the arbitrariness objection that Adams's first two versions of the divine command theory (DCT) permit him to make. I present criticisms of Zagzebski's answer to the arbitrariness objection, but argue that her theory can be defended against this objection on grounds other than those she gives. I also defend my divine preference theory against Zagzebski's criticism that it does not afford an adequate answer to the arbitrariness objection. I don't attempt to show that God has reasons for all of his preferences. Rather, I argue

that, even if God's preferences are somewhat arbitrary, *we* have reasons to conform our wills to them. It is not a fatal objection to divine will/divine command moral theories if they imply that God's will/commands is/are arbitrary to some extent.

The Euthyphro problem

The arbitrariness objection is widely attributed to Plato's *Euthyphro*. In the *Euthyphro*, Socrates asks Euthyphro to define piety/justice (*hosion*).[2] Euthyphro proposes the following definition: 'What is dear to the gods is pious, what is not is impious' (Plato (2000), 7a). Socrates rejects this definition because he claims that it leads to a contradiction. According to the Ancient Greek religion, the gods often disagree – what some gods love other gods hate. Therefore, some things are both loved by some gods and hated by other gods. It follows that some things are both pious and impious, but that is impossible. Euthyphro concedes this objection (Plato (2000), 7e–8a). At Socrates' suggestion, Euthyphro endorses the following revised definition: 'the pious is what all the gods love, and the opposite, what all the gods hate, is impious' (Plato (2000), 9d–9e). This revised definition avoids Socrates' objection to the first definition since it doesn't commit one to endorsing contradictions (saying that certain acts are both pious and impious) in cases in which the gods disagree. In cases in which a particular act is loved by some gods and hated by other gods, the second definition implies that the action in question is neither pious nor impious. (However, the second definition raises another problem that is not mentioned in the dialogue and, to my knowledge, has not been noted by anyone else – if the gods disagree about everything or almost everything, then Euthyphro's definition implies that nothing, or almost nothing, is pious or impious. In this connection, it is important to ask how many gods 'all the gods' refers to – if the list of the relevant gods is very long and includes minor deities, then the likelihood of agreement among all the gods is greatly diminished.)

Socrates begins his criticism of this second definition by asking, 'Is the pious being loved by the gods because it is pious, or is it pious because it is being loved by the gods?' (Plato (2000), 10a). He asks Euthyphro to choose between the following:

(A)   The gods love what is pious because it is pious.
(B)   What is pious is pious because the gods love it.

Clearly A and B are inconsistent and Socrates and Euthyphro both assume this. More problematically, they both assume that one of these alternatives must be true.[3] Euthyphro answers Socrates' question by endorsing A and rejecting B. He doesn't explain or defend this answer nor does Socrates ask him to. Socrates seems to agree with Euthyphro's answer. Socrates then argues that Euthyphro's answer to the question is inconsistent with his definition of piety.

> If the god-loved and the pious were the same, my dear Euthyphro ... then the pious would be pious because it was being loved by the gods. But now you see that they are in opposite cases as being altogether different from each other. (Plato (2000), 10e–11a)

Socrates' argument against Euthyphro's second definition of piety can be summarized as follows:

1. If Euthyphro's (second) definition is true/correct, then B is true (and A is false).
2. B is false (and A is true).

Therefore,

> Euthyphro's (second) definition is false/incorrect.

This argument is valid (it has the form of a *modus tollens* argument). Note that the argument is valid if premise 2 reads simply 'B is false'.[4] Therefore, if both premises are true, the conclusion must be true. Premise 1 is clearly true; it is true by definition. If being pious *means* being loved by all the gods (or if piety is the property of being loved by all the gods), then what is pious is pious because it is loved by all the gods. The only serious question about the argument is whether premise 2 is true. Socrates doesn't give a serious argument for thinking that A is true (and B is false). He asks Euthyphro to choose between A and B. Euthyphro answers that A is true, but he doesn't give any reasons for this. Socrates agrees with this answer so he doesn't ask Euthyphro to defend it. But clearly this assumption needs to be defended. This glaring weakness of the *Euthyphro* argument is masked by the fact that the *Euthyphro* is written as a dialogue.

This problem notwithstanding, many contemporary philosophers side with Socrates; they think that B is clearly false. Here's the argument. The gods must have some reason for loving and hating the things they love and hate. Otherwise, their loves and hates are *arbitrary*. If the gods' loves and hates are arbitrary, then there is no reason to take them seriously as the ultimate standard for morality. This argument assumes that if B is true, then the loves and hates of the gods must be arbitrary. But this assumption is false. Given Euthyphro's definition, we can't say that the gods love what they love because it is pious. This rules out one possible way in which the loves and hates of the gods could be non-arbitrary. But the conclusion of this argument, that if B is true then the loves and hates of the gods are arbitrary, follows only if we accept something like the following:

> Either we must agree with Socrates that the gods love what is pious/right because it is pious/right, or else we have to say that the gods have *no reason whatever* for loving and hating the things they do.

This statement presents a false dichotomy. Given his definition of *hosion*, Euthyphro can't say that the gods conform their loves and hates to some

independently existing standard of *hosion*. But this leaves open many other possible reasons why the gods might love some things and hate others. There are all sorts of different reasons one can have for loving or hating something. So, at best, the Euthyphro argument is incomplete, and, if we extend it in the way analogous to many recent arguments to the effect that divine will/divine command theories make God's will arbitrary (see endnote 1), the argument clearly fails. However, even if we grant all of this, it still might be the case that the kinds of reasons that the gods could have for loving and hating things (given the truth of Euthyphro's theory) don't give *us* reasons to regard their loves and hates as a moral standard that is binding on us. The key question is this: given the kinds of reasons that the gods could have for loving and hating the things they do *that are compatible with the truth of Euthyphro's theory*, do *we* have reasons to conform our will to the loves and hates of the gods? (I will have much more to say about this point later.)

With only slight modifications, Socrates' arguments apply to monotheistic analogues of the positions discussed in the *Euthyphro*. Consider the following monotheistic analogue of Euthyphro's second definition (I will call this 'the divine attitude theory of right and wrong'):

> What is right is what God loves. What is wrong is what God hates. What is right is right *because* God loves it. What is wrong is wrong *because* God hates it.

Socrates' crucial question can be modified as follows:

> Does God love what is right because it is right, or are right acts right because God loves them?

This question asks us to choose between:

> A′.   God loves right acts because they are right.
> B′.   Right acts are right because God loves them.

The modified version of Socrates' argument can be stated as follows:

> 1′.   If the divine attitude theory of right and wrong is true, then B′ is true and A′ is false.
> 2′.   A′ is true and B′ is false.
>
> Therefore, the divine attitude theory of right and wrong is false.

This argument is valid (the conclusion follows from the premises), and premise 1′ is clearly true. (The argument is also valid if we modify 2′ to read simply 'B′ is false'.) The only thing in doubt is premise 2′. Why think that premise 2′ is true – why think that A′ is true and B′ is false?

If we follow the argument for A over B presented earlier, the argument would be roughly this: either we must say that God loves right acts because they are right

or else God has *no reason whatever* for loving and hating the things he does (or at least no reason that can appropriately ground moral/value obligations). This argument has the same problems as the earlier argument. It rests on a false dichotomy – it is possible that God loves or hates actions for reasons other than their being right or wrong.

A modified version of Socrates' argument also applies to the DCT. The DCT holds the following:

> What is morally obligatory/permissible/wrong is what God commands/ permits/ forbids. An act is morally obligatory because God commands it; an act is morally permissible because God permits it; and an act is morally wrong because God forbids it.

The analogue of Socrates' question for this theory is:

> Does God command/permit/forbid us to do what is obligatory/permissible/ wrong, because it is obligatory/permissible/wrong, or is what is obligatory/ permissible/wrong obligatory/permissible/wrong because God commands/ permits/forbids it?

This question asks us to choose between:

A″.   God commands/permits/forbids us to do what is obligatory/ permissible/wrong, because it is obligatory/permissible/wrong.

B″.   What is obligatory/permissible/wrong is obligatory/permissible/ wrong because God commands/permits/forbids it.

The DCT commits us to accepting B′ and rejecting A′. If we endorse A′ (or if we say that B′ is false), then we must reject the DCT. The argument for accepting A″ and rejecting B″ is similar to the argument for accepting A over B (and accepting A′ over B′). The argument is roughly as follows: either we must say that God commands/forbids right/wrong actions because they are right/ wrong or else God has *no reason whatever* for commanding/forbidding what he commands/forbids. This argument is open to the same objections noted earlier.

The *Euthyphro* argument has a glaring weakness; it gives us no reason to accept Euthyphro's answer to Socrates' question. I have formulated a subsidiary argument for Euthyphro's answer analogous to recent arguments against the DCT, but that argument is also seriously flawed. Surely it is possible that the gods have reasons to love/hate what they love/hate, even if there is no standard of piety independent of the gods' loves and hates. Similarly, it is possible that God has reasons to command/will what God commands/wills, even if there is no independent standard of right and wrong. But this alone doesn't save divine will/divine command moral theories from related objections. It might still be the case that the kinds of reasons God could have for willing/commanding what God wills/commands that are consistent with the truth of divine will/divine

command moral theories provide *us* with no reason to defer to God's will or obey God's commands. The key question is what *we* have reason to do in light of God's willing/commanding that we do certain things. More precisely, given that God's reasons for willing/commanding what God wills/commands don't include God's being moved by independent moral facts about what is right and wrong, can the fact that God wills or commands us to do certain things give *us* reasons for thinking that God's will/commands is/are binding on us? A closely related but distinct question is this: how can God's reasons for commanding/willing certain things (which according to divine will moral theories can't be *moral reasons* that involve God's recognition of independent moral facts) create *moral reasons* for us?

Different versions of divine will/divine command theories give very different accounts of what sorts of reasons God can have for willing/commanding as God does and, therefore, we need to look at the *details* of particular divine will moral theories and ask if they can provide us with good reasons to follow God's will/commands. We cannot answer these questions in the abstract without considering the details of particular versions of divine command/divine will moral theories. I now turn to five recent divine will/divine command moral theories and their answers to the Euthyphro objection.

Adams's three modified divine command theories

In his paper 'A modified divine command theory of ethical wrongness', Robert Adams defends a DCT of right and wrong, but holds that there are standards of good and bad, personal welfare, and virtue and vice that are independent of God's commands (Adams (1973), 114–116). Adams writes:

> When I say 'it is wrong to do X', (at least part of) what I *mean* is that it is contrary to God's commands to do X. But 'It is contrary to God's commands to do X' implies 'it is wrong to do X 'only if certain conditions are assumed – namely, only if it is assumed that God has the character I believe him to have, of loving his creatures.[5]

Adams can't say that God commands what he commands *because* it is morally right (independently of being commanded by God). But Adams is *not* committed to the view that God's commands are arbitrary or that God has no reason to command one thing rather than another. Adams can say that God commands what God commands because God is loving, omniscient, and desires our good; because God loves human beings, God desires what is good for us.[6]

Adams presents his second modified DCT in 'Divine command metaethics modified again'. He writes:

> My new divine command theory of the nature of ethical wrongness, then, is that ethical wrongness *is* (i.e., is identical with) the property of being contrary

to the commands of a loving God. I regard this as metaphysically necessary, but not an analytic or a priori truth. (Adams (1979), 139)

According to Adams, Atheists who reject the DCT understand the concept of wrongness, but, unbeknownst to them, the best theory about the nature of wrongness is that wrongness is the property of being contrary to the commands of a loving God. Similarly, our pre-scientific ancestors all understood the concept of water but they didn't know the correct theory about the nature of water – that water is H_2O (Adams (1979), 136).

In 'Divine command metaethics modified again', Adams says nothing about standards of good and bad and seems to continue to assume that there are objective standards of good and bad and welfare that are independent of God's will and commands. At least everything he says there is consistent with his holding that there are independent standards of good and bad. Thus, although he doesn't address the arbitrariness objection in this paper, his second theory is consistent with exactly the same reply to the arbitrariness objection that the first version of his theory permits. He can't say that God commands what he commands *because* it is morally right (independently of being commanded by God). But he is *not* committed to the view that God's commands are arbitrary or that God has no reason to command one thing rather than another. Adams can say that God commands what he commands because God is loving and, because he is loving, God desires what is good for us.

Can Adams give a satisfactory answer to the second question I noted near the end of the section on the Euthyphro problem? If there are no independent *moral facts* about what is right and wrong that help constitute God's reasons for commanding what God commands, how can God's commands create *moral reasons* for us? Adams has a perfectly good answer to this question. He can say that the fact that a loving God commands us to do X for certain reasons (together with the fact that God's reasons include independent facts about what is good and bad and independent facts about what is good and bad for human beings) creates *moral reasons* for us to do x. In this way, God's commands (*together with* those independent facts about what is good and bad and what is good and bad for human beings) can create moral facts about right and wrong.[7]

In his more recent book *Finite and Infinite Goods*, Adams proposes yet a third version of his modified DCT. He states his theory of right and wrong as follows:

> An action's being morally obligatory consists in its being commanded by God and . . . an action's being morally obligatory consists in its being contrary to a divine command. (Adams (1999), 250)

He qualifies his theory by saying that he endorses it only 'on the assumption that':

> the character and commands of God satisfy certain conditions. More precisely, I assume that they are consistent with the divine nature having

> properties that make God an ideal candidate, and the salient candidate, for the semantically indicated role of the supreme and definitive Good . . . It is only the commands of a definitively good God, who . . . is not cruel but loving, that are a good candidate for the role of defining moral obligation. (Adams (1999), 250)

Adams stresses that he endorses the DCT as 'a theory of the nature of moral obligation only, and not of moral properties in general' (Adams (1999), 250).

The theory of right and wrong Adams defends in *Finite and Infinite Goods* is roughly the same as the theory he proposes in 'Divine command metaethics modified again'. However, *Finite and Infinite Goods* presents a substantially different theory of value than Adams's earlier writings. In this book, Adams defends a Platonistic theory of value according to which God is perfectly good and other things are good to the extent that they resemble God (Adams (1999), 14).

The two earlier versions of Adams's theory can answer the arbitrariness objection by appealing to the idea that God's commands are motivated and shaped by God's love and his concern to promote the good, where what is good is determined by standards of value independent of God. Since the divine resemblance theory of value is a component of the most recent version of his theory, Adams can no longer say that God's reasons for making certain commands rather than others include God's being moved by standards of value independent of God. This raises the question of whether Adams's most recent formulation of his theory can answer the arbitrariness objection.

Adams holds that, in order to count as the form or standard of the Good, God must have certain features; a deity who was cruel or who loved cowardice could not count as the form of the good (Adams (1999), 46).

> It might be thought that if God is the constitutive standard of excellence as I have claimed, the ascription of excellence to God will be trivial and without content, merely saying that God is like God. This does not follow, for the claim that God is the standard is not inscribed as the first line on a blank slate of ethical theory. It is made against the background of many substantive beliefs about what properties are excellences that must be reflected somehow in the character of any being that is the standard of excellence. One important excellence is justice. It clearly matters to the persuasive power of God's character, as a source of moral requirement, that the divine will is just. (Adams (1999), 253–254)

Adams says that God's commands cannot be the source of moral obligations unless God is just. Adams's DCT relies on what he calls a 'thin theory of justice' which 'does not presuppose moral obligation as such'.[8] According to Adams, standards of justice are independent of God's commands and God's will, but not

independent of God's nature. God is the constitutive standard of goodness/ excellence and justice is one of God's excellences. Things are good/excellent to the extent that they resemble God, but Adams does not say that every possible omniscient, omnipotent deity would constitute the standard of goodness:

> A deity would have to satisfy certain conditions (for instance, not being sadistic, and not loving cowardice) in order to be the salient candidate for filling the role indicated by our concept of the Good, though it is part of the point of my theory that such requirements do not completely determine what the deity would be like. (Adams (1999), 46)

This leaves us with questions about the criteria for determining whether a particular deity is or is not a suitable candidate for being the ultimate standard of goodness. Surely these criteria (which in Adams's theory are more basic than theories of good and bad and right and wrong) must be independent of God's actual nature. In Adams's theory, the meaning of conventional moral concepts is the starting point and foundation for everything else. He makes a sharp distinction between our conventional moral concepts and metaphysical theories about the nature of moral properties. He says that his divine resemblance theory of good and bad and his DCT of the right and wrong are the theories that best account for the role assigned to goodness and rightness by our conventional moral concepts (Adams (1999), 15–18 and 250–258).

As I argue below in my discussion of Zagzebski, the concept of love presupposes the concept of welfare – a necessary condition of loving someone is that one desires the welfare/*good* of the other person. This raises the question of whether Adams's theory involves some kind of vicious circularity. Adams claims that something's being good or good for someone consists in its resembling a loving God and that the property of being morally obligatory is identical to the property of being commanded by a loving God. But the statement that God is loving presupposes a prior independent notion of welfare/goodness; see below, where I argue that this kind of vicious circularity is a serious problem for Zagzebski's theory. To the extent that this is a serious problem for Zagzebski, it is very likely a serious problem for Adams as well. But I won't insist on this point. I am much more concerned to show that I have a good answer to the arbitrariness objection than to show that the third version of Adams's theory does not. If Adams's most recent theory is viciously circular, he can avoid this problem by modifying along the lines of my theory – see below.

In the two earlier versions of his theory, Adams defends a hybrid theory according to which some, but not all, moral notions are determined by God's commands. These two theories appeal (or can appeal) to realist standards of goods and bad, welfare, and virtue and vice independent of God's will. Clearly divine will/divine command theories of morality that anchor standards of good and bad as well as standards of right and wrong in God's will/commands cannot

make use of the quick and ready answer to the Euthyphro objection that is open to defenders of the DCT who think that standards of good and bad are independent of God, and, to that extent, such theories are much more vulnerable to the arbitrariness objection.

Zagzebski's divine motivation theory

Zagzebski defends what she calls 'exemplarism'. She defines moral concepts such as right and wrong and good and bad in terms of the motivations of moral exemplars or paradigmatically good people. Moral theorizing, she contends, should begin by pointing to and identifying moral exemplars. Identifying exemplars is prior to having a theory about what makes someone an exemplar. Similarly, identifying paradigm cases of gold or water is prior to analysing or discovering the essence of gold or water. She states two versions of her theory, a naturalistic version that doesn't presuppose the existence of God and a full theory (her preferred version of the theory) for Christian philosophers that presupposes the divinity of Christ. Her full theory claims that God (Jesus) is the perfect moral exemplar. She defines fundamental moral concepts in terms of the exemplar's (God's) motivations.

To have a good life is to have the kind of life that the exemplar would want for herself (Zagzebski (2004), 116 and 134). An act is right (permissible) provided that: (1) it is not the sort of act that an exemplar would characteristically refrain from doing in like circumstances, and (2) an exemplar would not feel guilty for doing it in like circumstances. An act is wrong provided that an exemplar: (1) would characteristically not do it in like circumstances and (2) would feel guilty about it if s/he did it. An act is a moral duty provided that: (1) an exemplar would characteristically do that act in like circumstances, (2) an exemplar would insist on doing it, and (3) an exemplar would feel guilty for not doing it.[9]

Zagzebski claims that her theory gives a good answer to the arbitrariness objection, but that Adams's theory and my divine preference theory do not.

According to Zagzebski, arbitrariness is a problem for traditional versions of the DCT, which imply that we have a duty to follow God's commands even if he commands cruel and hateful acts (Zagzebski (2004), 259–260). Adams's theory doesn't have this consequence, since if God were to make such commands he would show himself to be unloving and thus, according to Adams, we wouldn't be obligated to follow God's commands in that case. She contends that Adams's answer to the Euthyphro objection is still unsatisfactory because it is 'ad hoc'.

> It is possible that Adams's proposal succeeds at answering the objection that it is designed to address, but even so, it seems to me that it is unsatisfactory because it is ad hoc. There is no intrinsic connection between a command and the property of being loving, so to tie morality to the commands of loving

> God is to tie it to two distinct properties of God. In DM theory [Zagzebski's divine motivation theory], there is no need to solve the problem of whether God could make it right that we brutalize others by making any such modifications to the theory, since being loving is one of God's essential motives. The right thing for humans to do is to act on motives that imitate the divine motives. (Zagzebski (2004), 260)

In this respect, Zagzebski's theory may be preferable to the DCT on grounds of theoretical simplicity, but this is hardly a decisive objection to Adams's theory. It is unclear why it is a serious problem for Adams that he ties morality to two distinct properties of God instead of one. (Adams appeals to the commands of a loving God; Zagzebski appeals to God's loving motives). More needs to be said here. However, I won't dwell on this point, because if my defence of my own theory against the arbitrariness objection works, it can also be extended to defend Adams's theory against the objection – see the discussion of my divine preference theory below.

Zagzebski claims that her own theory completely avoids the problem of arbitrariness:

> The arbitrariness problem may or may not be answerable in the DC [divine command] theory, but the problem does not even arise in DM theory. That is because a will needs a reason, but a motive *is* a reason. . . . Motives are both initiating reasons and justifying reasons. If we know that God acts from a motive of love, there is no need to look for a further reason for the act. On the other hand, a divine command requires a reason, and if the reason is or includes fundamental divine motivational states such as love, it follows that even DC theory needs to refer to God's motives in order to avoid the consequence that moral properties are arbitrary and that God Himself is not good. Such a move makes divine motives more basic than the divine will even in DC theory. (Zagzebski (2004), 261)

If her theory avoids the arbitrariness objection, that would be no small thing. It would show that divine will moral theories can answer the Euthyphro objection even if they don't hold that standards of good and bad are independent of God and God's will/commands.

Does Zagzebski's answer to the arbitrariness objection succeed? I think not. Zagzebski to the contrary, not every motive is a justifying reason. There are bad motives. The motive of hatred (or unjustified hatred) is not a justifying reason. Here, Zagzebski could reply that the motive of love is a justifying reason and that everything God does is motivated by his love. Even if everything God does is motivated by love,[10] it is not so clear that love is always a good motive, i.e. a motive that provides justifying reasons. Not every act motivated by love is a good action that is worthy of imitation. Wrong acts and stupid acts can be motivated by love. Murders and other crimes are sometimes motivated by love. Why couldn't God's

loving motives go awry in similar ways? Although Zagzebski doesn't address this question, a promising answer is available to her. In cases in which wrong actions are motivated by love, e.g. the parent who harms another person's child in order to help her own child win a competition, the agent fails to show love for those she harms. Inasmuch as God loves everyone (equally), his loving motives can't be objectionable in the way that human actions motivated by very partial and limited love can be.

Even if we accept her claim that God's love is (always?) a justifying reason, there is another serious problem for Zagzebski. Her answer to the arbitrariness objection depends on her appeal to the idea that God is loving and motivated by love. According to her, what is good and good for people is determined by the motives of a loving God. This creates serious problems for Zagzebski because it commits her to rejecting standard definitions of love, according to which a necessary condition of loving a person is that one desires that person's good/welfare. Standard definitions of love make the notion of good (or someone's good/welfare) conceptually prior to the notion of love.[11] If the standard definitions of love are correct, it is viciously circular to say that what is good (or good for someone) is determined by the motives of a loving God. This is a problem that I raised earlier for my own divine preference theory of value (Carson (2000), 246–248). Zagzebski is aware of this problem and quotes from my statement of the problem (Zagzebski (2004), 221). The passage she quotes from me is as follows:

> We cannot say that God's being loving consists in the fact that God desires what is good for earthly creatures. It would be viciously circular to say that what is good and bad is determined by what a loving God desires and then go on to define what it is for God to be loving in terms of what is good and bad for humans (independently of what God prefers). (Carson (2000), 246)

Zagzebski attempts to defend an alternative account of love that avoids this problem. She writes:

> Carson is right that it is common to think that love consists in desiring the welfare of the loved one, but it seems to me that such a desire is an effect of love, not a defining characteristic of the state of love itself. We do not think that it is constitutive of love in general that it includes desiring the welfare of its object, since it is perfectly reasonable to speak of loving objects that have no welfare. Someone may love philosophy, art, scientific truth, and many other things that are not such that something can be good for them. I assume that love of persons and love of philosophy are both love in that they consist in the same type of emotion toward their objects. Differences between them obviously arise because they are different kinds of objects. Nonetheless, I see no reason to think that one consists in a desire for the welfare of the object whereas the other does not. (Zagzebski (2004), 221–222)

Zagzebski claims that one can have love for something that has no welfare. One can love art, philosophy, or chocolate. But this is strained. We don't love persons in the same way we love inanimate objects or activities. These examples involve different senses of the word 'love'. Zagzebski claims that love consists in having a certain kind of emotion towards an object. Love is an emotion that is expressed in certain characteristic ways. Like pity and anger, love has characteristic ends (Zagzebski (2004), 222). God does not aim at our good *qua* good. Because he loves us, God wants us to be healthy and have friendships, etc. (Zagzebski (2004), 222)). God also wants us to be happy, not suffer, and be full persons (Zagzebski (2004), 320). God desires those things because that is what loving beings desire for those they love (Zagzebski (2004), 222).

According to Zagzebski, desiring the good of another person is an effect of love, not a defining characteristic of loving that person. Love is an emotion that leads to such things as a desire for a person's welfare, but concern for the other person's welfare is not necessary for love. This is doubtful. Suppose that I feel what Zagzebski calls the emotion of love for someone but my emotion doesn't have its usual causal consequences, and I fail to have *any* concern for her welfare (or any concern for the things that constitute her welfare). Shouldn't we conclude that I don't really love the other person? I won't pursue this further here. As Zagzebski observes, everything depends on the proper definition of love – 'the real issue is what love consists in' (Zagzebski (2004), 221). However, since she rejects the standard view, Zagzebski has the burden of proof – a burden she hasn't fully met.

My divine preference theory

In *Value and the Good Life*, I defend a divine preference theory of value by appealing to a divine preference theory of rationality. I argue that what is good is what it is rational to desire and that (on the assumption that the 'right' sort of God exists) the best theory of rationality is that it is rational (in the sense of 'correct') for someone to desire what God prefers that she desire. I didn't propose a theory of right and wrong action in *Value and the Good Life*. In fact, I believe that if a suitable God exists, then right and wrong are in some way determined by God's will, but I haven't defended a preferred version of the divine will theory of right and wrong.[12]

The preferences of a cruel, hateful deity are not a plausible standard of good and bad. In order for God's will to constitute the standard of good and bad, God must have certain virtues or psychic traits. I hold that if a suitable God exists, then God's preferences are the standards of good and bad, but, if such a God doesn't exist, then some other standards determine what is good and bad.[13] Here is a summary statement of my view from *Value and the Good Life*:

> *If* there is an omniscient God who designed and created the universe and
> human beings for certain purposes/reasons, who cares deeply about human

beings, and is kind, sympathetic, and unselfish (in the ways explained above), then God's preferences are the ultimate standard for the correctness/ rationality of human preferences and for the goodness or badness of things. (If such a God exists, it is rational (correct) for person S to have a certain preference, p, if, and only if, God prefers that S have p.) If such a God does not exist, then the correctness or rationality of human preferences (and the goodness or badness of things) is determined in some other way.[14]

According to my theory, if a kindly and omniscient (etc.) God exists, there is no standard of good and bad independent of God's will. Since I think that love needs to be defined in terms of desiring the other person's good, I can't say that what is good depends on the will of a loving God. That would be viciously circular. I need to describe the kind of God whose preferences determine what is good/bad without any reference to God's being good or loving. I describe God as an omniscient person who created the universe and human beings for certain purposes/reasons, and is caring, kind, sympathetic, and unselfish. If Zagzebski is right about the concept of love (and I very much hope that she is), then I can modify my theory to say that it is the preferences of a loving omniscient God that determine the goodness and badness of things. However, if she is not right about this, then her theory has the same problem as mine. In that case, she can't appeal to the idea that God's motives are loving, and she may have to modify her description of God's motivations along the lines of my theory. (These comments also apply to Adams's theory. If his theory is also viciously circular in this way, he can avoid this problem by modifying it along the lines of my theory.)

Zagzebski poses the following dilemma for me:

> What determines that God prefers that we prefer A rather than B, when both A and B are equally compatible with – and even equally expressive of – his underlying kindness, sympathy, unselfishness, and so on?
>   This question generates a dilemma. Either God's preference of A over B is determined by His underlying motives, or it is not. If it is, the theory is really a form of Divine Motivation theory, and if so, I have no objection. If instead, God's underlying motives are insufficient to determine His preferences, then we need a reason for the fact that He prefers A to B. A preference, like a command, needs a reason. By contrast, a motive is a reason. When we give the motive, we give the reason.[15]

Zagzebski says that if God's preferences are determined by *his underlying motives*, e.g. God's love/kindness, then my theory is really a divine motivation theory, and she doesn't object to it. But she insists that God's motivations provide a deeper explanation of why things have the moral properties they have. God's motivations provide 'the rock bottom features of the divine psyche out of which arise obligation or moral correctness' (Zagzebski (2004), 269). Zagzebski is correct that motives often determine preferences and commands. She is also correct to claim that, when God's motives determine his commands and preferences, his motives constitute a deeper level of explanation of moral concepts than his

commands or preferences. However, I strongly disagree with what she says about cases in which God's preferences are not completely determined by his loving motives.

I accept the second horn of the dilemma she poses (I think that there probably are cases in which God's preferences aren't fully determined by his underlying loving/kindly motives), but I don't accept her view that this saddles me with a serious objection about the arbitrariness of God's will. It seems possible that God could have preferences that aren't determined by his underlying virtues (his kindness and sympathy, etc.). Consider a case in which God's virtues and essential nature underdetermine his preferences for our choices and preferences. Suppose that God's preferring that I prefer A to B and God's preferring that I prefer B to A are both equally compatible with and equally expressive of God's underlying kindness, sympathy, and knowledge, etc. It is possible that there are many such cases. Many different sorts of preferences seem to be compatible with the description of God that I use in my theory. On the other hand, it is also possible that God's purposes might hang together in ways we don't understand and that many things that seem arbitrary to us are not.

Suppose that God prefers that I prefer A to B (and that I choose A over B), but this preference doesn't derive from any of God's virtues. Let us grant that this implies that God doesn't have any reason to prefer that I prefer A to B (or choose A over B). Nonetheless, *I* still have reason to prefer A to B (and choose A over B). Even if God's preferences are in some sense arbitrary, they give *us* reasons to conform our preferences to God's preferences concerning them. Divine preferences for which God has no reasons can generate reasons for *us*. Sometimes friends and loved ones have whims that involve wanting us to do or prefer certain things. For example, suppose that, on a whim, my father wants me to prefer X to Y (or choose X over Y). Suppose that independently of this, it is not objectionable for me to prefer X to Y (or choose X over Y). This gives me some kind of reason to have the preferences that he wants me to have. My reason(s) might include reasons of gratitude. I love the Chicago Cubs and strongly desire that they play in the World Series. My having this desire is optional. Even though I do not have any good or compelling reasons for having this desire, my having it gives my wife some sort of reason to prefer or desire that the Cubs play in the World Series. My wanting her to want or do something gives her some reason to want or do it. It is not a problem for my view that it implies that moral facts can be created by arbitrary divine preferences (preferences God has for which he has no reason).[16]

But there is a serious problem with the forgoing. On my view, something is intrinsically good provided that it is rational for people to have an *intrinsic* desire that it occur or be the case. But my examples don't show anything about the rationality of intrinsic desires/preferences. They only show that the arbitrary desires of others can give us reasons to have instrumental desires for things they want us to desire. In the case of my wife and my desire that she desire that the

Cubs win the World Series, my desire gives her reasons to desire that the Cubs win the World Series *as a means* to my happiness and *as a means* to pleasing me, but it does not give her any reason to have an intrinsic preference that the Cubs win. If God has an arbitrary preference that I prefer X to Y, that gives me a reason to prefer X to Y as a means to pleasing God, but it does not give me reason to have an intrinsic preference for X over Y.

Here is my answer to this objection. What it is correct/rational for me to desire is what God wants me to desire. What it is correct for me to desire as an end is what God wants me to desire as an end. There may be things that God desires/prefers me to desire but that God does not desire me to desire as ends. God might simply want me to desire these things as a means to pleasing God and deferring to God's judgement. In such cases, God's preferences don't give me reasons to have intrinsic preferences for the things that God wants me to desire. But if God desires that I desire X for its own sake, then I have reason to desire X for its own sake, and X is intrinsically good. If God desires that I intrinsically prefer X to Y, then it is correct for me to prefer X to Y for its own sake, and X is intrinsically better than Y.

In order to adequately address Zagzebski's thought experiments about cases in which God has desires that aren't fully motivated by his virtues (his kindness/lovingness, etc.), it would be helpful to contrast my theory with what I call the 'ideal observer version of the divine preference theory of value' (IODPT). The IODPT can be stated as follows:

> IODPT. The correctness/rationality of human preferences and the goodness and badness of things should be judged by the preferences that a kindly/loving and omniscient God would have if such a God existed. Something is good (bad) if a kindly/loving omniscient God would prefer that we prefer that it exist (not exist).

God's being kind and omniscient underdetermines God's preferences. These characteristics are compatible with different kinds of psychological characteristics and inclinations. There is no reason to think that all possible kindly (loving?) omniscient deities would have the same preferences about *everything*.[17] In light of this, we probably should revise the IODPT by saying that X is good/bad provided that *all possible* kindly omniscient Gods would prefer that X exist/not exist.

Several standard versions of the ideal observer theory (IOT), including those defended by Roderick Firth and Richard Brandt, hold that the reactions of an ideal (but human) observer determine the truth or falsity of normative judgements. These theories hold roughly the following: a favourable (unfavourable) moral judgement about X is true provided that, *if s/he existed*, an ideal human observer (or all human ideal observers) would have a favourable (unfavourable) attitude about X. These versions of the IOT do not posit the actual existence of an ideal observer. IODPT is similar to the ideal observer theory in that it takes the

preferences that God would have (if God existed) to be the ultimate standard of rationality. There may be truths about what God would prefer (if God existed), even if God doesn't actually exist. According to the IODPT, what things are good or bad is independent of whether God actually exists. The IODPT differs from Firth's and Brandt's versions of the IOT in that it takes God, rather than a super-duper human being, to be the ideal observer. By contrast, my DPT makes God's *actual preferences* the ultimate standard of rationality, only if a loving and omniscient creator of the universe actually exists.

I think that we should prefer my DPT to the IODPT. The following scenario makes clear what is at stake between my theory and the IODPT. Suppose that there is a kindly (loving?) omniscient creator of the universe who has clear preferences about how human beings live their lives. Suppose also that certain other possible kindly (loving?) omniscient creators of the universe would have different preferences about how humans should live their lives. According to my theory, God's *actual preferences* determine what it is correct to prefer (and thus what is good and bad). But the IODPT does not imply that God's actual preferences have any priority over the preferences of other possible Gods. In such a case, we should accept the implications of my DPT and say that what is correct for us to prefer is determined by God's *actual preferences.* Human beings have reasons to harmonize their preferences with God's. But aren't God's actual preferences arbitrary since they could have been different without God's losing any of God's perfections? God might have had different preferences from those God actually has. God might have designed the cosmos and our place in it for different purposes from those God has in fact designed them for. God might not have created the earth or human beings at all – God might have chosen to create other forms of intelligent life instead. But the contingency of (at least some of) God's actual preferences and designs does not alter the fact that we have reasons to submit ourselves to those plans and preferences. Given the existence of actual divine purposes within a particular plan of creation (which has its source in a kindly omniscient being), we have reasons to fulfil those purposes. If there exists a kindly omniscient (etc.) creator of the universe who has a plan for the universe in which human beings play a role, then it makes sense for us to fulfil our purpose in that plan. If such a God exists, then it is reasonable for us to try to live in accordance with God's plan and God's preferences and not resist our natural ends and the design of the universe. It would be reasonable for us to defer to the superior knowledge and wisdom of such a God, just as it is reasonable for small children to defer to the wishes of their parents, provided that the parents are loving. Our own vantage point for judging normative moral questions is seriously deficient.[18]

In this connection, I need to raise an earlier question in connection with my view. Since, on my view, there are no moral facts of any sort that are independent of God's will, none of God's reasons for preferring or commanding that I do one

thing rather than another can be moral reasons. How can God's non-moral reasons create moral reasons for us? Very briefly, the answer is that God doesn't have independent moral reasons to will as he wills, but God's willing as God wills creates moral reasons *for us* to do as God wills. The fact that a kindly, omniscient creator of the universe wants/commands me to do X gives me an overriding reason to do X. God's preferring/commanding that I do X creates/constitutes the moral fact that my doing X would be good/right. Granted that God's preferring/commanding that I do X creates *a reason* (perhaps a strong reason) for me to do X, why think that those reasons are overriding? The answer is that there is no one else (including me) whose standpoint for judging these things is nearly as good as God's. God has perfect knowledge and a full understanding of how what I do fits with the purposes for which God created the universe. To the extent that my own preferences might conflict with the preferences God wants me to have, it is reasonable for me to defer to God's will and God's superior standpoint for judging these things. Because of this, and, very importantly, because there are no independent moral facts that could give us reasons to defy God's will,[19] it is reasonable for us to defer to God's will to the extent that we can discern it.

An Important note

This argument (and any argument for divine will/divine command/divine motivation moral theories) depends on the assumption that moral realism is false.[20] Given non-realism, there are no independent moral facts, e.g. X is good and Y is bad, that dictate what we should prefer. Moral realist facts, if they exist, are facts that at least partly determine the correctness/fittingness of *God's preferences*. Given non-realism, there are no independent moral facts in terms of which God's preferences (and God's preferences regarding our preferences) could possibly be mistaken. Any divine will/divine command/divine motivation moral theory is a non-realist moral theory according to my definition of realism. Such theories deny that there are any moral facts that are independent of God. Divine will/divine command/divine motivation moral theories *must deny* the existence of independent moral facts that shape God's commands/will/motivations and to which God's preferences must conform on pain of error. The case for divine will, divine preference, divine command, divine motivation moral theories depends on the view that moral realism is false, i.e. divine will/preference/command/motivation theories commit us to denying that there are any moral facts of this sort. Proponents of the DCT need to say that actions are right because God commands them (and deny that God commands right actions because they are right). Defenders of the divine preference theory of value need to say that things are good because God desires/prefers them (and deny that God prefers/desires those things because they are good). It would be far beyond the scope of this article for me to defend moral non-realism here, although I have defended it at length elsewhere.[21]

Conclusions

If successful, my arguments in this article show that arbitrariness is not nearly as serious a problem for divine will moral theories as is usually supposed. Such theories do not need to claim that God has reasons to will everything that he wills or command everything that he commands. However, it is still important that it not be the case that God has compelling reasons *not to will as he wills*. If my answer to the arbitrariness objection succeeds, then Adams and Zagzebski can avail themselves of it, and the possible deficiencies in their own attempts to deal with the objection noted earlier are not serious objections to their theories.[22]

References

ADAMS, ROBERT (1973) 'A modified divine command theory of ethical wrongness', in Adams (1987a), 97–122.
   (1979) 'Divine command metaethics modified again', in Adams (1987a), 128–143.
   (1987a) *The Virtue of Faith* (Oxford: Oxford University Press).
   (1987b) 'Moral arguments for theistic belief', in Adams (1987a), 144–163.
   (1999) *Finite and Infinite Goods* (Oxford: Oxford University Press).
ARRINGTON, ROBERT (1989) *Rationalism, Realism, and Relativism* (Ithaca: Cornell University Press).
ARISTOTLE (1962) *Nicomachean Ethics*, Martin Ostwald (tr.) (Indianapolis: Bobbs-Merrill).
BRINK, DAVID (1989) *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press).
BROWN, ROBERT (1987) *Analyzing Love* (Cambridge: Cambridge University Press).
CARSON, THOMAS (2000) *Value and the Good Life* (Notre Dame IN: University of Notre Dame Press).
FRANKENA, WILLIAM (1973) *Ethics*, 2nd edn (Englewood Cliffs NJ: Prentice Hall).
HAMLYN, DAVID (1989) 'The phenomena of love and hate', in Alan Soble (ed.) *Eros, Agape, and Philia* (New York: Paragon House), 218–234.
HARE, JOHN (2007) *God and Morality* (Oxford: Blackwell).
HARE, R. M. (1989) 'Ethics and ontology', in Hare, *Essays in Ethical Theory* (Oxford: Oxford University Press), 82–98.
MANN, WILLIAM (1998) 'Piety: lending a hand to Euthyphro', *Philosophy and Phenomenological Research*, **58**, 123–142.
MCTAGGART, JOHN MCTAGGART ELLIS (1927) *The Nature of Existence*, II (Cambridge: Cambridge University Press).
OUTKA, GENE (1972) *Agape: An Ethical Analysis* (New Haven: Yale University Press).
PINKER, STEVEN (2008) 'The moral instinct', *New York Times Magazine*, 13 January.
PLATO (2000) *Euthyphro*, in *The Trail and Death of Socrates*, 3rd edn, G. M. A. Grube (tr.) (Indianapolis: Hackett).
RACHELS, JAMES & RACHELS, STUART (2007) *Elements of Ethics*, 5th edn (New York: McGraw Hill).
SHAFER-LANDAU, RUSS (2010) *The Fundamentals of Ethics* (New York: Oxford University Press).
SIDGWICK, HENRY (1966) *The Methods of Ethics*, 7th edn (New York: Dover).
SOBLE, ALAN (1990) *The Structure of Love* (New Haven: Yale University Press).
SULLIVAN, STEPHEN (1993) 'Arbitrariness, divine commands, and morality', *International Journal of Philosophy of Religion*, **33**, 33–45.
   (1994) 'Why Adams needs to modify his divine-command theory one more time', *Faith and Philosophy*, **11**, 72–81.
TAYLOR, GABRIELE (1975–1976) 'Love', *Proceedings of the Aristotelian Society*, **86**, 153–54.
ZAGZEBSKI, LINDA (2004) *Divine Motivation Theory* (Cambridge: Cambridge University Press).

Notes

1. Among the philosophers who endorse this argument are Rachels & Rachels (2007, 54–58) and Shafer-Landau (2010, 60–64). William Frankena says that Euthyphro's view is 'hard to hold' and that 'Euthyphro's answer to Socrates' question seems to be the natural one, and it implies that what is right is so independently of whether God commands it or not' (Frankena (1973), 29). See Sullivan (1993) for a detailed and helpful discussion of the arbitrariness objection.
2. The Greek word '*hosion*' is usually translated as 'piety'. Some contend that it has a broader meaning than the English word 'piety'. According to G. M. A. Grube, *hosion* also means something like

'righteousness' or 'justice' (Plato (2000), 1). William Mann dissents from this translation and claims the debate in the *Euthyphro* is a debate about the specific religious virtue of piety, not a discussion of any general moral notions like rightness or justice (Mann (1998), 123–124). I venture no view on the scholarly question of the meaning of *hosion*. For my purposes, all that is important is that Socrates' argument *can be* extended to generate objections to theories that make moral properties dependent on God's will or commands. Whether or not Socrates/Plato had anything like this in mind, arguments very similar in form to his have been formulated as objections to divine will/divine command moral theories.

3. Monotheists and atheists who reject polytheism would claim that A and B are both false (since A and B seem to imply that a number different gods exist). At the very least, monotheists and atheists would want to say that A is true only if it means something like the following: 'if it is the case that what is pious is coextensive with what the gods love, then it must be the case that the gods love what is pious because it is pious'. As I argue below, this is not a serious problem for Socrates' argument – see below and endnote 4.

4. So, this argument works, even if we don't accept Euthyphro and Socrates' view that many gods exists and that in their loves and hates they conform to independently existing standards of piety. This is important because very few of us accept this assumption – see endnote 3.

5. Adams (1973), 100. Later, Adams gives a simpler statement of his theory as the view that 'X is ethically wrong' means 'X is contrary to the commands of a loving God' (Adams (1973), 101), but he says that he prefers the earlier formulation from p. 100 just quoted.

6. Adams doesn't mention or speak to the Euthyphro problem in this paper, but clearly his theory permits him to give a perfectly good answer to this question.

7. For a more critical view of Adams's first two theories and their ability to avoid the arbitrariness objection, see Sullivan (1994).

8. Adams (1999), 254. Adams's theory would be viciously circular if he made justice depend on rightness or wrongness (obligatoriness), which, according to his theory, depend on the commands of a loving and just God.

9. Zagzebski (2004), 140–141. She stresses that exemplars only fix the reference of moral terms and she leaves open whether the correct theory about the nature of moral properties is an exemplarist theory. She says that her

> definitions are neutral on the metaphysical question of whether the *phronimos* [exemplar] has *phronesis* because of the morally good properties of his acts or whether his acts have positive moral properties because they are done by a person with *phronesis*. As indexical definitions, they require that we be able to pick out the *phronimos* independently of evaluating his acts, but there is no commitment to the idea that our basis for doing that tells us anything about what is metaphysically basic. (Zagzebski (2004), 143)

This qualification only holds for the naturalistic version of her theory, not the full theological version of the theory that she states later. In her theological theory, the ultimate moral property is the resemblance to the motives of God.

> With God as the ultimate exemplar, we can give a more complete answer to the question of what makes a paradigmatically good person's emotions good ones. The answer given in motivation-based virtue theory is that those emotions fit their intentional objects. The more complete answer is that those emotions imitate the emotions that God has toward those objects. All moral properties are ultimately grounded in the emotional experiences of God. (Zagzebski (2004), 226)
>   God is good in the same way that the standard meter stick is one meter long. God is the standard of goodness. The aspect of God from which God's goodness diffuses is His motives . . . God's motivations are the ground of all value . . . (Zagzebski (2004), 285)

10. I am not sure I should grant that this is possible. Some things seem to be deeply arbitrary. Why are there nine planets instead of ten? Why do we have five fingers on each hand instead of six? How can God's decision to make the universe in these ways have been motivated by his love? Why can't God sometimes be motivated by something other than love? Why can't God have whims and arbitrary preferences so long as they aren't contrary to his being loving?

11. Many philosophers and theologians have proffered analyses of 'love'. Almost all of these analyses imply that at least a *necessary* condition of someone's loving another person is that she desires (or seeks to promote) the good of the other person. Aristotle defines *philia* (love, friendship) as a mutual feeling of goodwill. According to Aristotle, two people love each other provided that each desires the other's welfare for its own sake and each is aware of the other's good will (Aristotle (1962), 1155b28–1156a5). David Hamlyn says that loving someone entails wishing her well, desiring her welfare (Hamlyn (1989), 224). Gabriele Taylor claims that a necessary condition of one's loving another person is that one wants to benefit her (Taylor (1975–1976), 153–154). Sidgwick holds that love always involves 'a desire to do good to the object beloved' (Sidgwick (1966), 244). Robert Brown says that loving someone 'implies wishing to benefit the person and advance the person's welfare' (Brown (1987), 30); also see (Soble (1990), 257–258). Gene Outka characterizes *agape* (love) as: '[A]n active concern for the neighbor's well-being which is somehow independent of particular actions of the other' (Outka (1972), 260).

    McTaggart's definition of love is a notable exception to this. McTaggart defines love as a 'liking' for another that is intense and passionate (McTaggart (1927), 148). Although he says that we generally desire to do good to any person we love, he denies that benevolence (or desiring to do good to the other) is a necessary condition of love. He goes so far as to say that it is possible to desire the ill-fare (bad) of those we love (McTaggart (1927), 150). It's not clear that McTaggart's conception of love can serve the purpose that the requirement that God be loving serves in Adams's theory – restricting the kinds of preferences that God can have so that God's preferences can provide the basis of a plausible normative theory.

12. My divine preference theory of value can be extended to provide a divine preference theory of right and wrong, but endorsing the divine preference theory of value doesn't commit me to endorsing a divine preference theory of right and wrong. My divine preference theory of value is perfectly consistent with a DCT of right and wrong.

    It would be useful and instructive to formulate a divine preference theory of right and wrong. I think that this theory is much more defensible than Adams allows. Here is a first stab at a formulation of a divine will/divine preference theory of obligation. Since we can define obligation and permissibility in terms of wrongness (an act is obligatory if, and only if, it is wrong not to perform it and an act is permissible if, and only if, it is neither wrong to perform it nor wrong not to perform it), we only need to explain what it is for an action to be morally wrong. Consider this:

> It is morally wrong for S to do X provided that: 1. on balance, God prefers that S not do X; 2. on balance, God prefers that other people prefer that S not do X, and 3. one or both of the following: A. God prefers that S feel guilt and remorse for doing X and God prefers that other people resent and disapprove of S for doing X, or B. God feels angry with S for doing X.

In the case of supererogatory actions, condition 3 won't be satisfied. Therefore, this version of the divine preference theory of right and wrong allows us to make a satisfactory distinction between what is obligatory and supererogatory and thus answers Adams's objection that divine will theories can't give an adequate account of this distinction (see Adams (1999), 258–62).

    One reason to favour a divine preference/divine will theory of obligation is that there are serious problems created by saying that moral obligations derive from God's *commands*. The notion of a command is difficult to explicate – I take it that a command is some kind of performative act. Typically proponents of the DCT take the theory to be a straightforward objectivist theory according to which there is a single set of objective moral standards deriving from God's commands that apply to everyone. However, because it appears that God's commands have not been successfully communicated to all human beings, it is dubious to claim that God has issued a single set of commands to all human beings. I take it that if God's commands create moral obligations for me to do X, then God must have commanded *me* to do X. What conditions must be met for us to say that God has commanded *me* to do X? Must God's command be communicated to me? Must I believe that God has commanded me to do X? Can we say that in issuing the Ten Commandments to Moses, God has commanded *all human beings* to obey the Ten Commandments? I don't think that this should count as commanding a person who has never heard of the Ten Commandments or even conceived of a monotheistic God. More controversially, I don't see how this can count as a command to an atheist who has heard of the Ten Commandments, but doesn't believe that they came from God. (The natural law tradition holds that God writes the moral law on our hearts in our individual consciences. However, given how notoriously fallible the human conscience can be, and, given

the many horrible acts people have performed in good conscience, this is a very problematic view.) Adams is aware of this problem (the problem that commands need to be communicated to people in order to be binding on them). He claims that some divine commands are revealed to people in the moral codes and social requirements of societies, even if the people in question don't believe in God and don't recognize the requirements as divine commands. Adams also says that it is likely that some of God's commands are addressed to restricted groups of people and not all human beings (Adams (1999), 262–270).

 I think that divine preference theories of obligation can more easily make sense of the idea that God's will creates objective universal standards of obligation/right and wrong binding on all people than divine command theories. However, I don't want to go on record as endorsing the divine preference theory of right and wrong. I am strongly moved by Adams's argument that obligations require the existence of actual demands made on us by other persons (see Adams (1999), 232–233 and 242–346). Divine command theories of right and wrong make much more sense of this than divine preference theories of right and wrong. But, in any case, the divine preference theory of right and wrong sketched above is more promising than Adams supposes such theories can be, and it merits further examination and discussion.

13. This parallels Adams's statement of his theory in Adams (1979).
14. Carson (2000), 250. Zagzebski quotes this passage (Zagzebski (2004), 266).
15. Zagzebski (2004), 267. She makes a similar argument about Adams's DCT (see the passage from Zagzebski (2004), 261, quoted above in the section on Zagzebski).
16. God's decision to create human beings, rather than other sorts of rational beings, may be arbitrary. Maybe God just prefers us to other kinds of alternative rational persons without having any reason for that preference. God's creation of human beings raises the question of how an omnipotent God could be loving/kindly, given that he chose to create such imperfect creatures who are so liable to suffering, misfortune, and moral vices.
17. Adams agrees that different possible loving Gods could have very different preferences. See Adams (1987b), 148–9, and Adams (1999), 255.
18. The preceding two paragraphs closely follow Carson (2000), 243–244.
19. If moral realism were true and certain things were good or bad or right or wrong independently of God's will, then those independent moral facts might conceivably give us reasons not to do as God commands/wills.
20. *Moral realism* is the view that there are moral facts in virtue of which moral judgements are objectively true or false and that these facts are logically independent of the beliefs, attitudes, emotions, or preferences of rational beings and independent of the beliefs, attitudes, emotions, or preferences that rational beings would have in hypothetical situations (e.g. the moral beliefs that someone would have if s/he were fully informed about relevant facts). Moral non-realism is the view that there are no independent moral facts. For similar definitions, see Arrington (1989), 179, Hare (1989), 84, and Brink (1989), 17 and 7. For a defence of this definition see Carson (2000), ch. 7.

 Moral realism, as it is usually defined, is incompatible with divine will moral theories, since moral realism holds that there are moral facts independent of facts about God and God's will. Zagzebski defines moral realism more broadly as the view that there are objective moral facts and counts her own divine motivation theory as a non-realist theory, see Zagzebski (2004), 11–15 and 81–82. (Note that if we define moral realism as the view that there are moral facts in virtue of which moral judgements are objectively true or false and that these facts are logically independent of the beliefs, attitudes, emotions, or preferences of *human beings*, then divine will moral theories count as realist theories.)
21. See Carson (2000), ch. 7.

 For a considerably more extended argument defending my divine preference theory of rationality and my divine preference theory of value, see Carson (2000), chs 6–8. Here is a brief outline of that argument. First, I think that 'good' means roughly 'desirable' or 'worthy of desire'. If moral realism were true, we could say that it is appropriate to desire something because it is good – the goodness of something makes it a fitting object of desire. But, since I argue at length that moral realism is false, I can't say anything like this. The only tenable alternative for non-realists like me is to that say that things are good (worthy of desire) because it is rational to desire them. Moral non-realism makes theories of rationality central to morality.

 Given the falsity of moral realism, there is a very strong case for thinking that the best theories of rationality are those developed by various versions of the ideal observer theory and the ideal desire-satisfaction theory of value. Roughly, these theories say that what it is rational for us to desire is what

we would desire if we were fully and vividly informed, or if we possessed all a vivid awareness of all relevant information.

However, these theories of rationality are all fatally flawed. Being fully informed, or even possessing all relevant information, greatly exceeds human capacities. It is impossible for any human being to possess all the information relevant to the morality of the Second World War, or even all information relevant to questions about the goodness or badness of my choice of careers (having all information relevant to those questions would require having a vivid knowledge of what my life and life history would have been like if I had pursued any one of the thousands of different careers I might have chosen). Because of this, counterfactual statements of the sort upon which such theories of rationality depend, e.g. 'if I were fully informed, vividly aware of all relevant information, then I would (would not) prefer career X to career Y', have no determinate truth value, indeed the antecedents of such statements are arguably incoherent. (See Hare (2007), 283 and 288, for an endorsement of this argument.)

My divine preference theory of rationality avoids all of these problems. Being fully informed does not exceed God's capacities and, if the sort of God described by my theory exists, there are determinate facts about God's preferences. Together, the prima facie arguments for the ideal observer/ideal desire theory of rationality and the serious objections to those theories point us strongly in the direction of the divine will theory of rationality. The notion of an ideally rational human being is deeply problematic, but if there exists a God of the sort described by my theory, then God *is* the ideal observer.

22. Thanks to Linda Zagzebski, Aaron Bartolome, Steve Evans, Joe Mendola, Paul Moser, Marcella Russo, and an anonymous referee for helpful comments.