Assessing Arguments from Moral Objectivity for God's Existence

Martin Dunkley Smith

Lady Margaret Hall

DPhil Theology

Acknowledgements

Many thanks to the Arts & Humanities Research Council, the OCCA (Oxford Centre for Christian Apologetics), the University of Oxford, and Lady Margaret hall for their generous funding toward this project,

to my supervisors, Dave Leal and Roger Crisp for their guidance,

to Vince Vitale and Brian Leftow for their belief in me,

to my parents for their sacrificial support

and to my dear wife, Sarah, who reminds me that thought is but a shadow of the real

[Abstract]

Most moral arguments for theism claim that the best answers as to whether objective moral properties are reducible, what they might reduce to, whether their instantiation is explainable, and what might do the explaining, are provided by theological voluntarism. Theists have two major arguments for preferring the reductive version of voluntarism to the reduction of moral properties to natural properties. The first holds that if moral properties reduce to natural properties, there must be a particular conceptual connection between moral concepts and natural concepts. The charge is that this connection is absent. The second holds that if moral properties reduce to natural properties, moral properties must supervene on natural properties in the right way. The charge is that they do not so supervene. Theists also have a prominent argument for preferring the causal version of voluntarism to non-reductionist views of moral properties. It holds that voluntarism provides a deeper and simpler explanatory stopping point.

I argue that all three arguments fail to show an advantage for voluntarism. There is no conceptual connection that is plausibly a requirement on reduction that excludes natural properties and permits theistic properties. Theists cannot show that natural properties fail to satisfy the supervenience requirement without leaving voluntarism vulnerable to a powerful objection. And causal voluntarism requires positing at least one brute moral fact about God's moral authority. Non-reductive views can also posit one brute fact – one about the respect due to persons. And attempts to show that voluntarism's one brute fact has greater internal simplicity fail. Since theists have not shown voluntarism to have advantages over particular competitors, they have not shown voluntarism to have advantages over *all* competitors. Hence it is not been shown that voluntarism provides the best answers to questions about the reduction and explanation of moral properties.

Contents

1	**Introduction (Extended Abstract)**	5
2	**Assessing Moral Arguments**	13
2.1	Types of Moral Argument for Theism	
2.2	Objectivity	
2.3	Theism & Atheism	
2.4	The Necessity of Comparison	
2.5	Restricted Theological Voluntarism	
3	**Accommodating Moral Properties**	37
3.1	Reduction	
3.2	Explanation	
3.3	Surveying the Options	
4	**The Conceptual Gap Argument: Interpretation**	61
4.1	Introducing the Major Objection	
4.2	The Conceptual Gap Argument	
4.3	Textual Evidence	
5	**The Conceptual Gap Argument: Assessment**	85
5.1	Compound-to-Constituent Connection	
5.2	The Neo-Classical Picture	
5.3	Role-Satisfaction: Harrison	
5.4	Role-Satisfaction: Adams	
5.5	A Modest Argument	
6	**The Divergence Argument**	123
6.1	Interpreting the Argument	
6.2	Normative Commitments	
6.3	The Making Relationship	
6.4	The Counterfactual Argument	
7	**The Explanatory Depth Argument**	164
7.1	Causal Voluntarism	
7.2	Brute Realism	
7.3	The Explanatory Depth Argument	
7.4	No Normative Ex Nihilo	
7.5	Irreducible Rights	
7.6	Brute Rights	
8	**Conclusion**	208
9	**References**	215

1

Introduction
(Extended Abstract)

Moral arguments for theism attempt to show that the best account of some aspect of morality commits us to the existence of God. This thesis assesses only one type of moral argument for theism, those I call *metaphysical* moral arguments for theism. The aspect of morality they focus on is the existence and objectivity of moral properties. These arguments assume that a good independent case can be made for the existence and objectivity of moral properties. They then claim that the best answers we can give to whether these properties are reducible, what they might reduce to, whether their instantiation is explainable, and what might do the explaining, commit us to theism. Hence these arguments claim that there is a route from the objectivity of morality to theism.

The judgment this thesis renders on arguments of this type is that they fail, at least in their most prominent contemporary versions. These versions claim that the best answers to the above questions about objective moral properties, reduction, and explanation, are provided by divine command theories. In particular, they are provided by the kind of divine command theories developed by Robert Adams (1979b; 1999) and Philip Quinn (1978). Following Murphy (2012), I will label such theories *restricted theological voluntarism*. Voluntarism of this sort comes in a reductive version and a causal version (I elaborate on these differences in the next chapter).

Theists have two major arguments for preferring reductive voluntarism to the reduction of moral properties to natural properties. The first, more central argument holds that if moral properties reduce to natural properties, there must be a particular conceptual connection

between moral concepts and natural concepts. The charge is that this connection is absent. The second argument holds that if moral properties reduce to natural properties, moral properties must supervene on natural properties in the right way. The charge is that they do not so supervene. Theists also have a prominent argument for preferring the causal version of voluntarism to non-reductionist views of moral properties. It holds that voluntarism provides a deeper and simpler explanatory stopping point.

I argue that all three arguments fail to show an advantage for voluntarism. There is no conceptual connection that is plausibly a requirement on reduction that excludes natural properties and permits theistic properties. Theists cannot show that natural properties fail to satisfy the supervenience requirement without leaving voluntarism vulnerable to a powerful objection. And causal voluntarism requires positing at least one brute moral fact about God's moral authority. Non-reductive views can also posit one brute fact – one about the respect due to persons. And attempts to show that voluntarism's one brute fact has greater internal simplicity fail.

Since theists have not shown voluntarism to have advantages over particular competitors, they have not shown voluntarism to have advantages over *all* competitors. Hence it is not been shown that voluntarism provides the best answers to questions about the reduction and explanation of moral properties. Restricted theological voluntarism may not be worse than its rivals, but it is not better, and so it cannot be the *best* account we have. Metaphysical moral arguments for theism that utilise these divine command theories need them to be the best, and so these arguments fail.

There are good reasons to assess these moral arguments for theism. Evans states that moral arguments for theism "are important because of their prominence in popular apologetic arguments for religious belief (2018)." While he is speaking of the wide variety of moral arguments, what he says is particularly true of the metaphysical version. As well as being defended at the scholarly level by J P Moreland & William Lane Craig (2003), C Steven Evans (2013), David Baggett & Jerry Walls (2011; 2016), Robert Adams (1979b; 1999), and others besides, several of these philosophers have defended the argument at a popular level. None have done more than William Lane Craig in this regard, who in addition to his scholarly material, has promoted the argument in popular level debates, articles, books, and videos. The argument seems to have persuasive force; Craig remarks, "in my experience of speaking to university students, the moral argument I present is probably one of my most persuasive arguments, eliciting a positive response among students (Baggett, Walls, Copan, Craig, 2011)."

In my own experience, the metaphysical moral argument looms large over the Christian apologetics community. This is the community of Christians who intellectually defend their religious beliefs in discussions with atheists or agnostics. This is a wider community than Christian scholars in philosophy of religion. While some members of this community are scholars, most are not. Most do not get their contact with theistic or atheistic arguments from scholarly-level books and articles in philosophy of religion.[1] In my experience, members of this community have endorsed the metaphysical moral argument for theism more widely and confidently than any other argument for theism. It has been recirculated across social media, blogs, and podcasts, has been presented to undergraduates in short talks around the globe, and is sometimes viewed as a silver bullet argument for theism.

[1] This is not intended as a criticism of this community. This is just a factual statement about how the community extends far beyond specialists in philosophy of religion.

So, one reason to assess the metaphysical moral argument is its influence among educated non-philosophers. Indeed, this influence makes the relatively small size of the scholarly literature on the argument striking. It is no surprise that when Brian Leftow (Leftow, Anderson & Schellenberg, 2015, p.442) suggests areas of philosophy of religion that ought to be developed, he picks out metaphysical moral arguments as one such area. And one reason to assess the version of them that relies on restricted theological voluntarism is that it is the most prominent version defended by philosophers. Indeed, all the scholars that I listed above as defending the metaphysical moral argument do so by appeal to restricted theological voluntarism. But the popularity of this voluntarism is no mere fad. As I explain in the next chapter, this voluntarism is well equipped to respond to the Euthyphro Dilemma. And Adams' particular formulation of it makes sophisticated use of developments in philosophy of language wrought by Kripke (1980) and Putnam (1973). Thus, voluntarism of this sort is popular because it is a serious contender in contemporary meta-ethics.

As I said earlier, my assessment of metaphysical moral arguments for theism is that they fail. But of course, negative assessments of these arguments and the voluntarism they rely on are not new. What need is there to offer another negative assessment? Consider that most criticisms of voluntarism target its independent merit as a meta-ethical theory. They do not target its role in motivating a metaphysical moral argument for theism. Of course, if voluntarism is independently implausible then that counts against its use in a metaphysical moral argument for theism. But there can be criticisms of its use that are not criticisms of its independent plausibility. Metaphysical moral arguments attempt to show that voluntarism is the *best* account of objective moral properties. The claim is that it has advantages over its non-theistic competitors. In response, one can object that no such advantages have been shown. That is my

focus in this thesis. I do not argue that voluntarism is implausible, just that attempts to show that it is the best account fail. My assessment focuses on the comparative aspect.

In particular, if proponents of theistic moral arguments need to show that voluntarism does a better job than non-theistic accounts, then the first thing we should check is whether the objections they pit against non-theistic accounts rebound against voluntarism. That is, forget about the objections that distinctly target voluntarism, like the Euthyphro dilemma. The first hurdle is to show that voluntarism does not face the very same problems that non-theistic views are accused of having. And one needs to make sure that the premises of arguments attacking non-theistic views do not have entailments that cause other problems for voluntarism. Too little attention has been given to this task. I remedy this by focusing my comparative assessment on these initial hurdles.

My comparative assessment can also be distinguished from the kind that Wielenberg offers. He argues that voluntarism cannot be shown to improve upon the supposed errors of his own non-reductive view of moral properties (2009; 2014). What he leaves unaddressed is whether voluntarism improves upon the supposed errors of views in which moral properties reduce to natural properties – views I call natural reductionism.[2] No critic has said much on this score. This is unfortunate since natural reductionism remains a prominent non-theistic account. Proponents of metaphysical moral arguments for theism devote a lot of attention to critiquing it. And at the popular level, it is often considered the only non-theistic account. Further, all else being equal, natural reductionism is the most attractive view for non-theists. It allows one to accommodate moral properties using the natural properties one already acknowledges to exist, rather than positing new irreducible properties.

[2] The more standard name for these views is 'ethical naturalism'.

This thesis addresses this gap in the literature by focusing more on comparing voluntarism with natural reductionism than with non-reductionism. Further, I interpret the objections theists have to natural reductionism with more detail than any other interpreter so far. Perhaps due to lack of prior critical engagement, these objections have tended to remain at a rough, intuitive level. I make these objections explicit, arguing that the most important objection bears similarities to an objection against reductionism in the philosophy of mind. David Chalmers (1996) is famous for arguing that there is a *conceptual gap* between phenomenal concepts and physical concepts, a gap that precludes the reduction of phenomenal states to physical states. I argue that, likewise, theists claim that there is a conceptual gap between moral concepts and natural concepts, precluding natural reductionism.

By seeing parallels between arguments against reductionism in the philosophy of mind and argument against reductionism in meta-ethics, I join a recent body of work that explores such parallels. Bedke (2012), Yetter-Chappel & Chappel (2013), and Mehta (2017), have all argued that one can abstract from the features of anti-reductionist arguments in specific fields to see a basic re-occurring argumentative structure, one that identifies supposed conceptual gaps and draws metaphysical conclusions from them. My contribution to this body of work is in arguing that proponents of metaphysical moral arguments for theism are implicitly doing the same. This is novel: nobody has yet brought the philosophy of mind parallels into conversation with theistic meta-ethical accounts.

Interpreting theists' major argument as appealing to conceptual gaps make sense of the textual evidence in their writings. Speaking personally, it also seems to capture what was in my own mind when I endorsed metaphysical moral arguments for theism. There was a time when this

argument for theism struck me as extremely strong. While I remain a theist, as I have thought about the argument over the last decade, I have become increasingly dissatisfied with it. This thesis is an attempt to explain what now dissatisfies me about it. I hope that in my position as a theist and a former proponent of the argument, my interpretation and assessment of it will be sensitive to its force while being incisive as to its shortcomings.

The thesis will take the following route. In §2 I distinguish metaphysical moral arguments for theism from other moral arguments for theism. I also unpack some key concepts that metaphysical moral arguments use and defend the importance of truly comparing voluntarism with its non-theistic rivals. It is not enough to poke holes in non-theistic theories, one must show voluntarism to do the better job. In §3 I develop some of the metaphysical background that needs to be in place in order to assess these moral arguments for theism. In particular, I elaborate on what reduction and explanation amount to.

Next I move on to the arguments for preferring restricted theological voluntarism to its non-theistic competitors. In §4 I interpret the major argument from preferring Adams' reductive voluntarism to natural reductionism. This is the Conceptual Gap argument that parallels Chalmers' argument in philosophy of mind. In §5 I assess this argument and show that there is no account of what the relevant conceptual gap amounts to that affords voluntarism an advantage over natural reductionism.

In §6 I interpret and assess the other major argument for preferring reductive voluntarism to natural reductionism. The Divergence Argument holds that moral properties do not track or supervene upon natural properties in the way that their reducing to natural properties would require. I argue that motivating this argument would require theists taking on significant

commitments in first-order normative ethics. Further, defending such commitments would leave voluntarism vulnerable to an important objection, which I call the Counterfactual Argument.

In §7 I move on from reductive voluntarism to consider whether Quinn's causal voluntarism has advantages over non-reductive non-theistic views. The Explanatory Depth argument claims that causal voluntarism offers a more elegant explanatory stopping point. I argue, however, that in light of the need for causal voluntarism to posit a brute moral fact about God's rights to command, no such explanatory advantage can be retained. In §8 I tie everything together to show that, at present, we should conclude that metaphysical moral arguments for theism fail.

2

Assessing Moral Arguments

In the introduction, I picked out metaphysical moral arguments for theism as the arguments this thesis is concerned with. These arguments claim that the best answers we can give to questions about the reduction and explanation of moral properties commit us to theism. In this chapter, I distinguish these arguments from other moral arguments for theism. I also elaborate on the structure of metaphysical moral arguments and some key concepts they use. The concepts covered are *restricted theological voluntarism*, the *objectivity* of moral properties, *theism*, and *atheism*. After elucidating the argument, I vindicate the need for a comparative assessment of how theism and atheism answer questions about reduction and explanation. For I argue that even those metaphysical moral arguments that seem to bypass comparison, like William Lane Craig's, are forced to offer comparison within the dialectical context.

§2.1 TYPES OF MORAL ARGUMENT FOR THEISM

Many philosophers believe that some actions or attitudes have moral properties. I will usually focus on actions, as they are the focus of most moral theorising, but much of the same could be said for attitudes like intentions, desires, beliefs. Some actions are morally obligatory to perform, like keeping one's promises. Others are morally impermissible to perform, like torturing innocents. These properties are *deontic* moral properties. They concern what is morally right and wrong – what we are obligated to do or refrain from doing. Note that while in ordinary language, talk of what is 'morally right' is ambiguous between what is merely permissible and what is obligatory, I will always use it in this thesis to refer to what is obligatory.

Some actions are morally good to perform, like donating your bonus to charity. Others are moral bad to perform, like lying. These properties are *axiological* moral properties. They differ from deontic properties in that they concern moral values, not obligations. At least, if something is morally *good* it does not follow that it is morally obligatory. You could perhaps spend your bonus on a holiday for yourself rather than giving it to charity without doing anything morally wrong. (It is more controversial as to whether something could be morally bad without also being morally wrong).

There are other moral properties, but many philosophers consider the deontic (obligation-focused) and axiological (value-focused) to be the central ones. Further, these properties are what moral arguments for theism typically focus on. So this is sufficient background to begin demarcating different types of moral arguments. The argument types can be distinguished along the divisions between different sub-fields of moral philosophy. There are, firstly, distinctions between moral arguments for theism from practical ethics, from normative ethics, and from meta-ethics.

Practical ethics, as a field of study, looks at particular actions (or attitudes) and tries to figure out what moral properties, if any, those actions have. One might try to figure out whether abortion is morally wrong, or whether donating to a particular charity is morally good. A moral argument for theism in practical ethics attempts to show that it would be morally good or obligatory (at least in some circumstances) to believe theism. There are not many philosophers who have developed arguments of this type. But there are some examples; Mark Wynn (1997) has argued that due to the trusting relationships some children have with their parents, and the

moral goodness of epistemic partiality, it is morally good for these children to adopt the theistic beliefs of their parents (when the evidence against these beliefs is not too strong).

Normative ethics attempts to figure out which properties *make* actions right or wrong. Perhaps ultimately, there is just one property that makes an action right, like maximising happiness (as utilitarianism claims). But there could also turn out to be multiple. Normative theories are projectible; when all the non-moral facts are known it should tell us how moral rightness and wrongness are distributed across actions. For example, if maximising happiness is what makes actions right, then once we know the non-moral facts about which actions have maximised happiness, we know which actions have been right. Of course, we also have pre-theoretical intuitions about how rightness and wrongness are distributed. And the predictions of a theory can conflict with these intuitions. If a theory predicts that torturing innocents for fun is morally right, that is a mark against the theory since intuition says otherwise.

The possibility of divergence between pre-theoretical intuition and prediction opens up a couple different kinds of argument for theism from normative ethics. Both involve claiming that a normative theory only makes predictions that cohere sufficiently with pre-theoretical intuition if God is introduced. The first kind claims that prediction fails unless a theistic property is posited as one of the properties that makes an action right or wrong. The second kind claims that prediction fails unless a theistic property, or some fact about God, is posited among the non-moral facts. Presently, the most popular normative arguments for theism claim that prediction about which kinds of being it is wrong to harm fail unless we posit some beings as *made in the image of God,* or *loved by God* (Baggett & Walls, 2016, pp.115-144; Wolterstorff, 2008, pp.115-144). Such theistic properties make actions wrong or should be included in the non-moral facts.

I say more about normative ethics and the making relationship in §6. For now, note that I do not presume that the making relationship has anything to do with reduction, or anything to do with the kind of explanation that is of concern to metaphysical moral arguments for theism. Thus, normative arguments for theism should not be equated with metaphysical moral arguments for theism. To claim that a theistic property makes an action wrong is not necessarily the same as claiming that wrongness itself reduces to a theistic property. There has sometimes been confusion, especially at the popular level, over how metaphysical moral arguments relate to issues in first-order normative ethics. Does endorsement of one of these arguments require one to, say, oppose utilitarianism? Answers to these questions have not always been clear.

The reason why this has been unclear is, I think, because different accounts of the making relationship yield different answers to those sorts of questions. There are accounts of the making relationship that would erase a distinction between normative and metaphysical theistic moral arguments, but others preserve it. And as I argue in §6, most proponents of metaphysical moral arguments implicitly assume that these arguments do *not* carry substantial commitments in normative ethics. The intuitive thought is probably this: normative ethics concerns what our obligations *tell us to do,* whereas moral metaphysics concerns what obligations *are,* or how they are brought about. Claiming that obligations are divine commands or brought about by God's will does not tell us what our obligations demand of us. For all God's involvement, they may very well tell us to maximise happiness. I will assume, then, that these two kinds of theistic argument are distinct, at least until accounts of the making relationship are examined later on.

Moving on, meta-ethics, as a field of study, covers investigation into the semantics, epistemology, rationality, and metaphysics of morality. The semantics of morality concerns the

analysis of moral language and moral assertions. Issues here concern whether, say, assertions like "it is wrong to torture innocents" are truth-apt[3], or whether moral predicates like 'wrong' can be analysed in other terms. There is no contemporary tradition of arguments for theism from moral semantics *alone.*[4] This is not to say that moral semantics is irrelevant to arguments for theism. To the contrary, it is extremely important and much of the discussion in §3 and §4 will focus on it. It is just to say that moral semantics enters arguments for theism only insofar as it bears on other parts of meta-ethics, like moral metaphysics. So, strictly speaking, there are no arguments for theism from moral semantics.

The epistemology of morality concerns whether any beliefs we have about the existence or distribution of moral properties enjoy the status of knowledge, and if so, how such knowledge is achieved. (Or they might be concerned with positive epistemic statuses other than knowledge). Epistemological moral arguments for theism attempt to show that if moral properties are objective in a certain mind-independent sense (more on objectivity below) then the best account of how it is that we can know about such properties or their distribution commits us to the existence of God. Such arguments attempt to remain neutral on whether the very existence or nature of the moral properties themselves is best accounted for by God – they focus only on our epistemic access to those properties. Arguments of this sort are defended by Linville (2009) and Ritchie (2012) and tend to build on Sharon Street's (2006) work on the evolutionary debunking of moral beliefs.

Studies of the rationality of morality concern issues like whether we always have decisive reasons to be moral, or what sort of connections there are between forming a moral judgment

[3] Or whether they express propositions which are true or false.
[4] Adams (1973) once defended the view that what 'morally wrong' means in the moral discourse of Judeo-Christian religious believers is 'commanded by God'. He now rejects this view, but even when he held it, he did not take it to support an argument for theism.

like “it is wrong to torture innocents” and being motivated to act in conformity with it. Immanuel Kant famously formulated arguments for theism concerning moral rationality and these remain defended in some form or other today, most notably by John Hare (1997). The idea here, roughly, is that moral practice commits us to a certain goal, such as the ultimate alignment of happiness and duty. Then the argument is that this goal can only be rational to pursue if we believe that God exists.

With moral metaphysics, we reach the kind of issues this thesis is concerned with. Moral metaphysics seeks answers for questions about whether moral properties are objective, whether they are reducible or not, what they might reduce to, whether their instantiation is explainable or not, and what might do the explaining. As said before, moral metaphysical arguments for theism assume that there is a good independent case to be made for the existence and objectivity of moral properties. They then claim that the best answers we can give to whether moral properties are reducible, what they might reduce to, whether their instantiation is explainable or not, and what might do the explaining, commit us to theism. From here on I will often refer to these metaphysical moral arguments as simply *theistic moral arguments* or just *moral arguments.* And I shall refer to their proponents as *theistic moral arguers* or just *moral arguers.*

For the sake of simplicity, we can call a set of answers as to whether moral properties are reducible, what they might reduce to, whether their instantiation is explainable, and what might do the explaining, an account of how moral properties are to be *accommodated.* The more plausible a set of answers, the more plausible the accommodation account. Different views, like theism or atheism, can provide different answers to these questions and so provide different accommodation accounts. Thus, theistic moral arguments claim that theism has a more plausible accommodation account of moral properties than atheism has. Or put even more

simply, they claim that theism better accommodates moral properties than atheism. We can flesh out the structure of such arguments as follows:

P1. There are objective moral properties.

P2. It is plausible that theism accommodates objective moral properties and more plausible that it does so than that atheism does so.

P3. If P1 and P2, then there is *pro tanto* reason to believe theism.

Therefore,

There is *pro tanto* reason to believe theism.

Consider how this structure covers the wide variety of metaphysical moral arguments. First of all, this structure leaves open which moral properties are being targeted, for moral arguments differ on that score. Craig (2008) argues that theism best accommodates axiological moral properties (goodness and badness) *and* deontic moral properties (rightness and wrongness), whereas C. Steven Evans (2013) focuses just on the deontic. Moral arguments also differ in the degree of superiority they argue theism enjoys. Craig (2008) claims that theism is the only view that plausibly accommodates objective moral properties.[5] Others, like Baggett & Walls (2011; 2016), concede that non-theistic views can accommodate moral properties with some

[5] This seems to be the upshot of Craig's claim that there simply would not be an objective morality if atheism were true.

plausibility, but not as much plausibility as theism can. The above structure is neutral as to just how much of an advantage theism has.

Additionally, moral arguments differ in their explanations of *why* theism has an advantage over atheism. One of those differences is that some, like Adams' version, claim theism has an advantage when it comes to reduction. Others, drawing on Quinn's causal voluntarism, may claim that theism has an advantage when it comes to causal explanation. They also differ in how the argument for the advantage should be *framed*; deductively, abductively, or probabilistically. The above structure is neutral as to what sub-argument for P2 should be given. And the conclusion is neutral as to the *strength* of the pro tanto reason to believe theism moral arguments attempt to offer, from very strong to relatively modest. Let us call this structure the *moral argument schema.* Let us now get clearer on some basic concepts used in the schema.

§2.2 OBJECTIVITY

The first premise [there are objective moral properties] is a shorthand way of expressing a commitment to moral realism, which includes two claims. The first is the metaphysical claim that some moral properties are instantiated. The second commitment of moral realism is that these moral properties are relevantly independent of human mental attitudes. This is their objectivity. Moral realism might also be expressed in terms of the obtaining of objective moral facts rather than in terms of moral properties.[6] I will continue to focus on properties, bringing in facts only when the discussion requires it. But the reader should note that unless otherwise

[6] These would be first-order moral facts rather than any meta-moral facts that might obtain, like *moral claims are not assertoric* and *moral facts are reducible to physical facts.* Error theorists who deny moral facts will still accept some meta-moral facts, so it is important to distinguish the two; error theorists do not get to count as moral realists just because they accept meta-moral facts.

stated the discussion can be applied to facts with small tweaks. Keeping the discussion focused on properties will keep the discussion streamlined, and will mirror the focus of moral arguers, which tends to be on properties rather than facts.

So, the objectivity of moral properties is some sort of attitude-independence. But fleshing out what the relevant sort of attitude-independence that moral realism is committed to is not a trivial undertaking. It has traditionally been thought sufficient to flesh this out as the claim that there are at least some moral properties instantiated that would still be instantiated even if human intentional attitudes or conceptual schemes were different than they actually are. That is, a counterfactual independence from human attitudes is appealed to. Thus, William Lane Craig (2008, p.173) says,

> To say that there are objective moral values is to say that something is good or bad independently of whether anybody believes it to be so. Similarly, to say that we have objective moral duties is to say that certain actions are right or wrong for us independently of whether any human being believes them to be so. For example, to say that the Holocaust was objectively wrong is to say that it was wrong even though the Nazis who carried it out thought that it was right; and it would still be wrong even if the Nazis had won World War II and succeeded in exterminating or brainwashing everybody who disagreed with them so that it was universally believed that the Holocaust was right.

It is a matter of some controversy, however, as to whether this is sufficient to characterise the kind of mind-independence moral realism requires. Some avowedly non-realist views like Simon Blackburn's (1993; 1998) quasi-realism claim to capture this sort of counterfactual independence. My preferred way to flesh out realism's mind-independence is to agree with Christian Miller (2009) that we should add to counterfactual independence a further condition about what *explains* this independence. That is, the explanation of it cannot be that human intentional attitudes and conceptual schemes in the actual world are rigidified and then used to assess the moral status of actions even in worlds where those attitudes and conceptual schemes

are not present. This seems to rule out Blackburn's view counting as realist. Moral arguers should not find this addition objectionable since it makes it harder for atheism to accommodate objective moral properties, not easier.[7]

Whether moral realism is true is, of course, a live and complex debate in meta-ethics that is of interest independently of moral arguments for God's existence. I said earlier that moral arguments assume that a good independent case for objective moral properties – for moral realism – can be made. This suggests that moral arguers do not themselves present any arguments for moral realism. But moral arguers *do*, in fact, present arguments for moral realism. It is just that these arguments do not normally advance the debate over realism. They are the standard arguments for realism that non-realist meta-ethicists have not found persuasive (rightly or wrongly). And that is not surprising. Advancing the debate over realism would leave one with little space to say anything uniquely pertaining to a moral argument for theism. Moral arguments can thus be interpreted as attempts to give either those who are new to meta-ethics reason to believe realism and so ultimately to believe theism, or to give veterans of meta-ethics who are already committed to realism reason to believe theism.

Since moral arguments for theism do not attempt to advance the debate over realism, I will not step into that debate either. I will assume that we have good reason to believe moral realism and concern myself more with whether theism or atheism best accommodates it. Or at least, I grant that we have good reason to believe moral realism prior to assessing the second premise of the moral argument. That caveat is necessary because, as I explain in more detail in §2.3, an assessment of the ability of theism and atheism to accommodate moral realism may, in the end,

[7] Atheists should not find this addition objectionable either. I am arguing in this thesis that even when objectivity is understood so strictly, theistic moral arguments fail.

overturn our reasons to believe moral realism. If it turns out that theism and atheism *both* do an implausible job of accommodating moral realism, then we might have good reason to deny moral realism. But this reason to reject the first premise of the moral argument will not be an independent objection to the argument, since it will depend on the second premise having already been shown false (theism will have been shown not to do a plausible job of accommodating objective moral facts.) What I am granting for the sake of argument, then, is that we have good reasons to believe the first premise and that these are not overturned by any objections that are *independent* of objections to other premises of the argument.

§2.3 THEISM & ATHEISM

The second premise compares theism to atheism. What *is* theism, and so what is atheism denying? If theism is defined narrowly, then atheism will be an expansive view. For instance, suppose theism is defined as the following narrow package, what we might call *western monotheism*:

> Necessarily, one God exists and is perfect in power, knowledge, and moral goodness.
>
> God is distinct from and created the universe.

It will then be the case that atheism allows, say, that there is a being perfect in power, knowledge, and moral goodness for whom the universe is a constituent part. In other words, atheism will be compatible with panentheism. That is a very expansive understanding of atheism. Alternatively, suppose that theism is defined very expansively, such that it is, say, the claim that there is some immaterial agent without a natural cause. Then atheism will be much

more restrictive in scope, and close to naturalism.

Moral arguers tend to have western monotheism in mind as the claim they are arguing for. And yet, the claim they tend to criticise is naturalism. There is very little discussion, if any, as to whether western monotheism has any advantage in accommodating moral realism over panentheism. So, moral arguers tend to defend a narrowly defined theism and object to a narrowly defined atheism. How can they have it both ways? Implicit in extant moral arguments is the assumption that the two live options for contemporary thinking people are western theism and naturalism. This assumption is widespread in contemporary analytic philosophy of religion.[8] The assumption might be true but assessing whether it is would take us far afield into broader metaphysical issues. Thus, I will simply grant the assumption that our options are western monotheism or naturalism. And from here on I will take theism to be equivalent to the former, and atheism equivalent to the latter.

Most moral arguers are also sympathetic to *perfect being theism* as an approach to natural theology. According to perfect being theism, God is perfect; God is the greatest conceivable being, or the greatest possible being. As such, we can infer what properties God has by reflecting on what properties would be perfections in a being, those properties that have the maximal degree of greatness for a being. Through this thesis I will sometimes consider how arguments from perfect being theology bear on moral arguments.

Now to say something more about naturalism. How do moral arguers understand naturalism? It is at least understood as the claim that there are no immaterial agents without a natural cause.

[8] Though the field has, in recent years, been challenged to expand its scope. As Anderson (Leftow, Anderson, Schellenberg, 2015, p.445) says, "We have too often been extremely narrow, thinking within the bounds of Christian theism and avoiding more urgent, less materially comfortable matters to do with religions."

Oppy (2018, p.2) gives a rough definition that seems in spirit with how moral arguers treat it:

> (a) there are none but natural causal entities with none but natural causal powers and (b) well-established science is the touchstone for identifying natural causal entities and natural causal powers.

According to this definition, naturalists could accept abstract objects. They could also accept not just moral properties but irreducible, brute moral properties. Such entities are admissible either because they are natural or because they are non-causal. And moral arguers tend to allow naturalists this sort of leeway. When Wielenberg posits such entities, moral arguers do not spend time objecting that his view is not sufficiently naturalist.

One might, though, want a better grip on what properties count as natural properties. At the very least, they are micro-physical, macro-physical, and human mental properties – both individual and social. These are the ones that moral arguers tend to see as at the naturalist's disposal.[9] We can call these the *core resources* of naturalism. But we should not forget that, unless moral arguers want to debate not just moral properties but normative properties in general, naturalists can also appeal to non-moral normative properties. These include aesthetic properties, epistemic properties and so on. We can call the addition of non-moral normative properties, and any other admissible properties not explicitly mentioned here, part of the *expanded resources* of naturalism.

§2.4 THE NECESSITY OF COMPARISON

[9] Of course, some moral arguers will not agree that human mental properties are in fact natural properties. Perhaps they are the properties of a soul, and this excludes them from being natural. But most moral arguers do not assume that their moral argument succeeds only on the condition that some form of dualism about mind is correct. So they should be happy to allow naturalists to appeal to mental properties as natural properties.

So far, then, I have explained some key concepts used in the moral argument schema. I have explained what is meant by the objectivity of moral properties, and what is meant by theism and atheism (naturalism). Now it is time to comment further on the premises themselves and their relationship to the conclusion.

The argument has a clearly valid form. If the argument is not good it is because at least one of the premises is false, or without support. I have already explained that I will not be assessing the first premise of the moral argument schema; the claim that moral realism is true. I will assume the truth of that premise. So in finding fault with the argument, I must focus on either the second or third premise.

The third premise can be assessed very quickly: it should be accepted. If moral realism is true, and theism more plausibly accommodates it than naturalism does, that does indeed provide a *pro tanto* reason to believe theism. As Adams (1979a, pp.144-5) claims when giving his own moral argument, "One of the most generally accepted reasons for believing in the existence of anything is that its existence is implied by the theory that seems to account most adequately for some subject matter." If the best answers to whether moral properties are reducible, what they might reduce to, whether their instantiation is explainable, and what might do the explaining, entail theism, there is reason to believe theism. Therefore, the only premise left in need of assessment is the second premise:

> P2. It is plausible that theism accommodates objective moral properties and more plausible that it does so than that atheism (naturalism) does so.

This premise is not self-evident, it stands in need of motivation. In order to motivate it, moral

arguers need to present some accommodation account – some set of answers to the questions about reduction and explanation – and show that it is plausible and beats its naturalistic competitors. There is, however, an objection one might have to moral arguers really being committed to P2. That is, despite the wide coverage of the moral argument schema, one might object that it still excludes some arguments that, intuitively, ought to count as metaphysical moral arguments for theism. One might object that not all moral arguments feature *comparisons* between theism and atheism. For Craig's (2008, p.172) moral argument seems not to. He structures his argument as follows:

1. If God does not exist, then objective moral values and duties do not exist.
2. Objective moral values and duties exist.

   Therefore,
3. God exists.

This argument does not seem to involve anything like the second premise of my moral argument schema. It does not seem to involve any *comparison* at all. Its first premise states an alleged entailment of atheism. But there is no contrasting premise about what *theism* entails. The argument seems to run on a purely negative critique of atheism. In response, Craig's argument may not match the schema on the surface, but I will argue that given the dialectical context of his argument, he will need to make moves that ultimately collapse his argument into one that fits the schema.

Consider that Craig's first premise, the conditional claim about what follows for moral realism if God does not exist, is not self-evident. So it requires supporting argumentation. And Craig provides some. Suppose, however, that the kinds of considerations that support the first premise

also support:

4. If God *does* exist, then objective moral values and duties do not exist.

That is, suppose that the alleged deficiencies that atheism has in accommodating moral realism are also deficiencies that theism has. In which case, there is just as much support for Craig's moral argument as there is for a parallel atheist argument that derives from (4), and the claim that moral realism is true, the conclusion:

5. God does not exist.

In that situation, we would have two options. We might hold fast to there being an objective morality but be agnostic as to either atheism or theism really fail to accommodate it. And so the moral argument for God's existence would not work for there would be no rational reason to prefer its conditional premise over the conditional premise of the atheist argument. Or alternatively, we might hold that both theism and atheism entail that there is no objective morality and so, since theism and atheism exhaust the options, conclude that there is no objective morality after all. This would be the equivalent of *double incompatibilism* in the free will debate - the view that free will is incompatible with both determinism and indeterminism. Here the double incompatibilism would be in holding that moral realism is incompatible with both atheism and theism. And obviously Craig's moral argument for God's existence gets no grip on moral nihilists who reject its second premise. So either way Craig's moral argument would be in trouble.

The lesson is that it is important to show that theism does a *better* job than its competitors of

accommodating an objective morality. Pointing out deficiencies in competitors does not help if theism has the same deficiencies. Or even if theism does not have the *same* deficiencies, it is important to show that it does not have other, equally detrimental deficiencies. The Euthyphro dilemma, for instance, is thought to uniquely target some theistic attempts to accommodate an objective morality. It is no good if those theistic attempts avoid the deficiencies of competitors only to have the Euthyphro dilemma show up their own equally serious deficiencies. So in dialectical context, Craig must offer some *comparison* of theism and atheism. And in practice he does so, by defending restricted theological voluntarism as preferable to any atheist-compatible alternative (Kurtz & Craig, 2009).

It is also worth noting here that it is not enough to show that theism is merely better than its competitors at accommodating moral realism. After all, it could still be implausible that theism accommodates moral realism, just a bit less implausible than that its competitors accommodate it. And if all options suffer that implausibility, then we have reason to be believe that no view can accommodate moral realism and so reason to deny moral realism. Being the best candidate for a job does not entail that you are qualified for it – every candidate might be awful and so the job might not get filled. So moral arguers need to show that theism not only accommodates better than its competitors, it also accommodates plausibly.

Now, one might grant the need to compare theism with its competitors but still hold that there will be asymmetry in how they are assessed. And that is because, the argument might go, theism *conceptually entails* moral realism. So we know that theism entails *that* moral realism is true even if we do not know *how* theism accommodates it. But if theism entails moral realism, then there must be *some* plausible accommodation account compatible with theism, even if we do not know what that account is. And so, while ordinarily a lack of extant plausible

accommodation accounts compatible with a worldview would be strong evidence against that worldview's ability to accommodate moral realism, that would not be so for theism. It does not matter how poorly the extant theism-friendly accommodate accounts fare. Theism is guaranteed some plausible accommodation account or other. But naturalism, on the other hand, has no such guarantee. For naturalism does not conceptually entail moral realism. And so naturalism is in the ordinary circumstance in which a lack of extant plausible accommodation accounts compatible with it would be strong evidence against naturalism's ability to accommodate moral realism. All the moral arguer needs to do, then, is show that naturalism has no extant plausible accommodation accounts available to it. No assessment of theism's accounts need enter the picture.

There are a couple problems with this argument. Firstly, I argue in §5 that theism does not have any relevantly interesting conceptual relationship to moral realism. Secondly, even if theism has all the relevant conceptual entailments, they do not help, for theism could have *impossible* entailments. The free will debate will again be illustrative. Traditionally, the concept of God in western monotheism is that of a being with free will. So theism conceptually entails that there is at least one being with free will, God himself. But it could turn out to be impossible that any being have free will; double incompatibilism could turn out to be true. Theism entailing that there is a being with free will does not preclude this. If double incompatibilism turns out to be true, then theism will turn out to have an impossible entailment. And so theism itself will turn out to be impossible. Indeed, if we have no extant plausible accounts of *how* any agent could be free, on either determinism or indeterminism, we start to have good reason to *believe* double incompatibilism, and so good reason to believe that theism is false.

Likewise, then, it could turn out to be impossible that there is any plausible way for theism to

accommodate moral realism. Theism entailing moral realism does not preclude this. If it does turn out to be impossible for theism to plausibly accommodate moral properties, then theism will turn out to have an impossible entailment and will itself be impossible. And again, if we have no extant plausible accounts of how theism could accommodate moral realism, we start to have good reason to *believe* that theism cannot accommodate it, and so good reason to believe that theism is false. It is important, then, to show that theism does have a plausible accommodation account. And so both theism and naturalism need to have their accounts assessed: there is no asymmetry. Craig's argument fits the schema. P2 really is the premise in need of assessment.

Now as I said in the introduction, the accommodation accounts most moral arguers appeal to are divine command theories. Particularly, a form of them we can describe as restricted theological voluntarism. Moral arguers attempt to show that this form of voluntarism is plausible and superior to its naturalistic rivals. But I have not yet said much about what this voluntarism looks like or why so many moral arguers have been drawn to it. Thus in the final part of this chapter, I elaborate on restricted theological voluntarism. Note that from here on I will shorten the name of the view to just *restricted voluntarism.* And when context does not require further specificity, I will refer to the view simply as *voluntarism* or *divine command theory.*

§2.5 RESTRICTED THEOLOGICAL VOLUNTARISM

Proponents of restricted voluntarism claim that what is morally right and wrong depends on God's will or commands. For simplicity, I will focus on commands. It being morally right to perform ϕ depends on God prescribing ϕ in a command. It being morally wrong to perform ϕ

depends on God proscribing ϕ in a command. For brevity we can say that moral rightness depends on divine prescription and moral wrongness depends on divine proscription.

Restricted voluntarists do not claim that what is morally *good* or *bad* depends on God's commands. Rather, goodness is said to depend on God himself, his nature, or character. The rough idea is that God himself serves as the paradigm of moral goodness. As voluntarists often put it, God is "The Good itself." Creatures are morally good insofar as they resemble God. An analogy often invoked is that of the Paris meter stick. We are to imagine that there is a particular, concrete stick in Paris that determines the length of a meter.[10] Whether some object is a meter in length depends on how it resembles the length of the Paris stick (Alston, 2002; Craig, 2009). Likewise, whether a creature is morally good depends on its resemblance to God. (Moral badness would depend on failure to resemble God).

It should be noted that some restricted voluntarists hold that God is the paradigm not just of *moral* goodness but also of other sorts of goodness. Adams claims that God is the paradigm of *excellence,* of which moral goodness is a part, but only a part. Thus, specifically moral goodness does not depend on resembling God *simpliciter* but on a certain kind of resemblance to God. Different theorists will spell out the relevant resemblance in different ways. I will take no stand on how it should be spelt out. Since my objections to the use of restricted voluntarism in a theistic moral argument do not depend on these details, I shall simply gloss over them. Thus, I will continue talking about moral goodness as resembling God *simpliciter,* while recognising that voluntarists make more nuanced claims than that. (And likewise for moral badness as failing to resemble God.)

[10] We can leave aside the question as to whether there is any such stick that performs this role in the actual world.

Giving separate accounts for moral obligations and moral values is the distinctive mark of restricted voluntarism. Splitting the accounts in this way has offered voluntarists responses to objections often thought to decisively defeat any divine command theory. Consider Euthyphro objections to divine command theories. Some versions object that on divine command theory, God cannot have any reason to command us to act a certain way. The thought here is that the only reason God could have for commanding a particular way is a moral reason. But, supposedly, according to divine command theory, there are no moral reasons prior to God's commands. Restricted voluntarists deny that there are no moral reasons prior to God's commands. They grant that there are no moral *obligations* prior to God's commands but God can still command on the basis of what is morally good to command. For on their view what is morally good depends on God's nature, not God's commands.[11]

Other Euthyphro objections allege that on divine command theory, God could make it obligatory to rape or torture. The thought here is that divine command theory denies that there are any moral norms external to God's commands to which the commands must conform. With no moral restrictions on commanding, God could command horrendous actions and, worse, those actions would become obligatory. But again, restricted voluntarists deny that there are no moral norms external to God's commands. They grant that God's commands are not restricted by moral *obligations,* but they are restricted by God's morally good character. Due to God's goodness, there are certain actions that it is impossible for him to command. This restriction on what God can command is why the view is called restricted voluntarism.

[11] Of course, it is possible to challenge the idea that God's reasons for commanding must be moral reasons but the point here is that restricted voluntarists can grant that they must be moral reasons and still escape the objection.

Splitting moral obligations and moral goodness in this way also helps voluntarists tackle another objection. Theists want to say that God is morally good. But if both goodness and obligatoriness depend on God's commands, then God's goodness amounts to his obedience to his own commands. This seems an unattractive view of God's goodness. For one, it is difficult to see what is so great about a being who keeps his own commands. Restricted voluntarism avoids that awkward account of God's goodness by denying that goodness depends on God's commands. There are, then, several different problems that restricted voluntarism seems to resolve. This greater resilience to traditional objections is no doubt a large reason for the popularity of restricted voluntarism.

In describing restricted voluntarism, I have been using the language of dependence. Moral obligations *depend* on God's commands, and moral goodness *depends* on God's nature. But there are lots of different ways that *x* can depend on *y.* There are two main kinds of dependence voluntarists appeal to. Adams appeals to reduction. He claims that moral obligations reduce to divine commands. Thus the property *moral rightness* reduces to *prescribed by God,* and the property *moral wrongness* reduces to *proscribed by God.* He also claims that moral values reduce to resembling God. Thus the property *moral goodness* reduces to *resembling God,* and the property *moral badness* reduces to *failing to resemble God.*

There is a puzzle about how to interpret God's own moral goodness on Adam's view. If moral goodness reduces to resembling God, then God's moral goodness will need to be God's resemblance of himself. The problem here is not so much whether one can resemble oneself. Suppose that resembling God is a matter of having some of the same properties as God. If that's so, then God trivially resembles himself. The issue is that this view puts the focus of God's goodness on his perfect resemblance of himself rather than his substantive character traits like

faithfulness, kindness etc. God is supposed to be the paradigm of goodness because of these character traits, not because of how he resembles himself. But Adams can get around this by saying that while the moral goodness of creatures is their resemblance to God, the moral goodness of God is not God's own self resemblance.

There are at least two different ways to flesh this out. Both agree that God's goodness reduces to his substantive character traits like faithfulness, kindness etc. Where they differ is in what to say about the relationship between God's goodness and creaturely goodness. The first strategy is to deny that there is any such thing as moral goodness *simpliciter.* There is only divine goodness and creaturely goodness. God's goodness reduces to his character traits, and creaturely goodness reduces to resembling God. The second strategy is to affirm that there *is* such a thing as moral goodness *simpliciter.* Moral goodness is God's goodness. That is, moral goodness reduces to God's own faithfulness, kindness etc. One then adopts a Platonic metaphysic in which creatures get to be morally good by *participating,* in some sense, in God's goodness. The claim could then be that participating in God's goodness (moral goodness simpliciter) reduces to resembling God.

Since Adams' theory is explicitly influenced by Platonic ideas, the latter reading is plausibly what he actually has in mind. But nothing I say hinges on which particular strategy is in view. When I talk about creaturely goodness, I will leave it open as whether that should be understood as its own qualified type of moral goodness or whether it should be understood as the participation in God's moral goodness *simpliciter.* Indeed, since in most of this thesis it will be creaturely moral goodness in view, I will normally omit the "creaturely" qualifier. I will only make explicit that I am dealing particularly with creaturely goodness when the context requires it.

To recap, we are looking at the two main ways to understand the dependence of moral obligations on God's commands, and moral values on God's nature. Adams claims that the dependence relationship is a reductive relationship. This is easily the more popular way of understanding the dependence. The majority of theistic moral arguers who appeal to restricted voluntarism deploy this kind of reductive voluntarism. But although it has not attracted much of a following, Quinn appealed to a causal relationship rather than a reductive one. At least, he did for moral obligations. Quinn claimed that God's commands cause there to be moral obligations. Put in terms of properties, God causes moral rightness and moral wrongness to be instantiated on certain actions. (It is not so clear how Quinn understood the relationship between God's nature and moral value.) This causal voluntarism should not be neglected in assessing moral arguments.

Clearly, both reductive voluntarism and causal voluntarism provide answers to some of the questions posed in moral metaphysics. They provide accounts as to how moral properties are to be accommodated. According to Adams, moral rightness is reducible, and reducible to properties about God's commands. According to Quinn, the moral rightness is explainable, and is explained by the causal power of God's commands. Moral arguers claim that at least one such voluntarist account is better than all its naturalistic competitors. They claim that one of these accounts can help establish P2, the key premise in theistic moral arguments. The task of this thesis is to assess whether any moral arguer has succeeded in using one of these accounts to establish P2. As already stated, I argue that no moral arguer has so far succeeded.

In the next chapter, I elaborate on what reduction and explanation involve, thus providing the metaphysical backdrop for the assessment of P2 in the remainder of the thesis.

3

Accommodating Moral Properties

In the previous chapter, I argued that the key premise of all theistic moral arguments that stands in need of assessment is P2 [it is plausible that theism accommodates objective moral properties and more plausible that it does so than that naturalism does so.] This is the claim that theism has plausible and better answers to questions about whether moral properties are reducible, what they might reduce to, whether their instantiation is explainable, and what might do the explaining, than naturalism. But I have not yet said much about what reduction or explanation involves. In this chapter, I elaborate on these concepts, arguing that there are three different forms of reduction: identity, constitution, and realisation. I also argue that reducing some phenomenon is not the same as explaining it, at least not in a robust sense of explanation. All of this metaphysical backdrop will be drawn upon throughout the thesis in assessing P2. In addition, I show that there are a limited number of ways that moral arguers can construct an accommodation account of moral properties.

The questions that an accommodation account attempts to answer are split into two kinds. There are questions about reduction and questions about explanation. I will say that answers to the questions about reduction form an *ontological* account of moral properties. And I will say that answers to the questions about explanation form an *explananda* account of moral properties.[12] A full accommodation account is therefore some combination of an ontological account and an explananda account. What P2 claims, then, is that theism has a plausible and better combination of an ontological and explananda account than any combination that naturalism has available to it. A given moral argument may claim that theism's advantage lay

[12] Note that these are stipulative definitions. I am not claiming that answers to different sets of questions might not also happily be labelled as ontological or explananda accounts.

in its ontological account or in its explananda account, or both, but most focus on just one.

§3.1 REDUCTION

Let us start, then, by explicating what is involved in offering an ontological account of moral properties. To offer such an account is to answer the question: **do moral properties *reduce* to some other sort of property, or not (and if so, which sort)?**[13] One might, for instance, try and reduce moral properties to psychological properties (perhaps about the satisfaction of human desires). Or one might try and reduce moral properties to theistic properties (perhaps about God's commands). Alternatively, one might deny that moral properties can be reduced at all and claim that moral properties are irreducible.[14]

Sometimes in popular-level discussion, reduction is taken to be an enterprise that only naturalists countenance. This is probably because reduction is popularly conceived as what happens when one does injustice to a phenomenon through using too austere an ontology to account for it. Thus, describing a view as "reductionist" is often an insult to the view, an accusation that it distorts the phenomenon or leaves out some crucial aspect of it. And many believe that the ontology naturalism is limited to is very austere (the stuff of micro-physics) and inadequate to account for rich phenomena like morality, consciousness, love, beauty etc. Thus, naturalism is thought to commit one to a reductionist (pejoratively understood) enterprise, but theism – with its richer ontology – is not. And the thought would go that theists ought to avoid reduction since reduction is bad, and theism's ontological resources provide a

[13] As I explain below, for some types of reduction, this description of reduction as involving some properties reducing to *other*, distinct properties is strictly inaccurate. But for the time being I am just painting an intuitive picture of reduction.

[14] Perhaps there are a variety of questions one might have about moral properties where accounts that provide answers could happily be labelled accounts of the ontology of moral properties. I am using the label just to designate accounts that answer the question about whether moral properties are reducible (and if so, to what).

way around it.

As a corrective, it may or may not turn out that naturalism cannot account for things like morality, consciousness, and so on, by appeal to micro-physics or anything else. But a naturalistic view is not to be described as a reductive view only on the condition of such failure. A naturalistic view could be reductive even if it accounts for the phenomenon in a full, satisfying, non-distorting way. There are good, successful reductions, and bad, unsuccessful reductions. Most people would be happy to say that water reduces to H_2O and would not say that there is anything distorted or missed out in this reduction. Of course, attempts to reduce some *a*-properties to some *b*-properties *can* distort the *a*-properties, but in principle they need not. There is nothing about reduction that is inherently distorting and so reductive claims should be assessed for adequacy on a case by case basis.

Therefore, the description of a naturalistic view as reductive should be seen as a neutral description, one applied simply because the view claims that *a*-properties reduce to *b*-properties. And since there is nothing inherently bad about reduction, theists can also have good reason to attempt it. Many *do* attempt it. But neither theists nor naturalists *need* offer reductive accounts. Both parties can pursue accounts according to which the *a*-properties are irreducible features of reality. Overall, then, from the perspective of popular-level discussion, it can be surprising that major theistic proponents of moral arguments (Adams, 1999; Craig, 2009; Evans, 2013) are reductionists about moral properties, whereas the most prominent defender of naturalism against moral arguments (Wielenberg, 2014) holds a non-reductionist position.

What, though, is reduction? Reduction seems to involve claiming that *a*-properties are not

robustly distinct, ontologically, from *b*-properties. But we can get clearer on what reduction involves. I will argue that it suffices for moral properties to have a reductive relationship to some *b*-type properties if one of the following conditions obtains:

(i) every token moral property is *identical* to a token *b*-type property.

(ii) every token moral property is *realised* by a token *b*-type property.

(iii) every token moral property is *constituted* by some token *b*-type properties.

Identity is the tightest of the relationships that could hold between a token moral property and a token *b*-type property. Indeed, it would be misleading to speak of a relationship between *two* token properties if there is identity. For there would just be the *one* property picked out in two different ways, by moral terms in one case and by *b*-terms (e.g. psychological terms, or theistic terms etc.) in the other case.[15] This would be similar to how "Clark Kent" and "Superman" pick out the same individual through different terms. It is clearly apt to describe an identity relationship as reductive since it claims in the strongest, most literal sense that there is no robust distinction between the *a*-properties and the *b*-properties: the *a*-properties *are* the *b*-properties.

Sometimes, however, philosophical views are considered non-reductive *even when* they identify every token *a*-type property with a token *b*-type property. Consider Anomalous Monism, a view in the philosophy of mind that is sometimes labelled non-reductive even though it identifies every token mental property with a token physical property.[16] It is labelled

[15] Note that in §2.3 I permitted irreducible, brute moral properties as counting as natural. Obviously, irreducible and brute moral properties are self-identical. If identity is a reductive relationship, it might look, then, that even views that posit irreducible, brute moral properties could count as reductive. As a corrective, we should only count an identity relationship as reductive if the relevant property is typically referred by more than one type of concept. If moral properties are irreducible and brute, then plausibly they are only typically referred to by moral concepts, and thus should not count as reductive.

[16] Donald Davidson (1970), the originator of Anomalous Monism, targeted mental *events* rather than properties, and targeted mental events with intentional descriptions in particular. But these details do not matter for our purposes.

non-reductive because the physical properties that the mental properties are identified with are not all of the same type of physical property. It claims that an instance of *believing that p* could be identical to an instance of *being in brain state* s_1 in one case, and in another, be identical to an instance of *being in brain state* s_2. It might be argued, then, that condition (i) does not suffice for reduction.

I grant that there is *a* sense in which Anomalous Monism is non-reductive. Physical properties, as a broad type, can be divided into more specific subtypes, like *being in brain state* s_1, *being in brain state* s_2, and a whole host of other brain-related (and not brain-related) properties. Likewise, mental properties, as a broad type, can be divided into more specific subtypes, like *believing that p, believing that q, desiring that p* etc. And Anomalous Monism denies that the identity reduction base of all tokens of a mental subtype are of *just one* physical sub-type. One could say that Anomalous Monism denies that mental *sub-types* can be reduced to physical *sub-types*. But it still affirms that every property in the general type of the *mental* is identical to a property in the general type of the *physical.* And intuitively, this is still a reductive claim, even if it is a more modest reductive claim than one that attempts to identify mental and physical sub-types.[17]

Sometimes it is thought that if *a*-properties reduce to *b*-properties, then the *b*-properties are in some sense prior to or more basic than the *a*-properties. If that is so, then an identity relationship would seem to be non-reductive since there are not two distinct properties that can have any relations of priority or basicness to each other. And I grant that sometimes the word "reduction"

[17] This means that moral particularists could also be reductionists about moral properties. Particularists claim that any feature of an action that is morally relevant in one case may fail to be morally relevant in another. They could be reductionists by claiming that one feature of an action is morally relevant and *identical* to the moral property of the action in one case, and in another case, some other feature is morally relevant and identical to the moral property.

implies a one-way relationship. It is natural to say that water reduces to H_2O but less natural to say that H_2O reduces to water. And this one-way implication of reduction might explain why it is odd to say that Superman reduces to Clark Kent. Perhaps, then, identity is not strictly speaking a reductive relationship. Nonetheless, it is a relationship where the *a*-properties are *nothing over and above* the *b*-properties. This is what identity shares with realisation and constitution. All are relationships where the *a*-properties are nothing over and above the *b*-properties. It is just that identity is the only one where the relationship is two-way, where the *b*-properties are also nothing over and above the *a*-properties. Nonetheless, I will continue to describe all of these relationships as reductive, for the simple fact that reductive language is less cumbersome than nothing-over-and-above language.

The nothing-over-and-above claims behind my reductive language may suggest to some that reduction involves claiming that the phenomenon does not exist – that it is literally *nothing.* But I deny that reduction entails eliminativism. What is distinctive about an eliminativist position is the claim that the concept of the *a*-properties fails to refer to anything. Perhaps *a*-properties were posited to explain some phenomenon, but then we discovered that what actually explains the phenomenon is so different that the concept of the *a*-properties just does not apply. The classic example is the positing of phlogiston to explain combustion. What actually explains combustion is so different to what phlogiston was imagined to be, that the concept of phlogiston has no application. I am assuming, when talking about identity, realisation, and constitution, that the concept of the *a*-properties *does* have application.

Moving on from identity relationships, realisation occurs when some property is instantiated by virtue of a functional role being occupied. For instance, perhaps the concept of pain can be analysed as: the state that performs the role of registering bodily tissue damage, causing

wincing, motivating one to move away from the cause of the damage etc. And perhaps the brain state of c-fibres firing performs this functional role; in which case, c-fibres firing *realises* pain. More specifically, a case of role-occupation is one of realisation rather than one of identity when the role-occupier is not *rigidly designated* by the concept that has a functional analysis. For instance, the concept of water can perhaps be given a functional analysis: the stuff that comes out of our taps, flows in our rivers, appears translucent blue etc. H_2O occupies this role, but it is identical to water rather than merely realising water because, intuitively, H_2O is still water even in worlds where it is not performing the water role. At least, this is the intuition that Twin Earth thought experiments seem to prompt (Putnam, 1973).

So c-fibres firing merely realises pain rather than being identical to it if the state is not pain in worlds where it is not performing the pain role. If c-fibres firing merely realises pain, then plausibly there are some worlds where there is pain but it is realised by some other state, like a silicon state or an ectoplasmic state. Realisation seems reductive because, intuitively, in a world where c-fibres firing realises pain, there is no robust distinction between the fact that S is in pain and the fact that S's c-fibres are firing. A moral property, then, can be realised by some *b*-type property, and thus reduced to it, if a moral concept can be given a functional analysis, the *b*-type property suffices for the functional role to be occupied, and the *b*-type property is not rigidly designated by the moral concept.

Now to the final relationship; constitution is easiest to grasp first through consideration of ordinary physical objects. The chair I am sitting on is constituted by the assembly of its legs, seat, and back. The chair is not identical to or realised by any one of these parts, yet assembled they constitute it. When it comes to moral properties, then, one might hold that the property of moral wrongness is constituted just by theistic properties: God's commanding and God's

having moral authority. Alternatively, it might turn out that some property is constituted by properties of *more than one type*. For instance, perhaps wrongness is constituted both by properties about God's nature and the nature of human organisms. In this sort of case, the property is reducible, it is just not reducible to a single type of property. Such moral properties will be dualistic composites, akin to how on some views, human persons are composites (of mind and body).

Constitution seems to be a reductive relationship because, to borrow language from Kripke (1980) and Chalmers (1996), after God creates some legs, a seat and back, and then assembles them, there is nothing more God needs to do to get a chair.[18] And indeed, God could not create the legs, seat, and back, assemble them, and then prevent there being a chair. Intuitively, the fact that there is a chair does not seem to be a robustly distinct fact from the fact that there is a set of legs, seat, and back assembled. The chair is nothing over and above the assembled legs, seat, and back.

I have argued, then, that it suffices for a relationship between *a*-type properties and *b*-types properties to be reductive if the relationship is one of identity, realisation, or constitution. Of course, we must be careful to not always assume that the properties being considered for reduction are the whole, broad class of moral properties. As discussed in §2.1, there are different sub-types of moral properties. These different sub-types may have different ontological accounts. One might, for instance, claim that moral rightness and moral wrongness are reducible but that moral goodness or badness or not. There is no guarantee, then, that moral properties as one single general type of property have an ontological account shared by all its

[18] This is not an infallible guide to picking out a case of reduction. It could be that when God creates *x* there is nothing more he needs to do to create *y* but only because *x* causes *y* of necessity. And as I argue below, causation is not a reductive relationship. But where we are confident that *x* is not the cause of *y,* it being the case that God gets *y* for free when creating *x* is evidence (even if defeasible) that *y* reduces to *x.*

members. Accordingly, I will be careful in assessing moral arguments to specify, when necessary, whether it is only a particular subtype of moral property that is in view.

There are other sorts of metaphysical relationship that might obtain between one type of property and another, such as causation or supervenience. One might wonder whether these suffice for reduction. I deny that they do. Intuitively, causation – or at least what is sometimes referred to as *efficient* causation – is a relationship that occurs between two types of property that have a robust ontological independence from each other. To illustrate, consider that according to western monotheism the physical world has a robust ontological independence from God. Such theists do not want to say that the physical world reduces to God. Such a claim would be closer to pantheism or panentheism. Yet theists want to say that God has *some* metaphysical relationship to the world. The relationship that seems to preserve robust ontological independence between God and world is causation. Causation is the paradigmatic metaphysical relationship that occurs between robustly distinct things. So, if moral properties have a merely causal relationship with some properties, then moral properties do not reduce to those properties.

Supervenience is a relationship where *a*-type properties track *b*-type properties across situations. The scope of the situation, and whether the tracking occurs merely within a world or across worlds, makes for distinctions in the strength of supervenience claimed, distinctions sometimes captured by the language of local vs global, weak vs strong supervenience. But the same point against supervenience being a reductive relationship applies to all its varieties. Consider the supervenience claim that no worlds that are identical in all of their non-moral properties are divergent in any of their moral properties. The claim that moral properties track the non-moral properties in this way is just that – a claim about one type of property *tracking*

another. There is nothing inherently reductive about this claim. There is nothing incoherent in the claim that *a*-type properties track robustly distinct *b*-type properties.

Of course, many instances of supervenience cry out for explanation. If mental properties supervene on physical properties, we want some explanation as to why, and one explanation might be that mental properties reduce to physical properties. If mental properties are not robustly distinct from physical properties, then of course mental properties track physical properties. But a reductive relationship is not the only sort that can explain supervenience. Perhaps, for instance, mental properties track the physical properties because the physical properties necessarily cause mental properties to be instantiated. And I argued above that causal relationships are not reductive. So not only is supervenience not itself a reductive relationship, it also does not entail any reductive relationship. It is a metaphysically indeterminate relationship – a tracking relationship to be explained by a more determinate sort of relationship like identity, constitution, realisation, or causation.

Are there any other metaphysical relationships that are candidates for being reductive? One might suggest that I have neglected to consider a *grounding* relationship. There is considerable debate about what grounding is, or whether it is any *one* thing. If grounding turns out to be constitution, then by my account grounding is a reductive relationship. If it turns out not to be constitution, and to be some sort of non-causal explanatory relationship, then it will probably not be reductive. For below, when I discuss the relationship between accounts of the ontology of moral properties and accounts of their explanation, I will argue that reductive relationships are not robustly explanatory.

There is, however, little need to take a stance on grounding. When the language of grounding

is used by moral arguers, the context reveals that one of two things are meant. Firstly, some sort of reductive relationship like identity is meant. Craig (2010) claims that moral properties are "grounded" in God's commands, but as he makes clear elsewhere (2009), he follows Adams in claiming that the property *being commanded by God* is identical to the relevant moral property. Secondly, grounding is used to describe the making relationship that is of interest to first-order normative ethics. When some non-moral property of an action makes that action wrong, the non-moral property is sometimes said to "ground" the wrongness of the action. This seems to be what Baggett & Walls mean when they claim that human value is grounded in humans being made in God's image (2016, pp.115-144). I will have more to say about the making relationship in §6. Meanwhile, avoiding the language of grounding will help to keep distinct issues distinct.

There might be other metaphysical relationships that exist between some *a*-type properties and some *b*-type properties but since the literature on moral arguments does not appeal to others, I will not consider any others as candidate reductive relationships. In sum, identity, constitution, and realisation suffice for reduction, and causation and supervenience do not suffice (causation in fact *excludes* reduction). An ontological account of moral properties that claims that moral properties are reducible is one that will appeal to identity, realisation, or constitution. Now it is time to talk about explananda accounts.

3.2 EXPLANATION

Giving an explananda account of moral properties is answering the question: **does the instantiation of moral properties have an explanation? (If so, what is it?)** And the account either offers an explanation of their instantiation or declares that there is no explanation to be

had; it declares that the properties are *brutely* instantiated. Of course, accounts that declare bruteness are still addressing moral properties as properties that are *prima facie* candidates for having an explanation, these accounts just claim that there turns out not to be any explanation (perhaps even that there could not possibly be one). Hence why it is still felicitous to call them explananda accounts.

The main thing I will argue in this section is that giving an ontological account is not the same as giving an explananda account. Giving answers to questions about moral properties and reduction does not necessarily settle questions about moral properties and explanation. That they are distinct accounts is *prima facie* plausible. After all, one is asking whether moral properties are reducible or not. The other is asking whether moral properties have an explanation or not, and if so, what that explanation is. Roughly, ontological accounts tell us *what* moral properties are (they are irreducible, or natural, or theistic) and explananda accounts tell us *why* moral properties are instantiated (or tell us that there is no why). These seem like different questions being answered.

But one might object that even though these accounts might answer different questions, answers can be provided by appeal to the same set of properties. That is, *b*-properties can be both the reduction base and the explanation of the *a*-properties. Thus, picking out the reduction base of the *a*-properties gives one all the ingredients needed to offer an explananda account of them. And so, while there is an in-principle distinction between an ontological account and an explananda account, in practice, the former account brings along with it the latter.

There are a few problems with this, however. Firstly, not all ontological accounts are reductive. One might hold that the *a*-properties are irreducible. In which case, the ontological account

does not invoke any further *b*-properties that can then be invoked once more in an explananda account. This problem might be surmountable if irreducible properties were necessarily brute. That way, settling the ontological account as non-reductive would settle the explananda account one will offer. But it is just not true that *sui generis* properties are always brute. On theism, physical properties are irreducible. They do not reduce to properties about God, or anything else.[19] But they nonetheless have an explanation. God causes the physical properties to be instantiated in his act of creation. So, offering an ontological account need not get you close to an explananda account, if only because some ontological accounts are non-reductive.

Secondly, not even a reductive account that *does* invoke some further *b*-properties will be able to appeal to those same *b*-properties in an explananda account. Very plausibly, some further *c*-properties will need to be appealed to, or one will opt for bruteness. For if properties have explanations, they are always (with a caveat addressed below) explained by some robustly distinct properties (or facts, events etc – but I will focus on properties for ease of exposition). That the window is smashed is explained by its being hit by a cricket ball, a robustly distinct property. But reduction is precisely an attempt to eradicate a robust distinction between *r*-properties and *b*-properties.

At the most extreme end, a reductive account that posits an identity relationship eradicates any distinction at all between *a*-properties and *b*-properties. Since on such an account, the *a*-properties *are* the *b*-properties, the claim that *b*-properties explain the *a*-properties amounts to the claim that these sorts of facts *explain themselves.* But again, explanations always appeal to some robustly distinct properties. Nobody would accept that the fact that the window is smashed explains itself. So, it is highly likely that at least an identity reduction account will

[19] At least, physical facts do not reduce to anything else if we assume for the moment that panpsychism is false.

need to appeal to some further *c*-properties if an explanation of the *a*-properties is to be offered.

A less extreme reductive account – one that posits a constitutive relationship between *a*-properties and *b*-properties – seems more likely to allow that there also be an explanatory relationship between the *a*-properties and *b*-properties. For there is a sense in which, say, a chair's existence is explained by the existence of its parts and their assembly. But even so, this is not the sort of explanation we are normally most interested in discovering. Somebody who asks for an explanation of the chair's existence is not normally satisfied by the answer that it exists because its parts exist and are assembled together. One wants to know *why* the parts exist and are assembled together. And once again, if there is an explanation of this, it is always to be found in some further, distinct property (or substance, fact, event etc). In the case of the chair, the further fact might be that its parts were put together by machines in a factory. Constitution offers a sort of very modest explanation, but a more robust explanation requires appeal to robustly distinct facts.

Of course, in the discussion above I noted a caveat when claiming that explanation appeals to robustly distinct properties in every case. For one might hold out the possibility that some very special entities or properties are indeed *self-explanatory.* Some have thought that about God or the property *being God*. God is considered by some to be neither brute nor explained by anything outside God; God existence is explained by God own's nature. This is a highly controversial idea, and many theists will be content to say that God's existence is brute rather than self-explanatory. But some theists hold it.[20] And so, if the good is identical with God, as Adams (1999) claims, and God's existence is self-explaining, then a reduction base of a moral

[20] O'Connor (2008) presents a recent defence of this sort of idea.

property also explains that moral property.[21] I will not rule this out but the point remains that *prima facie,* there is good reason to believe that an ontological account will not provide an explananda account.

Of course, I do not deny that having an ontological account can assist in building an explananda account. Suppose you know that God's commands are explained by God's willing them to be. If you then learn that moral obligations are identical to God's commands, you are in a position to learn that moral obligations are explained by God's willing them to be. For if God's commands are identical to moral obligations, then the explanation of the former is the explanation of the latter. The point, however, is that even in this sort of case, the property that acts as the reduction base (being commanded by God) is not the same as the property that does the explaining (being willed by God). And this is just one sort of case, others will provide less assistance for building an explananda account; consider that one could know that water is H_2O without having any clue as to what explains the existence of H_2O.

I have said more about the relationship between reduction and explanation than about what explanation itself is. I do not intend to give a general account of what explanation is. Sometimes, as Murphy says, a common understanding of explanatory language is at least partly epistemic, "to explain is, on this understanding, to give new information (2011, p.45)." It should be clear from the above that I am not using explanatory language in an epistemic sense. Rather, it describes a metaphysical relationship between two properties (or facts, events etc). An explanation answers why a property is instantiated, where this answer is one "that is not observer-relative – one that is not dependent on one's mastery or failure of mastery of concept-

[21] Although strictly speaking, this is a moral substance, not a moral property. God is not a property. So anything that is identical to God cannot be a property either.

use and that is not dependent on one's knowledge or ignorance of the relevant facts (2011, p.45)." A paradigmatic explanation of this sort is a causal explanation. But I will not provide any further details than that.

Since reduction and explanation are distinct relationships, a full account of how moral properties are to be accommodated requires giving an account of both what explains them (if anything) and what they reduce to (if anything). Such a full account will adopt one of the following four structures:

(1) Moral properties are reducible and have an explanation.

(2) Moral properties are reducible and are brute.

(3) Moral properties are irreducible and have an explanation.

(4) Moral properties are irreducible and are brute.

Now recall what P2 of the moral argument schema claims: [it is plausible that theism accommodates objective moral properties and more plausible that it does so than that naturalism does so.] That turns out to be the claim that theism has an accommodation account matching one of those structures that is more plausible than any account naturalism has using any of those structures. But in fact, these structures can be further sub-divided based on the type of properties appealed to in reduction or explanation. In the next section, I explain these further sub-types and argue that only a select few of them are useful to theists who want to run a moral argument for theism. Most of those turn out to correspond to the two main versions of restricted voluntarism.

§3.3 SURVEYING THE OPTIONS

Recall that I am assuming, along with moral arguers, that theism and naturalism are our two options for big metaphysical outlooks. This limits the properties that can be appealed for reduction or explanation to theistic properties and natural properties. What we need to do, then, is divide the four structures given above for accommodation accounts even further, to reflect that one could appeal to theistic properties for one purpose and natural properties for another. The different combinations are given in this table:

	moral properties	
	reduce to these properties …	are explained by these properties…
1	natural	natural
2	natural	theistic
3	natural	none (brute)
4	theistic	natural
5	theistic	theistic
6	theistic	none (brute)
7	none (irreducible)	natural
8	none (irreducible)	theistic
9	none (irreducible)	none (brute)

One might think that theists are restricted *qua* theists to only using those accounts that appeal to theistic properties. But that is not obviously the case. Any natural properties that could serve as reduction bases or explanations are seemingly compatible with theism. Theism does not deny natural properties, it just adds theistic properties. Nor is there any explicit contradiction between theism and irreducible or brute moral properties. So there is a lot of logical space for developing an accommodation account compatible with theism. But some of these options turn out to be useless in motivating a moral *argument* for theism.

We should distinguish two aims that a theist might have in developing an accommodation account. One aim is to develop the account that best satisfies the theist herself in light of everything she believes. Another aim is to develop the account that best provides the *naturalist*

a reason to switch over to theism. It is those accounts that have promise for the latter aim that are of interest when assessing moral arguments. But those accounts are a smaller subset of all the accounts that the theist might be interested in if she puts aside arguing with naturalists. Below, I rule out those accounts that are implausible on theism or have little chance of motivating a moral argument. Of course, we should also be interested in which options are available to naturalism. Obviously, any options that appeal to theistic properties will be unacceptable. But there are others that will be implausible on naturalism. I will also weed those out.

Let us start with the most obviously implausible combinations. Theism cannot really countenance the option that moral properties are reducible to natural properties *and brute.* For on theism, no instantiated natural properties that could serve as a reduction base are brute. On theism, all physical properties, for instance, are either caused by God or by something else that God caused.[22] Neither can theism countenance the option that moral properties are theistically reducible but have a natural explanation. This would imply that features of the natural world cause God to, say, command a certain thing. But even if God's knowledge of the natural world constrains what God is able to command (maybe God only commands what is physically possible for us to do), the natural world does not cause God to command.

There is also an option for naturalism that can be dismissed: that moral properties are reducible to natural properties yet are brute. Naturalism might have to countenance *some* brute natural properties, like the properties of the initial state of matter or energy, but the kind of natural

[22] Note that causation does not need to be deterministic. So examples of indeterminacy in the physical world do not necessarily counter-example the claim that physical properties are caused by God or other things caused by God. Now some might hold that the free actions of embodied creatures are literally without cause and propose this as a counter-example. But for one, this strikes me as an implausible view of free will since it has difficulty explaining how the agent has any control over the action. And secondly, actions are not the sort of things one will be tempted to reduce moral properties to. So this view of free will can be ignored for our purposes.

properties that are candidate reduction bases for moral properties are not plausibly brute, on naturalism. For these properties are all high-order or late-arriving properties as far as the structure and development of the physical universe is concerned. Consider that if moral properties are reduced to properties of humans and their capacities, it is implausible that these facts are brute, since human beings are not brute. Human beings are caused and explained by an evolutionary process.

Next we can weed out those options that even if plausible, cannot motivate a moral argument for God.

Consider the option that theism is true, moral properties are naturalistically reducible, and have a naturalistic explanation. This option does not appeal to any distinctly theistic properties, so it is hard to see how it allows for the claim that theism has some advantage over naturalism in accommodating moral properties. After all, naturalism would surely be able to appeal to those same natural properties. What moral arguers need to do is appeal to the unique ontological resources of theism. Otherwise they cannot claim any advantage for theism and the moral argument fails.

One might protest that the plausibility of moral properties being reducible to natural properties is greater on theism than naturalism. How could this be? Well, perhaps certain moral concepts have a functional analysis. So suppose that moral wrongness is whatever property has a certain functional profile. It might turn out that the property that satisfies that profile is a natural property but that it only has that functional profile due to its interactions with some theistic property. Here is a comparison to clarify: suppose that the concept of water has a functional analysis. Water is whatever the predominant substance is that falls from the skies, comes out

of our taps, appears translucent blue etc. Now H_2O satisfies that functional profile but it only does so because of its interactions with other substances and properties. It would not fall from the sky if it were not for the earth's gravitational pull, for instance. Now suppose, somewhat implausibly, that gravitational theory is hopelessly riddled with problems such that the best explanation of falling bodies is the direct intervention of God. It would then turn out that H_2O is water but that the best account of how it could be H_2O entails theism. Theism works as a background enabling condition of H_2O being water. Although this argument with respect to water is not very good, perhaps there is a promising equivalent with respect to moral properties and theism?

I do not think so. Consider that it is either controversial that the background condition requires theism, or it is not. If it is controversial, then probably the moral argument collapses into some other argument. Suppose for instance that some natural property *n* satisfies the functional role of moral obligatoriness only if the *purpose* of human organisms is to *n*. Some might argue that human organisms have a purpose only if God designed them. But this is controversial. Others argue that human organisms can have a purpose by virtue of their evolutionary history. Theists might be right that God is needed for purpose, but it takes some argument to show that. And if the argument is successful, any *moral* argument that is enabled by that would be secondary. Purpose is of independent and wider interest beyond any role it has in moral concepts. The heart of the debate would not really be about anything particular to moral theorising, rather it would be purpose and teleology.

On the other hand, if it is not controversial that the background condition requires theism, then it is not clear that we have a case of *natural* reductionism. Let us consider a case where it would be uncontroversial that the background condition requires theism (at least on the assumption

that theism and naturalism are the only options). Suppose the concept of moral wrongness can be analysed, in part, as that property of actions that one is guaranteed to be retributively punished for instantiating. If naturalism is true, there is no supernatural agent who can guarantee that every action with such-and-such property is retributively punished. So, take a natural property like *inflicts pain*. Even if *inflicts pain* satisfies every other part of the role laid out by the concept of moral wrongness, it will not fulfil the role entirely unless theism is true. And so it is more plausible on theism than on naturalism that *inflicts pain* is the reduction base of moral wrongness.

In a case like this, it is tempting to say that really moral wrongness is just doing something that is guaranteed to be retributively punished by God. But that is a straightforwardly theistic property. There seems to be no advantage in saying that really, it is the natural property that is the reduction base. So, this may as well be treated as a case of reducing moral properties to theistic properties. Therefore, there does not seem to be much promise in a moral argument for theism defended via natural reductionism, which is perhaps why no moral arguments use that strategy.

Still, one might claim that when it comes to moral properties being reduced to and explained by natural properties, theism has an advantage with respect to the explanation component. Perhaps it is more plausible that some natural property explains moral properties on theism than on naturalism. One might argue, say, that on naturalism, the natural property that does the explaining will either be itself brutely instantiated or will ultimately trace back to some brutely instantiated natural property. Perhaps on naturalism, the instantiation of all natural properties is ultimately brute, and this is implausible. But to go down this route is just to collapse the moral argument into some sort of cosmological argument against naturalism. And

in that case, the moral argument has lost any independent dialectical force. Therefore, we can safely rule out appealing to only natural properties in a moral argument for theism.

A similar point can be made against moral arguers relying on the option that moral properties are reducible to natural properties but theistically explained. If one conceded that moral properties were naturalistically reducible but denied that their immediate explanation is also natural, one would be arguing for some direct intervention on God's part in creating part of the natural world. This would result in offering some argument other than a moral argument. For instance, if moral properties reduce to properties about human psychology, and human psychology can only be explained by God's direct intervention, then one's moral argument has collapsed into something else. It might be, say, an argument from consciousness.

One might think that the option for theism of leaving moral properties without reduction or explanation would also preclude running a moral argument. For once more, no distinctly theistic properties are appealed to. But some theists seem to hold that even if moral properties are neither reduced to nor explained by theistic properties, such moral properties fit better, in some sense, on theism than on naturalism. Brute, irreducible moral properties (or facts) would be ultimate, "deep" parts of reality, and that is supposed to be strange on naturalism (Evans, 2013, p.115; Wainwright, 2005, p.67).

These arguments are interesting but underdeveloped. They are usually expressed by moral arguers who endorse some divine command theory but want to say that even if their theory fails, there is still promise for a theistic moral argument. I will say more about these arguments in §8, where I suggest that they should be interpreted as claiming that irreducible brute moral properties are more *probable* on theism than naturalism. Since I argue in this thesis that divine

command theories do not motivate a moral argument for theism, I suggest that these probabilistic arguments are where moral arguers could now direct their attention. Appealing to brute, irreducible moral properties remains a live option for moral arguers.

Having weeded out the implausible options or those useless for theistic moral arguments, here is a table summarising the remaining options:

	moral properties			
	on **theism**		on **naturalism**	
	reduce to	explained by	reduce to	explained by
1	theistic	Theistic	natural	natural
2	none (irreducible)	theistic	none (irreducible)	natural
3	none (irreducible)	none (brute)	none (irreducible)	none (brute)
4	theistic	none (brute)		

For a moral argument for God's existence to succeed, at least one of these theistic options must be shown plausible and better than all of these naturalistic options. That is what is required to show that the crucial P2 premise is true.

As already explained, this thesis assesses attempts to motivate moral arguments by appeal to restricted voluntarism. Between the reductive and causal versions of restricted voluntarism, three of these four options are covered. Adams' reductive voluntarism claims that moral wrongness reduces to the theistic property *proscribed by God* and is explained by the theistic property of *God's directly willing to proscribe.* That's option one in the table. Quinn's causal voluntarism claims that moral wrongness is irreducible but explained by the theistic property of *God's directly willing to proscribe.* That's option two in the table. Further, as explained in §2.5, Adams' view also claims that God's goodness reduces to his own character traits, like faithfulness, kindness etc. But these are necessary aspects of God's nature that are not explained by anything. That's option four in the table. The only option not represented is three.

And as I said above, that option is very underdeveloped by moral arguers. So voluntarism covers all of the plausible options that moral arguers have spent time developing.

Assessing P2 could involve assessing the plausibility of voluntarism or assessing whether it has advantages over all its naturalistic competitors. As previously indicated, I focus on the latter comparative element. Now here is an obvious point: a necessary condition of beating all of one's competitors is beating *some particular competitor*. Adams' reductive voluntarism needs to at least beat an equivalent naturalistic view that reduces moral properties to natural properties. Quinn's causal voluntarism needs to at least beat an equivalent naturalistic view where natural properties cause moral properties. I argue in the rest of this thesis that moral arguers have not shown voluntarism to meet this necessary condition. They have not shown that their theistic views have an advantage over their naturalistic equivalents.

In the next chapter, I begin to show that Adams' reductive voluntarism does not have an advantage over its similarly reductive naturalistic equivalents.

4

The Conceptual Gap Argument: *Interpretation*

Theistic moral arguers aim to motivate P2 [it is plausible that theism accommodates objective moral properties and more plausible that it does so than that naturalism does so]. One way they have tried to do that is by arguing that Robert Adams' reductive voluntarism is more plausible than its naturalistic competitors. Obviously, that requires reductive voluntarism having an advantage over some *specific* competitors. It must have an advantage over natural reductionism, which claims that moral properties reduce to natural properties.[23] Moral arguers have two main arguments to show reductive voluntarism is superior to natural reductionism. The first of these arguments, however, has not been expressed in explicit detail. The task of this chapter is to interpret what the argument is so that it can be assessed in the following chapter.

I interpret the argument as claiming that in order for moral properties to reduce to properties picked out by non-moral concepts, there would have to be conceptual connections of a certain sort between those concepts and moral concepts. Moral arguers charge that there are no such connections between natural concepts and moral concepts but there are those connections between theistic concepts and moral concepts. So moral properties cannot reduce to natural properties, but they can reduce to theistic properties. This interpretation is evidenced by widespread spread use of conceptual considerations to attack reductionist positions, and by the textual evidence in the writings of moral arguers themselves.

I already explained in §2.5 some details of Adams' reductive voluntarism. It claims that moral

[23] This is the view more standardly called 'ethical naturalism' in the literature.

rightness reduces to being prescribed by God. And it claims that moral goodness reduces to resembling God. But I have not yet said which reductive relationship – identity, constitution, or realization – Adams appeals to. The standard interpretation is that he opts for identity claims, such that moral rightness is one and the same property as being prescribed by God. There are alternative interpretations around. Wainwright (2005, pp.88-9) holds that Adams is best interpreted as invoking constitution rather than identity. But since my discussion is unaffected by which particular relationship is in view, I will just adopt the more standard identity interpretation.

In addition, not much has been said yet about natural reductionism other than that it reduces moral properties to natural properties. One prominent view in this camp is Cornell Realism, championed by Richard Boyd (1988) and others. Roughly, on this view, moral goodness is a complex natural property made up of simpler natural properties. All else being equal, when enough of these simpler properties, like pleasure, or honesty, are present, goodness is present. The precise simpler properties that are present in each case of goodness can vary. Another well-known natural reductionism is espoused by Philippa Foot (2001). Roughly, moral goodness in human beings reduces to their functioning properly as human beings. This view appeals to properties and facts about natural teleological or biological normativity. However, my discussion will not need to draw on the details of any particular natural reductionist view.

§4.1 INTRODUCING THE MAJOR OBJECTION

As I said above, I will be interpreting the first of two main arguments for preferring reductive voluntarism to natural reductionism. This first argument has a place of prominence in the attack on natural reductionism, despite its lack of explicit formulation. It ties together a theme that

runs through the writings of several moral arguers. That theme is the apparent gulf between the normativity of moral properties and the apparent lack of normativity in natural properties. Below is a sampling of quotes from moral arguers commenting on this gulf.

Here are Moreland & Craig expressing what they explicitly label the "major objection" to natural reductionism, which is that

> it confuses an is with an ought by reducing the latter to the former. Moral properties are normative properties. They carry with them a moral "ought." If some act has the property of rightness, then one ought to do that act. But natural properties like the ones listed do not carry normativeness. They just are (2003, p.401).

Their objection seems to be that moral properties are essentially normative but natural properties are not normative. And so, since they miss out an essential quality of moral properties, natural properties cannot be the reduction base of moral properties.

A similar objection is expressed by Wainwright:

> Mackie has pointed out that, in the ordinary view, moral facts have "requiredness" or "not-to-be-doneness" built into them, and the property (or set of properties) picked out by "goodness" has "to-be-pursuedness" built into it. Yet it is difficult to see how purely empirical properties could have these characteristics built into them (2005, p.51).

Again, the claim seems to be that moral properties are essentially normative (here expressed as their possessing "requiredness" or "not-to-be-doneness") but natural properties are not. Or at least, empirical natural properties are not. And so it is difficult to see how moral properties could reduce to natural properties.

Baggett and Walls claim that (or at least attribute to C.S. Lewis the view that)[24]

> The source of ... moral obligation isn't likely to be mere matter. An evolutionary account of *feelings of* or *beliefs in*, say, moral obligation is certainly possible, but how could naturalism explain obligation itself? How collections of atoms could generate and issue genuinely binding moral commands is altogether mysterious, if not absurd (2011, p.11).

Here again, the claim seems to be that certain paradigmatic types of natural property, this time concerning the arrangement of atoms, lacks the features to constitute something normative like a *binding* obligation.

Evans says,

> If one thinks, as Mackie clearly did, that the world should be understood in a naturalistic manner, then values and other moral properties may indeed appear odd or peculiar. To focus on the deontic moral properties that are my major concern, how could one begin to explain the existence of actual moral obligations if the universe consists solely of subatomic particles configured in various ways? (2013, p.120)

Yet again, the claim is that there is a difficulty in seeing how natural properties – this time, properties about subatomic particles – can account for normative properties, like moral properties.

What I am calling the major objection to natural reductionism is an objection that seems to focus on the normativity of moral properties and the apparent lack of that normativity in natural properties. But of course, these quotes underdetermine any precise argument. My task will be to find a precise formulation of the objection that has the best chance of successfully defeating natural reductionism while leaving reductive voluntarism intact. Here is one formulation of the

[24] It is ambiguous as to whether Baggett and Walls endorse the precise views that they attribute to Lewis.

major objection, one that is most easily discernible in the Moreland and Craig quote:

The Major Objection

(1) All moral properties are normative.

(2) No natural properties are normative.

(3) If all moral properties are normative but no natural properties are normative, then moral properties do not reduce to natural properties.

Therefore,

(4) Moral properties do not reduce to natural properties.

A moral arguer would then add that theistic properties *are* normative, and so the same objection cannot be launched against theistic reductionism.

But this formulation does not make for a strong argument. It is problematic because it is plausibly question-begging. After all, it does not seem like the second premise is one that a natural reductionist would accept. The natural reductionist holds that some natural properties, individually or collectively, *are* morally normative.[25] The position of the natural reductionist is precisely that some natural properties, in a certain sense, *are* moral properties, and so bear moral normativity. The natural reductive thesis is precisely that there is no robust ontological distinction between moral properties and natural properties. So this argument is unlikely to be dialectically effective.

[25] The same sort of accusation of circularity to arguments like this is made by Finlay (2014, p.13)

The task of finding a dialectically effective version of the argument will take us to the sort of interpretation of this argument that I hinted toward at the beginning of this chapter. The argument will turn out to concern conceptual constraints on reduction. But to get there let us consider what moral arguers might say in defence of the claim that we should not *expect* to get a dialectically effective argument.

I imagine some moral arguers responding with a concession that the second premise is too close to the conclusion. But, they would say, it is a mistake to think that there is much hope for expressing their primary grounds for rejecting natural reductionism in the form of an effective argument. Their grounds for rejecting natural reductionism is just a very strong intuition that natural properties are not morally normative. That natural properties are ontologically distinct from moral properties is something one can just, in a sense, "see". Here moral arguers might approvingly quote Derek Parfit:

> Many kinds of thing, event, or fact are, however, undeniably in different categories. Rivers could not be sonnets, experiences could not be stones, and justice could not be—as some Pythagoreans were said to have believed—the number 4. To give some less extreme examples, it could not be a physical or legal fact that $7 \times 8 = 56$, nor could it be a legal or arithmetical fact that galaxies rotate, nor could it be a physical or arithmetical fact that perjury is a crime. It is similarly true, I believe, that when we have decisive reasons to act in some way, or we should or ought to act in this way, this fact could not be the same as, or consist in, some natural fact, such as some psychological or causal fact (2013, pp.324-5)

What we have is a bare appeal to a strong intuition about the distinction of moral and natural properties. And moral arguers are in good company in having this intuition. It is not just Parfit who seems to share it. It is widespread. Laskowski (2019, p.94) calls this intuition *the sense of incredibility in ethics* and cites several philosophers who express it.

> Huemer … claims, ‘‘value properties are radically *different* from natural properties.’’ Dancy … says, ‘‘There remains a stubborn feeling that [normative] facts about what is right or wrong, what is good or bad, and what we have reason to do have something *distinctive* in common, and that this common feature is *something that a natural fact could not have*.’’ And in a statement that has become a slogan, Enoch … writes, ‘‘Normative facts are *just too different* from natural ones to be a subset thereof” (emphases added by Laskowski).

Moral arguers might grant that what we *can* argue about is whether we have any overriding reason to distrust this intuition. But they would say that the initial reason that the intuition itself provides to reject natural reductionism cannot be expressed in any substantial argument. All they can do is pose the question to any onlookers to the debate: “don’t you also have this strong intuition? If you do, join us in rejecting natural reductionism.”

One might reply that there is just as much intuitive resistance to theistic reductionism as there is to natural reductionism. But as best as I can tell, that is just not true. Granted, *some* philosophers seem to have intuitive resistance to theistic reductionism. Philip Quinn quickly dismisses Adams’ divine command theory with this line: “it seems clear to me that being morally wrong and being contrary to the commands of a loving God are distinct properties (2013, p.82).” And Wielenberg says,

> Although naturalism and supernaturalism are strikingly different in some respects, they are importantly alike in that they are both reductive in nature, declaring that ethical properties ultimately turn out to be properties of some other kind. My main concern about such views is captured by William FizPatrick's remark that they "secure the 'reality' of ethical facts and properties only by turning them into something else and deflating them in the process". (2014, pp.14-5)

So Wielenberg rejects both natural and theistic reductionism on the same basis, that they both in same way fail to capture essential aspects of moral properties. This is not a far cry from the

complaint of Moreland and Craig that natural properties lack an essential aspect of moral properties – *moral normativity* – and so reducing moral properties to natural properties is a kind of distortion or confusion. It is just that Wielenberg applies the complaint across the board, to theistic reductionism also.

These philosophers, however, seem more exceptional. Waghorn presents the view that seems to be the mainstream one among philosophers writing on theistic moral arguments. He compares reductive voluntarism and natural reductionism and comments:

> In the case of the former, if we … leave Euthyphro problems aside, the difficulty is always liable to be the ontological claim (that God exists), rather than the grounding claim (that God grounds objective ethics). In the case of the latter, matters are reversed; the hard doctrine is not the ontological one (that natural properties exist), but the grounding claim (that natural properties can ground rightness and wrongness) (2015, p.158).[26]

The mainstream view is that Euthyphro objections are the main objections to theistic reductionism, a position which is otherwise a satisfying means of accommodating moral properties. The thought is that it is, in some sense, perfectly intelligible or transparent how it is that theistic properties could be the reduction base of natural properties. We might have concerns about what such an account would entail in worlds where God commands us to do horrendous things, but these concerns seem to strike divine command theory further afield than at the heart of its moral-to-theistic reductive claims. Whereas with natural reductionism, the struggle is to make the relevant moral-to-natural reductive relationships themselves intelligible or transparent. So when theistic reductionists make heavy weather about the intuitive resistance to natural reductionism, it is not obvious that the same weather blows back against them.

[26] Grounding here should be understood to mean reduction.

§4.2 THE CONCEPTUAL GAP ARGUMENT

I hold that the intuitive resistance to natural reductionism can be expressed in a substantial argument. As already hinted, the argument is one that concerns conceptual constraints on reduction. It moves from conceptual considerations to metaphysical considerations. Laskowski agrees that the sense of incredibility in ethics involves such conceptual-to-metaphysical moves; "Many influential concerns about Reductivisms from Robust Realists [non-reductionists] tend to move from premises about normative concepts, thoughts, or language to metaphysical conclusions about the falsity of Reductivism (2019, p.95)." That is, these opponents of natural reductionism are not simply directly intuiting a metaphysical truth – that natural reductionism is false. Rather their intuitions are explained, first and foremost, by relations between moral and natural *concepts* from which metaphysical conclusions are inferred, even if only implicitly. But why think this?

Firstly, some opponents of natural reductionism explicitly appeal to conceptual considerations. When Parfit explains his view of why the reduction of moral properties to natural properties does not succeed in the way that the reduction of heat to molecular motion does, he appeals to differences between the concept of heat and the concept of moral properties like rightness and wrongness (2013, pp.301-2). And below I will show evidence of conceptual considerations in the writings of theistic moral arguers.

Secondly, there are parallels between intuitive resistance to natural reductionism in meta-ethics and intuitive resistance to natural reductionism in philosophy of mind. As Laskowski observes: "… there is a seemingly analogous sense of incredibility that many philosophers of mind experience while attempting to believe claims concerning the reducibility of phenomenal

consciousness (2019, p.98).” Scholars are increasingly seeing parallels between resistance to reductionism in philosophy of mind and to reductionism in meta-ethics. For instance, Mehta proposes to subsume these arguments under one more general phenomenon, which he calls the explanatory gap:

> That we, like Descartes, can conceive of phenomenally unconscious physical duplicates of ourselves; that we, like Moore, can coherently ask whether what we desire to desire is truly good; … I regard these facts, and many more, as mere instances of a general and philosophically pervasive phenomenon that I call the **explanatory gap.** Explanatory gaps are often thought to foreclose reductive possibilities wherever they appear (2017, p.1).

But in philosophy of mind, this intuitive resistance is expressed in argument. Paradigmatically, it is expressed in Chalmers' conceivability argument, which explicitly moves from conceptual considerations to metaphysical, anti-reductive conclusions. Indeed, Mehta argues that what unites all cases of explanatory gaps – all cases where a sense of incredibility is prone to occur – is the nature of the concepts involved. In light of these parallels to philosophy of mind and other areas, it would not be surprising if the sense of incredibility in ethics could be articulated in conceptual-to-metaphysical arguments. Although there could be various ways to formulate this argument, they would all fit the following broad schema:

The Conceptual Gap Argument

Conceptual Constraint: Moral properties reduce to natural properties only if moral concepts have a certain kind of conceptual connection *c* to natural concepts.

Moral-Natural Gap: It is not the case that moral concepts have conceptual connection *c* to natural concepts.

Therefore, it is not the case that moral properties reduce to natural properties.

Differing instantiations of the schema will give a different account of the conceptual connection *c* at issue in the argument. Note that this argument is not question-begging. A natural reductionist could, in principle, accept either premise while still being a good natural reductionist. So here we have a substantive, potentially dialectically effective argument. Of course, in order for this argument to not be turned against reductive voluntarism, moral arguers will need to endorse:

> **Moral-Theistic Bridge:** It *is* the case that moral concepts have conceptual connection *c* to theistic concepts.

We can see echoes of this argument or its premises in the history of meta-ethics. Moore's argument against natural reductionism seems to fit this schema, although his focus was on words rather than concepts. Moore (1903) claimed that moral properties could only reduce to natural properties if moral terms could be defined in natural terms (a kind of conceptual constraint). But his 'open question' test seemed to demonstrate that moral terms and natural terms could not be so inter-defined (a kind of moral-natural gap claim). And Hume (1739) famously claimed that there was an inferential gap between *is* and *ought* – easily read as the claim that there a sort of conceptual connection missing between moral claims and non-moral claims.[27] This would imply a kind of natural-moral gap. As we shall see, theistic moral arguers make claims that have a strong affinity to these Moorean and Humean theses.

[27] Although what Hume himself actually meant by there being a gap between *is* and *ought* is contested among interpreters. An alternative interpretation is given by MacIntyre (1969).

Clearly, the argument makes a distinction between moral, natural, and theistic concepts. Moral concepts in this context are concepts like MORALLY GOOD, MORALLY BAD, MORALLY RIGHT, and MORALLY WRONG. These are the so-called *thin* moral concepts. Paradigmatic theistic concepts concern God and divine commands. And paradigmatic naturalistic concepts concern microphysics, macrophysics, and human mental states – individual and social. That is, natural concepts concern at least the entities included within the core resources of naturalism. But as I noted in §2.3, moral arguers should also allow naturalists to appeal to non-moral normative properties. These properties seem neglected in the above quotes from moral arguers.

If they were brought into view, one could not claim that natural properties entirely lack normativity. One could, however, still claim that our concepts leave a gap between the kind of normativity natural properties are capable of and full-blown *moral* normativity. And we will see below an example where Craig does this – he appeals to a gap between prudential normativity and moral normativity. But since my criticisms of the Conceptual Gap argument do not appeal to non-moral normative properties, the reader is welcome to restrict natural properties to the core resources of naturalism.

Before turning to the textual evidence from moral arguers themselves, I will say a bit more about what sort of connection the conceptual connection *c* would be in the argument. It is clearly supposed to be an *a priori* connection. We are supposed to see the connection between two concepts c_1 and c_2 from the armchair. So take some proposition *p* that employs c_1 and some proposition *q* that employs c_2. If the right connection is there, and *p* and *q* employ the concepts in the right way, then it is supposed to be *a priori* that $p \supset q$. Consider that it is *a priori* that if Harry is a bachelor, then Harry is single. That entailment is afforded by the way those

propositions use the concepts BACHELOR and SINGLE, and the relationship between those concepts. The exact relationship that the moral arguer posits for connection *c* might not be the exact one that obtains between BACHELOR and SINGLE, but it is something in that ballpark.

Levine distinguishes between three alleged sources of *a priori* knowledge:

> The first source is logical form. I know *a priori* that I cannot both quit and not quit my job, because I know that no statement of that form can express a truth; it's a logical contradiction … What's more controversial … is the next alleged source of *a priori* knowledge, what I'll call "semantic form," or, more generally, information concerning the application conditions of concepts. So if I know that all bachelors are unmarried *a priori* it's because I know that nothing counts as a bachelor unless it's unmarried… Kant thought there was a third type of *a priori* knowledge, the synthetic *a priori.* Both logical and semantic form yield analytic *a priori* judgments, so if there is a synthetic *a priori,* it must come from another source. The most compelling example is mathematics, where it seems as if there is a direct apprehension of what must be true in some domain and it isn't based obviously on either logical or semantic form (2001, p.41).

The kind of *a priori* knowledge of $p \supset q$ that connection *c* is supposed to provide is clearly the second kind Levine identifies. It is something to do with the semantical or informational properties of the concepts that establishes the *a priori* link between *p* and *q.* Connections of this broad sort are just what I will call conceptual or analytic connections or entailments.

It is clear that the argument would have little promise if either of the other two sources of *a priori* knowledge were in view. Consider the first source, based on logical form. While it would be easy to establish that natural propositions lack any connections with moral propositions purely by virtue of their logical, syntactical form, it would be similarly easy to show that theistic propositions lack those connections too. After all, whatever *a priori* connections there might be between [God prescribes ϕ] and [it is obligatory to ϕ] it is clear that none are due to the logical form of the propositions. Now consider the third source, the *synthetic a priori*. If there are any

such connections of this sort, it is plausible that they obtain between some natural propositions and some moral propositions. For moral truths like [if *A* tortures *B* for fun, then *A* has done something morally wrong] are, like Levine's example of mathematical truths, among the prime candidates for synthetic *a priori* knowledge. Charity therefore provides further reason to think that theistic moral arguers have conceptual connections in mind, insofar as they have this argument in mind at all. Below I provide evidence that they do have it in mind.

4.3 TEXTUAL EVIDENCE

To begin with, let us see where theistic moral arguers make links back to Humean and Moorean claims, which as we saw above, echo claims made in the conceptual gap argument.

Evans gives a very direct and approving nod to Moore:

> ... as G. E. Moore famously argued, the identification of the good with any particular natural property or cluster of properties can always be challenged. Just because human beings (or some human beings) desire some outcome, it does not follow that the outcome is in fact desirable (2010, p.118).

Notice what is wrapped up in this short paragraph. Evans is objecting to the identification of a moral property – goodness – with any natural property or properties. That is, he is trying to establish the metaphysical conclusion that natural reductionism is false.[28] To establish that conclusion he gives an example that he thinks will speak for any substitute example. From human beings desiring something, it does not follow that the outcome is desirable, i.e. good.

[28] One might think that Evans is only objecting to natural reductionisms that posit an identity relationship between moral properties and natural properties. Even if that were so, his argument matches the schema of the conceptual gap argument, it merely targets a narrower form of natural reductionism. But it is likely that his use of identity language is used expansively, to cover all reductions. After all, he uses identity to cover cases where a moral property is purported to reduce to multiple natural properties – a case that I would describe as constitution rather than identity.

But it is *propositions* that 'follow' each other. So what Evans is saying is that there is no entailment between a natural proposition like [human beings desire *x*] and a moral proposition like [*x* is good]. He seems to endorse the Moral-Natural Gap. And since this is taken to be a problem for natural reductionism, he implicitly endorses the Conceptual Constraint.

The very first statement of the major objection to natural reductionism that I listed in this chapter was in a quote by Moreland and Craig. In it, they give a nod to Hume's law when they say that natural reductionism confuses an *ought* with an *is*. Could they mean that natural propositions fail to have the right conceptual connection to moral propositions? There is good evidence that Craig, at least, means something like that.

Consider what Craig says in response to Mark Murphy's critique of his moral argument. Murphy suggests that a naturalist might accommodate the moral value of human beings by appeal to prudential value and the human capacity to take a reflective, objective stance. But Murphy anticipates how Craig might object to this suggestion:

> One might, on Craig's behalf, put the worry as follows. *There is such a thing as prudential value* and *humans have the capacities of reflective and objectivizing intelligence* do not together logically entail there is such a thing as moral value. In other words, one can, without contradiction, affirm both there is such a thing as prudential value and humans have the capacities of reflective and objectivizing intelligence while denying there is such a thing as moral value (2009, pp.125-6).

Murphy, however, claims that if Craig takes this line of attack, then he will undermine his own theistic position:

> But if the standard for explanation here is logical entailment - p explains q only if p logically entails q - then Craig's theistic account fails as well. For God exists plainly does not logically entail moral value exists ... (2009, p.126).

Let us remain neutral, for a moment, on how to interpret "logically entails" and see where Craig is left after these comments. He is left with a couple options: (a) he can deny that logical entailment (whatever that is) is the standard he would hold the naturalistic account of moral properties against or (b) he can deny that his own theistic account fails to meet that standard. Of course, he could also deny both. But tellingly, Craig goes for option (b) *only*; he claims that God's existence "most certainly" *does* logically entail that moral value exists (2009). He takes no issue with Murphy's interpretation that logical entailment is the standard he is applying to natural reductionism. This strongly suggests that Craig demands that natural reductionists present some natural proposition that "logically entails", in some sense of that phrase, a moral proposition.

What is this logical entailment, and why does Craig think it so obvious, and Murphy think it so implausible, that [God exists] logically entails [moral value exists]? We can make sense of all this by considering two of the sources of *a priori* entailment discussed above. The first source concerns logical, syntactic form and seems to be what Murphy has in mind. It is clear that the logical, syntactical form of [God exists] provides no reason to infer [moral value exists]. After all, inferring *q* just from *p* is not a valid logical move. This makes sense of why Murphy finds a lack of "logical entailment" there. The second source concerns conceptual, semantic content and seems to be what Craig has in mind. Craig is a theist who thinks that the concept of God is the concept of a being who, among having other attributes, is morally perfect. So the semantic content of the concept of God has some moral content. One can see why Craig would think, given that moral content, [God exists] conceptually entails [moral value exists]. (I make no comment, just yet, on whether Craig is correct in holding there to be such conceptual entailment.)

What we learn from this is that Craig accepts the Conceptual Gap Argument. He endorses the Conceptual Constraint, that is, he holds natural reductionism to the standard of showing a conceptual connection between natural propositions and moral propositions. And he endorses the Moral-Natural Gap, that is, he holds that natural propositions and moral propositions lack those connections. His conclusion: natural reductionism is false.

There is further evidence that Craig has this argument in mind. On his website, Craig (2011) responds to an argument cited by a questioner that features the following premise:

(a) You can't derive an 'ought' from an 'is'.

The argument is meant to challenge Craig's divine command theory, which might appear to violate this premise. Craig considers the suggestion that he respond by saying that this premise is true "only on naturalism". And Craig endorses this suggestion, saying that it is "absolutely correct that (a) is a problem for naturalism, not for theism… (2011)." Plausibly, Craig is endorsing the Moral-Natural Gap and the Moral-Theistic Bridge. The only *is* propositions from which we cannot derive moral *ought* propositions are natural propositions, according to Craig. And since this is cited as a problem for naturalism, it seems that Craig is endorsing the Conceptual Constraint. Here is further evidence that Craig endorses the Moral-Theistic Bridge:

> So how does Divine Command Theory derive an "ought" from an "is"? Well, it says that we ought to do something because it is commanded by God. That is deriving an "ought" from an "is." … it seems to me that Divine Command Theory's derivation of an "ought" from an "is," far from being objectionable, captures a central feature of moral duty and plausibly grounds it (2010).

Craig is saying that propositions about God's commands have the right sort of conceptual connection to moral propositions. And as a result, God's commands satisfy the conceptual

constraints on reduction. God's commands "ground" – are the reduction base – of moral obligation.

I am not alone in interpreting Craig along these lines. Consider what Baggett and Walls say about Craig's conditional premise in his moral argument:

> … most typically, when a conditional says "If God doesn't exist, then objective moral facts don't exist," its speaker means to convey approximately this, echoing Milliken: "Hold everything equal, assume the world just as it is, but exclude from the picture God's existence; in such a world, moral facts don't exist." To be conceivable, this scenario has to assume less than an Anselmian God, but go ahead and try to do it for a moment (2011, pp.98-9).

Note that Baggett and Walls (and Milliken) interpret this as to do with what is *conceivable,* which is typically interpreted as a matter of what is conceptually possible. This is backed up by Baggett and Walls' proviso that to be conceivable, we must assume less than an *Anselmian* God. But of course, the Anselmian God is precisely the God whose non-existence is supposed to be conceptually impossible. So the interpretation amounts to this. Were one to have a long conjunctive proposition with all the truths about the natural world, call it *n*, it would fail to conceptually entail any moral propositions (any *m*). That is, one could not conceptually entail any moral propositions from *n*. It is not *a priori,* on conceptual grounds, that $n \supset m$. This is the Moral-Natural Gap, and implicitly, the Conceptual Constraint.

One might note here a similarity between Craig's argument and Chalmers' conceivability argument in philosophy of mind. And indeed there *is* a similarity, and one to be expected if the intuitive resistance to natural reductionism in meta-ethics is akin to the intuitive resistance to physical reductionism in philosophy of mind. But one might think that this similarity throws up a couple problems for Craig or any moral arguer who wants to use this style of argument.

Firstly, Chalmers himself considers the prospects for running a version of his argument against the reduction of the moral to the natural and he considers the prospects to be dim. He has two reasons for thinking this: "First, there does not seem to be a conceivable world that is naturally identical to ours but morally distinct, so it is unlikely that moral facts are further facts in any strong sense. Second, moral facts are not phenomena that force themselves on us. When it comes to the crunch, we can deny that moral facts exist at all (1996, p.83)." We can dismiss Chalmers' second reason quickly. We are assuming in this thesis that a good independent case can be made for moral realism. Of course, that does not entail that the objective existence of moral properties is so obvious as to be *forced* on us. But it is not clear why their existence needs to be forced on us before we are willing to posit them and argue for their irreducibility to natural properties.

Chalmers' first reason seems to be that he simply denies the Moral-Natural Gap. He claims to not be able to conceive of a natural world like ours without moral properties. It is important, though, to keep in mind the different potential sources of *a priori* knowledge. It is compatible with the Moral-Natural Gap that there are synthetic *a priori* connections between natural propositions and moral propositions. So, if we take the question of whether a natural duplicate of our world without moral properties is conceivable to be the question as to whether $n \supset m$ is *a priori,* and we don't care about the source of the *a priori* knowledge in view, the moral arguer can concede, "yes, it is". The moral arguer can say that what is leading Chalmers to deny the Moral-Natural Gap is that he is focusing on what is *a priori* simpliciter, not what is *a priori* on conceptual grounds. It is only the latter that need be at stake in the conceptual gap argument against moral-to-natural reductionism.

Of course, this response depends on our being able to know the source of our *a priori* knowledge. And we might doubt that we have access to its source. At least, we might doubt whether we can tell the difference between analytic *a priori* knowledge and synthetic *a priori* knowledge (*a priori* knowledge based on logical form might be easier to discern). But for what it is worth, it seems to me that I *can* introspect, in some way, whether my a priori knowledge is analytic or synthetic. It seems to me that my knowledge that [if Harry is bachelor, then Harry is single] comes just from my competence with the concepts employed, whereas my knowledge that [If Harry tortures Sally just for fun, then Harry has done something wrong] does not. So, my criticisms of the conceptual gap argument will not hinge on finding any problem here. But anyone who does find a problem here will simply have extra reason to deny reductive voluntarism an advantage.

There is another problem that the comparison with Chalmers' argument might throw up. Chalmers' argument depends on a principle stating a connection between what is conceivable and what is metaphysically possible.[29] His argument is that it is conceivable that a physical duplicate of our world lacks phenomenal properties. And if that is conceivable, he claims it is metaphysically possible that a physical duplicate of our world lacks phenomenal properties. But if that is possible, then the phenomenal is something over and above the physical. It is, however, precisely on the grounds of the conceivability-possibility linking principle that Chalmers denies that God is metaphysically necessary:

> I think theist views of this sort are to be rejected. If the existence of such a god were knowable a priori, then such a god might exist necessarily, but if it is not, then one should conclude that such a being exists only at best contingently. I cannot respond to every argument for believing in a necessary god here, but they all rest on highly contentious premises, and once again, CP [the principle stating

[29] Chalmers actually distinguishes between several kinds of conceivability, and it is only some of those kinds that he claims has a connection to possibility. But those details need not concern us here.

> a link between conceivability and possibility] itself provides an argument against these views. The best way to defend the existence of a necessary god is to argue that a world without such a being is not even conceivable (2010, p.197).

Presumably, Chalmers does not find a world without God to be inconceivable. But Craig and others hold that God's existence is metaphysically necessary. So, if they want to run an argument like Chalmers' and maintain God's necessity, they have to claim that God's non-existence is inconceivable. But to support a claim like that is to run an ontological argument for God. No moral arguer wants their moral argument to depend on also running an ontological argument. This can either be taken as a problem for moral arguers or as a problem for the plausibility of my interpretation of them.

This problem provides an occasion to clarify an important difference between the Conceptual Gap Argument as I have framed it, and Chalmers' argument. The Conceptual Gap Argument does *not* depend on any principle linking conceivability and possibility. The Conceptual Constraint premise only asserts a link between conceivability (conceptual connection) and reduction. Chalmers' principle is a stronger principle than the Conceptual Constraint premise I have attributed to moral arguers. According to Chalmers, if it is conceivable that *x* exist without *y,* then not only does *y* not reduce to *x, y* is not caused to exist by *x* of necessity, nor does *y* co-exist alongside *x* of necessity. The Conceptual Constraint, on the other hand, merely claims that if it is conceivable that *x* exists without *y,* then *y* does not reduce to *x.* But it can still be that *y* is caused to exist by *x* of necessity, or that *y* simply co-exists alongside *x* of necessity. In other words, according to Chalmers, if it is conceivable that *x* exists without *y,* then *y* does not supervene on *x.* But as I argued in §3.1, supervenience is just a tracking relationship between *x* and *y* that can be explained by reduction but can also be explained by necessary causation as well (or just necessary co-existence). If there is a conceivability constraint on supervenience, then there is a constraint on reduction. But the reverse does not

follow. A conceivability constraint on reduction does not place such a constraint on supervenience.

Even aside from my different framing of the argument, Chalmers himself (2010, p.178) considers that anti-reductionist arguments could run on a narrower constraint than his constraint on possibility. That is, he considers how his argument against materialism could run even if he granted what he calls “strong necessities”. These are brute metaphysical necessities that do not seem logically necessary – that is, necessary *a priori.* If it is conceivable that God not exist in a world that is a physical duplicate of our world, yet God’s existence is necessary, then God’s co-existence with the physical universe is a strong necessity. Chalmers points out that the best candidates for strong necessities involve ontologically distinct entities. I.e., even if God exists (or abstract entities exist) in any world where there is a physical world, these entities are something over and above the physical world. So Chalmers’ linking principle could state that if it is conceivable that *x* exists without *y,* then either it is possible that *x* exist without *y* (and so *y* does not reduce to *x*) or there is a strong necessity linking *x* and *y* as two ontologically distinct entities (and so once again, *y* does not reduce to *x*).

In summary then, the similarities between the Conceptual Gap Argument and Chalmers’ argument neither show the Conceptual Gap Argument to fail in motivating a moral argument, nor show that it is implausible to interpret moral arguers as running the Conceptual Gap Argument.

The most explicit evidence that the Conceptual Gap Argument is in view comes from Craig's work. But this argument being in mind would also make sense of the various passages where moral arguers claim that it is "difficult to see" how some natural property could be the reduction

base of a moral property. Recall that Wainwright says, “it is difficult to see how purely empirical properties could have these characteristics built into them (2005, p.51).” Plausibly, what this is a report of is the difficulty in seeing any connection between the natural and the morally normative. That is, the two seem to easily come apart in our conceivings. What is being emphasised is the lack of transparent connection between natural propositions and moral propositions. This also explains why some of the passages cited at the beginning of the chapter are rather quick, and unelaborated. That makes sense if these authors just do not see any connection between the cited natural properties and the moral properties. It is hard to elaborate on just not seeing something. Suppose you tell me that there is a tight metaphysical connection between knowledge that *p* and feeling good about *p*. My reaction will be to announce that I just do not see the connection between those things.

Granted, Baggett & Walls (2016, p.54-78) claim to distance themselves from Craig’s approach. And it is true that, unlike Craig, their argument does not include a claim about a counterfactual atheist world. But the Conceptual Gap Argument does not require that when one checks for conceptual gaps between moral propositions and some natural proposition that the latter be a long description of an entire natural world. The natural proposition can be far simpler and localised than that. And Baggett & Walls seem to perform such checks between moral propositions and simple natural propositions. As an example, consider how they object to Frans de Waal’s account of moral obligations; “What he is referring to as “morality” does not seem to be any set of moral truths at all, but rather moral beliefs and practices at most. Although he identifies some necessary additions to animal behavior to arrive at “morality,” what he adds does not seem to be nearly enough (2016, p.156).” What they seem to be saying is that there is a gap between moral obligation and what this account describes. What is provided in the account is not enough to close that gap. So even Baggett & Walls plausibly appeal to conceptual

gaps.

I have shown that the Conceptual Gap Argument is plausibly what many moral arguers have in mind. Thus, there is more that can be said and appealed to beyond brute intuitions that are directly against natural reductionism. In the next chapter I assess the Conceptual Gap Argument's success in showing reductive voluntarism to have an advantage over natural reductionism. That will involve finding the most plausible type of conceptual connection *c* that a moral arguer could plug into the argument.

5

The Conceptual Gap Argument: Assessment

Recall that theistic moral arguers aim to motivate P2 [it is plausible that theism accommodates objective moral properties and more plausible that it does so than that naturalism does so]. One way they have tried to do that is by arguing that reductive voluntarism is more plausible than its naturalistic competitors. That requires it having an advantage over some *specific* competitors, like natural reductionism. In the previous chapter, I defended an interpretation of the major argument that moral arguers offer for preferring reductive voluntarism to natural reductionism. The argument moves from conceptual considerations to a metaphysical conclusion, as follows:

The Conceptual Gap Argument

Conceptual Constraint: Moral properties reduce to natural properties only if moral concepts have a certain kind of conceptual connection *c* to natural concepts.

Moral-Natural Gap: It is not the case that moral concepts have conceptual connection *c* to natural concepts.

Therefore, it is not the case that moral properties reduce to natural properties.

And I noted that in order for this argument to not be turned against theistic reductionism, moral arguers will need to endorse:

Moral-Theistic Bridge: It *is* the case that moral concepts have conceptual connection *c* to theistic concepts.

Much of the argument hangs on what the conceptual connection *c* is. I argued that it is an *a priori* connection. Further, it is an *a priori* connection grounded in the semantic, informational content of the concepts involved. So if proposition *p* has conceptual connection *c* to proposition *q,* it will be *a priori* that $p \supset q$, broadly similar to how it is *a priori* that if Harry is a bachelor, then Harry is single. That entailment is afforded by the way those propositions use the concepts BACHELOR and SINGLE, and the relationship between those concepts. Still, this leaves a lot underdetermined. *Prima facie,* there are lots of ways that concepts could be connected to each other. What we need is some more specific account of what the connection *c* is. Only then will we have a determinate version of the Conceptual Gap Argument to assess.

Morals arguers need to satisfy a few criteria when it comes to giving an account of *c.* Firstly, it needs to be independently plausible that *c* is a constraint on reduction. There is no use giving an account of *c* that succeeds in preserving the Moral-Natural Gap and the Moral-Theistic Bridge if there are widely accepted examples of reduction that do not feature *c.* Secondly, it needs to be plausible that *c* fails to obtain between moral concepts and natural concepts. There is no use giving an account of *c* that is lenient enough to cover all accepted examples of reduction, and lenient enough to preserve the Moral-Theistic Bridge, if it is so lenient that it fails to preserve the Moral-Natural Gap. Thirdly, there is no use giving an account of *c* that is covers the accepted examples of reduction and is stringent enough to preserve the Moral-Natural Gap if it is so stringent that it fails to preserve the Moral-Theistic Bridge.

I argue in this chapter that there is no account of *c* that satisfies these criteria. I start by

considering what I call a *compound-to-constituent* conceptual connection. I argue that either theistic concepts lack that connection to moral concepts, or they can be said to possess it only if natural concepts also possess it. I also argue that it is either implausible that compound-to-constituent connections are a constraint on reduction, or there is no motivation for including them within a constraint on reduction. Next, I consider *role-satisfaction* conceptual connections. These plausibly *are* part of a conceptual constraint on reduction. But to appeal to these connections, moral arguers need to provide an analysis of moral concepts. The analysis provided by Harrison is implausible, and the analysis provided by Adams fails to exclude natural reductionism.

§5.1 COMPOUND-TO-CONSTITUENT CONNECTION

Let us start with an idea introduced in the previous chapter that might seem promising. There we saw that according to Craig (2009), there is a conceptual connection between the propositions [God exists] and [moral value exists]. That is, there is a conceptual connection between a theistic proposition and a moral proposition. This seems to be a claim that the Moral-Theistic Bridge is satisfied. So what we can do is see what sort of conceptual connection there supposedly is between the relevant concepts in those propositions (GOD and MORAL VALUE) and see how the Conceptual Gap Argument fares if we let *c* be that connection.

Plausibly, what Craig has in mind is this: GOD is the concept of a morally perfect being.[30] So if that concept applies in a world, there is a morally perfect being in that world. If there is a morally perfect being in that world, then there is moral perfection in that world. If there is

[30] Of course, Craig thinks the concept has additional content, but it is at least the concept of a morally perfect being.

moral perfection in that world, then there is moral value. This is very similar to what someone would say to explain why [a bachelor exists] entails [singleness exists]. BACHELOR is the concept of a single (unmarried) being. So if that concept applies in a world, there is a single being in that world. If there is a single being in that world, then singleness exists in that world.[31]

Let us represent the connection from [God exists] to [moral value exists], and from [a bachelor exists] to [singleness exists] as $p \supset q$. Both the God case and the bachelor case have a similar structure. In both cases, what explains the connection is that the key concept in the proposition on the left-hand side of the conditional operator is constituted, at least partly, by the key concept in the proposition on the right-hand side. GOD is partly constituted by MORAL PERFECTION, i.e. MORAL VALUE, and BACHELOR is partly constituted by SINGLE. Let us call this kind of conceptual connection between propositions *compound-to-constituent* connection. For both GOD and BACHELOR here are compound concepts, constituted by more basic concepts, like MORAL VALUE and SINGLE. From a proposition using a compound concept, one can infer some propositions using a constituent concept of that compound concept. So let us treat *c* as compound-to-constituent connection and evaluate the Conceptual Gap Argument under that interpretation.

It might seem that Craig is clearly right that the Moral-Theistic Bridge is upheld on this interpretation. The moves leading from the concept of God as a morally perfect being to that entailing moral value seem unassailable. But it is less obvious that the concept of God *is* that of a morally perfect being than one might think. As Lambert argues, one way to test whether our concept of God is that of a morally perfect being, or even a morally good being, is to

[31] One might be uncomfortable with the implications of this talk of moral value or singleness *existing.* It might seem to treat these things as substances. I take it that Craig is speaking loosely here. One can read the above as concerning the *instantiation* of properties instead.

describe some being that is not morally good, but is otherwise as impressive as theists normally think God is, and ask whether we would still apply the concept GOD to that being. So Lambert (2020, p.4) asks us to imagine that

> There is a necessarily existing incorporeal being who is eternal, perfectly free, omnipotent, omniscient, and who created the universe and everything in it. This being sustains the universe from moment to moment. Moreover, this being is the object of religious worship, answers prayers, and occasionally performs miracles. This being also loves all created beings, is deeply concerned for their well-being, always tries to promote their well-being to the best of his ability, and so on.

Next he asks us to suppose that the world in which this being exists is one in which moral error theory is true, meaning that this being is not morally good. For on error theory, there are no moral properties. This being is still loving, concerned for the well-being of creatures, and still all-powerful, all-knowing, and so on, but this being is not morally good. Lambert holds that we would still apply GOD to this being. And that seems plausible. After all, it would be strange to say that a world with such a being, even given error theory, is a world in which atheism is true. So MORAL VALUE does not seem to be a constituent concept of GOD.

What Craig might say in response is that Lambert's argument tests the wrong concept of God. Clearly, there is some common, ordinary concept of a divine being that does not have MORAL PERFECTION built in, nor the concepts of any other divine attribute entertained by philosophical theologians. This common concept of a divine being is what allows, say, views to be described as polytheist when those views posit several immaterial beings that are superhuman in power or knowledge, but fall short of omnipotence, omniscience, or the other traditional perfections. Such views could only be described as positing many gods, and so counting as polytheist, if the concept of a god used in describing those views is very permissive. But Craig might say that the concept of a god that he is working with is not *that* concept. He is working with the

concept of the *Anselmian* God. That is, the concept of the greatest conceivable being. And, he might go on, it is a conceptual truth that the greatest conceivable being is morally perfect. So there is, after all, a conceptual connection between the concept of God and the concept of moral value, so long as we focus on the Anselmian concept of God.

It is, however, implausible that there is a *conceptual* connection between the greatest conceivable being and moral value. Or at least, if Craig maintains that there is a conceptual connection, it will be hard to preserve the Moral-Natural Gap. Let me explain. There is, I grant, some *a priori* connection between [*x* is the greatest conceivable being] and [*x* is morally perfect]. After all, it seems to be an *a priori* truth that moral goodness is great-making. So the upper limit of greatness would seem to require the upper limit of moral goodness, that is, moral perfection. But we should recall again that there are different sources of *a priori* connection. And given that Craig wants to uphold the Moral-Natural Gap, he will have to maintain that the connection between moral goodness and greatness is a synthetic *a priori* connection. For consider other *a priori* connections between certain properties and greatness, like that between power and greatness or knowledge and greatness. Power and knowledge are both naturalistically acceptable properties. And greatness is some sort of *normative* property. So these seem to be *a priori* connections between natural properties and normative properties. But then if these are *analytic* connections, why aren't the *a priori* connections between natural properties and *morally* normative properties *analytic*? It would seem unprincipled to hold that truths connecting natural properties are greatness properties are analytic, but truths connecting natural properties and moral properties are synthetic.

Craig can hold that both are analytic, in which case he seems to lose the Moral-Natural Gap, or he can hold that both are synthetic, in which case he seems to lose the Moral-Theistic Bridge.

He cannot have it both ways. So Craig still needs to offer some reason to think that GOD has MORAL VALUE built in – a reason that works in the dialectical context of running a moral argument for theism. What he might say is that the concept of God he is working with is neither the ordinary, everyday concept, nor the Anselmian concept. Rather, it is a concept that has emerged from a certain theological, philosophical tradition. The concept simply stipulates that God is a being with moral perfection. But if that is the move, then it is unclear why naturalists cannot simply stipulate their own concept that has MORAL VALUE built in. Perhaps their concept is ROD – the concept of an ordinary human being with a finite degree of moral goodness. Once again, the Moral-Natural Gap is jeopardised.

Suppose, however, that naturalists, for whatever reason, cannot make that stipulative move. If that's the case, then Craig has a concept of God that includes MORAL VALUE as a component. He can thus uphold the Moral-Theistic Bridge and maintain that connection *c* is the compound-to-constituent connection. He can seemingly do this without abandoning the Moral-Natural Gap. Even so, what we need to ask is whether *c,* so understood, can serve as an independently plausible constraint on reduction. And the answer is that it cannot. There are widely accepted examples of reduction that do not feature *c* between the relevant concepts and propositions.

Consider the reduction of water to H_2O. If compound-to-constituent conceptual connection were a constraint on reduction, then [H_2O exists] would have to entail [water exists] in the way that [Harry is a bachelor] entails [Harry is single]. But plainly no such conceptual connection obtains. WATER is not a constituent concept of the concept of H_2O. There could very well be alien chemists who have the concept of H_2O without ever having formed the concept of water. There may simply not be any H_2O in liquid form on their world. So Craig is unable to sustain the Conceptual Constraint on this interpretation of *c*.

Craig, however, might grant that compound-to-constituent connection is not a necessary condition on reduction. What he might say instead is that *c* should be understood disjunctively. That is, a necessary condition on reduction is that there be compound-to-constituent connection *or* some other connection yet to be spelt out. So the reduction of moral value to God nature satisfies the Conceptual Constraint via compound-to-constituent connection, and the reduction of water to H_2O satisfies it via some other conceptual connection.

It may very well be the case that there is a disjunctive conceptual constraint on reduction. The problem for Craig, however, is that there is no reason to include compound-to-constituent connection within that disjunctive constraint. There simply is not any previously accepted example of reduction where compound-to-constituent connection is present. Or rather – because that is not strictly correct – there are no examples that feature compound-to-constituent connection where the reduction goes in the right *direction* to be of help to Craig. Let us consider again the claims that Craig is making. He claims that [God exists] conceptually entails [moral value exists]. The direction of this entailment is $G \Rightarrow MV$, which is what makes it a case of compound-to-constituent connection. But Craig claims that moral value reduces to God, or his nature. According to Craig, God is the Good itself. So the direction of the desired reduction is $MV \rightarrow G$. Note, then, that the direction of the entailment is opposite to the direction of the reduction. But there are no other accepted cases of reduction where the directions are opposite in that way.

Accepted cases of reduction that feature compound-to-constituent entailment have the direction of the entailment and the direction of the reduction aligned. Consider the case of knowledge. It is widely accepted that [Maya knows that *p*] entails [Maya believes that *p*]. This seems to be a

case of compound-to-constituent entailment. For KNOWLEDGE is a compound concept that has BELIEF as a constituent component.[32] The direction of the entailment is $K \Rightarrow B$. It does not work in the other direction; [Maya believes that *p*] does not entail [Maya knows that *p*]. But if there is any reduction here, its direction is also $K \rightarrow B$. Knowledge reduces to belief (that is true, justified, and meets some other condition), not belief to knowledge. What we have here is an alignment between conceptual constitution and metaphysical constitution. KNOWLEDGE is constituted, in part, by BELIEF and knowledge itself is constituted, in part, by belief itself. There do not seem to be any cases where that alignment is violated.

The best that Craig can do is try and achieve alignment between his conceptual claims and his metaphysical claims. That is, he can switch to saying that [Moral value exists] entails [God exists], or he can switch to saying that God reduces to moral value. But neither of these options is attractive. It seems plainly false that [Moral value exists] entails [God exists]. It simply is not the case that MORAL VALUE includes GOD as a constituent. Craig seems to admit this when he states that divine command theorists like himself are not trying to give theological definitions of moral terms or concepts:

> Theists are *not* offering theistic definitions of moral terms. We are not saying, for example, that "obligatory" means "commanded by God" or some such thing. (2011).

Could Craig say that God reduces to moral value, rather than moral values reduces to God? Perhaps he could, but that would just be to give up on reductive voluntarism. He would cease to give an accommodation account of moral properties and would instead be giving an accommodation account of God! But perhaps what Craig might say is that I am wrongly

[32] At least, this is a fairly standard way to think about knowledge.

attributing to him the kind of reduction that has a direction of fundamentality. Not all kinds of reduction do. As I acknowledged in §3.1, identity relationships have no direction of fundamentality. So Craig could say that God is nothing over and above moral value, and moral value is nothing over and above God. Thus the direction of reduction is both $MV \rightarrow G$ and $G \rightarrow MV$. Craig gets alignment between the direction of conceptual entailment, $G \Rightarrow MV$ and one of the two directions of reduction. And since Craig explicitly uses the language of identity (2009), this looks like an attractive move.

The problem with this move, however, is that it simply creates another dissimilarity between Craig's attempted reduction and every other accepted case of reduction. Craig would be claiming that there is an identity relationship between God and moral value, and that GOD has a compound-to-constituent connection to MORAL VALUE. But there are no other cases of identity where one of the concepts has a compound-to-constituent connection to the other. Clark Kent and Superman are identical but even if CLARK KENT and SUPERMAN share some sort of conceptual connection, it is not a compound-to-constituent one. Neither of these concepts is a component of the other. Some people think that faith is identical to trust. But they wouldn't say that FAITH is a component of TRUST, or vice versa. Rather, they would say that the two concepts are univocal – they have equivalent semantic content. This is precisely what you cannot say of two concepts related by compound-to-constituent connection. KNOWLEDGE is not univocal with BELIEF. In sum, compound-to-constituent connection may be part of a disjunctive conceptual constraint, but if it is, that is no help to Craig.

I conclude that appealing to compound-to-constituent entailment as *c,* even if it is just one of the disjunctive options included in *c,* cannot support the Conceptual Gap Argument. Moral arguers will have to appeal to some other sort of conceptual connection. The most promising

one they can appeal to is role-satisfaction, which I explain below.

§5.2 THE NEO-CLASSICAL PICTURE

We saw above that the widely accepted reduction of water to H_2O causes problems for the claim that compound-to-constituent entailment is a necessary condition of reduction. In fact, the water case caused problems for a once widespread view about the conceptual constraints on reduction that we might call the *classical* view.[33] According to the classical view, there is a strong symmetry between conceptual relationships and metaphysical relationships. If two properties are identical to each other, their respective concepts must be equivalent to each other. They need to have the same semantic, informational content. If one property is going to be constituted out of several, then the concept of that one property will have to be composed of the concepts of the other properties. There will need to be a compound-to-constituent connection. So if faith is identical to trust, then FAITH is univocal to TRUST. And if knowledge is constituted by justified true belief, then KNOWLEDGE is composed of JUSTIFIED, TRUE, and BELIEF.

On the classical view, there is no other route to reduction except concepts being univocal or having compound-to-constituent connections. G E Moore (1903) held to the classical view. His objection to the identification of moral properties with natural properties was that moral concepts like GOOD do not seem univocal to natural concepts like PLEASURABLE. But the water case, brought to the centre of philosophical attention by Hilary Putnam (1973) and Saul Kripke (1980), seemed to provide a third way to reduction. As already observed above, WATER does

[33] And it is not just the water case but also the reduction of heat to molecular motion, and a host of other scientific reductions that came to be known as necessary *a posteriori* reductions.

not have the concept of H_2O as a constituent. But nor does it seem univocal to the concept of H_2O. Yet water does reduce to H_2O.

There has been fierce disagreement over the lessons to be drawn from the water case. One prominent view is that the water case shows that some reductions have no conceptual constraints whatsoever. This is perhaps the mainstream view, or at least it was for a while. On this view, the fact that WATER refers to H_2O has absolutely nothing to do with the semantic, informational content of WATER. *A priori* conceptual analysis of WATER does not reveal any necessary conditions of *x* being water. As Brandon-Mitchell (2008, p.24) puts it, Kripke taught a generation of philosophers that "All analysis did was make explicit fallible beliefs about things." The prevalence of this view re-opened the door to views like natural reductionism, which had seemed implausible on the classical view. The stance that Brink takes below became standard:

> Because the ethical naturalist rejects the semantic test of properties, the failure of synonymy between moral and natural terms does not interfere with the theoretical identification of moral and natural properties. There is nothing metaphysically queer about an ethical naturalist's identification of moral and natural properties. So construed, ethical naturalism requires theoretical identifications, for instance, water = H_2O, temperature = mean kinetic molecular energy, light = electromagnetic radiation. Hedonistic utilitarianism, for example, can be construed as making one such naturalistic claim; it claims that rightness = the maximization of pleasure (1989, p.176).

So the reduction of moral properties to natural properties was taken to be modelled on reductions like the water case, where this was understood as a case when there are no conceptual, *a priori* constraints.

There are, however, defenders of what we might call the *neo-classical* picture who claim that there still are conceptual constraints in the water case, they are just not as stringent as the

classical picture held that they were. Neo-classicalists can differ on the details of what this conceptual constraint is, but they all seem to agree on one minimal claim. That is, the concept of the property to be reduced must be "gappy". Here is Parfit expressing this view using the example of heat:

> In the relevant sense, 'heat' means, roughly:
>
> the property, whichever it is, that can have certain effects, such as those of melting solids, turning liquids into gases, causing us to have certain kinds of sensation, etc.
>
> This concept, we can say, has an explicit gap that is waiting to be filled, since this concept refers to some property without telling us what this property is. This concept refers to this property indirectly, as the property that can have certain effects, such as those of melting solids, etc. This feature of the concept of heat allowed scientists to fill this gap, by discovering that molecular kinetic energy is the property that can have these effects (2013, pp.301-2).

According to Parfit, heat could only be reduced to molecular kinetic energy because of the lack of specificity in the concept of heat. The concept left room for the relevant property that is the reduction base of heat to be discovered by scientists. Adams seems to take a similar interpretation of the water case: "What I must know, at least implicitly, about water in order to understand the sense of the word 'water', and so to be a competent user of the word, is that if there is a single chemical nature shared by most of the stuff that I and other English-speakers have been calling "water", then, of necessity, all and only stuff of that nature is water (1999, p.15)." That is, according to Adams, the concept of water leaves it open as to what the relevant chemical nature is. This is what allows water to be reduced to H_2O. Adams describes the concept of water as laying out a *role* for a property to satisfy. The concept does not specific which property satisfies that role, but the roles leave space for something to satisfy it.

The most well-developed neo-classical views are those of Chalmers (1996) and Jackson (1998),

though I will only draw from Chalmers. Like Adams, he holds that successful reduction requires satisfying some role left open by the relevant concept.

> What is it that allows such diverse phenomena as reproduction, learning, and heat to be reductively explained? In all these cases, the nature of the concepts required to characterize the phenomena is crucial. If someone objected to cellular explanation of reproduction, "This explains how a cellular process can lead to the production of a complex physical entity that is similar to the original entity, but it doesn't explain *reproduction,*" we would have little patience - for that is all that "reproduction" *means.* In general, a reductive explanation of a phenomenon is accompanied by some rough-and-ready *analysis* of the phenomenon in question, whether implicit or explicit (1999, p.44).

According to Chalmers, the concept of reproduction is (roughly) that of the production of a complex physical entity that is similar to the original entity. It leaves unspecified what the process of production is. All it specifies is an open causal role, and it is the job of scientists to discover what satisfies it. If they discover a cellular process that produces a complex physical entity that is similar to the original, they have successfully reduced reproduction to that cellular process. For Chalmers, the situation is different for the concepts of phenomenal states: "Phenomenal states, unlike psychological states, are not defined by the causal roles that they play. It follows that explaining how some causal role is played is not sufficient to explain consciousness. After we have explained the performance of a given function, the fact that consciousness accompanies the performance of the function (if indeed it does) remains quite unexplained (1996, p.47)." Reduction requires the relevant concept to have an open role to be satisfied and according to Chalmers, the concepts of phenomenal states do not meet that criteria.

Neo-classicalists like Chalmers hold to there being some conceptual constraint on reduction. But so far, nothing has been said about whether this preserves any sort of *a priori* connection between the concept of the phenomenon being reduced and the concept of the reduction base.

Indeed, for all that has been said so far, it might seem that such an *a priori* connection has been denied. After all, while it may be an *a priori* matter to determine whether the concept of water has an open role, neo-classicists seem happy to leave it to the empirical, *a posteriori* work of scientists to determine what satisfies that role.

While it is correct that neo-classicalists do not attempt to establish any *a priori* link between simple propositions like [water exists] and [H_2O exists] they can maintain – and Chalmers explicitly does maintain – that there are other relevant *a priori* links. What we can do *a priori* is conceive of some possible world *w.* And then we can ask whether, in *w,* there is some *x* that plays the water role. If there is, then water is *x* in that world. Where the *a posteriori* work steps in is in determining whether the actual world is *w,* or some other world where water is some different *x.* Only empirical investigation could tell us that the world we are in is one where H_2O is the *x* that plays the water role. There is*,* then, a form of *a priori* conceptual constraint that neo-classicalists can hold applies even in cases like the water case.

> **Role Satisfaction:** Moral properties reduce to natural properties via role-satisfaction in *w* only if moral concepts include an open role and it is *a priori* that given a relevantly full description of *w,* natural properties satisfy that role in *w*.

The above, however, has involves some over-simplification. Neo-classicalists need to say more in order to show that the water case does not constitute an exception to even this type of conceptual constraint. For it might seem that it is not *a priori* that if *x* satisfies the water role in some non-actual world *w,* then *x* is water. That, after all, seems to be the intuition that Putnam's twin earth thought experiment prompts. In that thought experiment, we are asked to imagine another possible world where there is a planet functionally identical to our earth. On

this twin earth there is a chemical structure that plays the water role (it appears translucent blue, falls from clouds etc.), only this chemical structure is not H_2O but rather XYZ. The question is whether our WATER concept refers to XYZ in that other world. And our intuition is supposed to be that it does not. WATER only ever refers to H_2O, no matter what possible world we are considering. So here we have a situation in which WATER includes an open role, and it is *a priori* that in a particular world, some natural properties, XYZ, satisfy that role, yet it is not the case that water reduces to XYZ in that world, or any other.

Chalmers' response to this objection is two-fold. Firstly, he claims that we are mistaken if we think that the open role in our WATER concept just includes things like *it appears translucent blue, falls from clouds etc*. It also includes some *indexical* content. Water is the stuff *around here,* that falls from *our* skies, fills *our* oceans, and so on. Indeed, Adams includes such indexical content in his own explication of the water role: "if there is a single chemical nature shared by most of the stuff that *I and other English-speakers* have been calling "water" … (my emphasis) (1999, p.15)." So water, in any world, is only ever the chemical structure that satisfies the indexical conditions of the possessors of the WATER concept. So when we are considering *our* WATER concept, water is only ever the chemical structure that fills the oceans *around here,* in the *actual* world. When that same chemical structure is present in some other world, even if does not fall from the sky, or fill any oceans in that world, that chemical structure is still water. And even if some other chemical structure in some other world does fall from the sky, or fill oceans, that chemical structure is not water. According to Chalmers this explains our intuitions about the twin earth case. Since we discovered that H_2O plays the water role on our planet in our world, intuitions about the application of the WATER concept reflect that the referent of WATER across worlds is tied to what is the case in our world.

Secondly, in light of the indexical content of some concepts like WATER, Chalmers claims that we need to distinguish two different kinds of *a priori* reflection on a world *w.* We can conceive of *w* as actual and we can conceive of *w* as counterfactual. Consider that before scientists discovered H_2O they could have conceived of different epistemic possibilities for how the world could turn out. They could have conceived of the possibility that H_2O would turn out to play the water role, or they could have conceived of the possibility that XYZ would turn out to play the water role. Now what seems *a priori* is that had XYZ turned out to be the chemical structure that fills our oceans, then XYZ would be water, in our world and every world. So when we conceive of a world where XYZ plays the water role, and when we conceive of that world as the *actual* world, it is indeed *a priori* that XYZ is water. It is only when we conceive of a world where XYZ fills the oceans of some twin earth, not as the actual world but as a counterfactual world, where XYZ fails to be water. So really, the conceptual constraint needs to be:

> **Role Satisfaction*:** Moral properties reduce to natural properties via role-satisfaction in *w* only if moral concepts include an open role and it is *a priori* that given a relevantly full description of *w considered as actual,* natural properties satisfy that role in *w*.

This modification concerning the indexical content of certain concepts will not be important in the below discussion, so the reader can simply ignore that complexity if he or she wishes. The important thing for our discussion will be the need for some open role or other in the concept of the phenomenon to be reduced. I bring up the issues with indexical concepts only to demonstrate that the water case does not obviously show that some forms of reduction entirely lack *a priori,* conceptual constraints. Chalmers' account is a sophisticated defence of such

constraints, carefully designed to handle cases like water and heat. It is the best account that theists could draw from in constructing an independently plausible conceptual constraint on reduction. Indeed, to pull together the above discussion, it seems like the best account of *c* and the Conceptual Constraint premise in the Conceptual Gap Argument will be the following disjunctive one:

> **Conceptual Constraint:** Moral properties reduce to natural properties only if moral concepts have a certain kind of conceptual connection *c* to natural concepts, where *c* is either:
>
> > **Univocation:** moral concepts are univocal in semantic content to some natural concepts, or
> >
> > **Compound-to-Constituent:** moral concepts are composed of natural concepts, or
> >
> > **Role Satisfaction:** moral concepts include an open role and it is *a priori* that in our world, when it is considered as actual, and given a relevant full description, natural properties satisfy that role.

I am not arguing that the Conceptual Constraint premise is true under this interpretation, although I suspect it is. Any further evaluation of the premise would take us far afield into a deep, independent investigation into the nature of reduction and conceptual analysis. I will not undertake that investigation here. So I will grant this premise from here on. Any further objections I have to the Conceptual Gap argument will be objections to the Moral-Natural Gap

or the Moral-Theistic Bridge.

Of course, above I argued that the most promising candidate for upholding the Moral-Theistic Bridge via compound-to-constituent connection fails. And no contemporary moral arguer holds that the Moral-Theistic Bridge can be upheld via univocation. So the only hope left is via Role-Satisfaction. I turn now to assess whether moral arguers can uphold the Moral-Theistic Bridge or the Moral-Natural Gap via role-satisfaction. I address Harrison's account of what the role that needs satisfying is and argue that he offers an implausible analysis of that role. I then address Adams' account of what the role that needs satisfying is and argue that natural properties trivially satisfy it. Thus Adams cannot uphold the Moral-Natural Gap.

§4.3 ROLE-SATISFACTION: HARRISON

According to some moral arguers, the concept MORALLY OBLIGATORY has an analysis that reveals an open role that only a theistic property can satisfy. The analysis reveals that moral obligations are commands or imperatives to do this or that. What is left open is whose commands they are, or more generally what the nature of these commands is. But when we reflect on the special nature of these commands, we can see that they would have to be God's commands.[34] The most sophisticated proponent of this argument is Gerald Harrison. In his *Normative Reasons and Theism* he argues that, in fact, *all* normative reasons, not just moral obligations, are identical to divine favourings. One sort of divine favouring is a divine imperative – a divine command – but other favourings are divine suggestions, divine exhortations, divine pleadings, and so on.

[34] The argument here is not, of course, that all commands of any sort must be divine commands. Thus, this argument does not collapse the moral argument into some sort of argument from the existence of commands.

His argument begins with an analysis of normative reasons. Normative reasons are, he claims, favouring relations (2018, pp.11-27). To have a normative reason to ϕ is to be favoured doing ϕ.[35] This analysis *qua* analysis leaves open whose favourings they are, if any, or what their nature is. But he argues that in any world we can conceive of, only a mind can favour something. And he argues that if we conceive of a world like ours, only a mind external to us humans could do the relevant sort of favouring. This is essentially a divine mind. So normative reasons are the favourings of a divine mind. Thus it seems that both the Moral-Theistic Bridge and the Moral-Natural Gap are upheld. A divine mind is no sort of naturalistically acceptable thing.

It seems to me that there is one major problem with Harrison's argument. It is that his analysis of normative concepts is incomplete, and when it is completed, it turns out to be circular, and so not an analysis of normative concepts at all. Consider what he takes to be the best evidence that normative reasons are to be analysed as favourings:

> But, for me anyway, far and away the most powerful evidence that normative reasons are favourings of doing and believing things is the phenomenological evidence. Talk of normative reasons as commands, requests, pleas, suggestions and so on, does no more than give expression to how things appear.
> I can talk with some authority about how things appear to me. And in addition to there appearing to be trees and mountains and sky and all manner of other physical phenomena, there also appear to me to be favourings of doing and believing things. And it is these curious favourings that I refer to as normative reasons and that I have hitherto assumed others have been referring to as well, for the way they talk about them corresponds to how I do.
> So, speaking personally, when I judge myself to have reason to do or believe something, I am most certainly judging that I am either being commanded, requested, pleaded with, advised, or something similar, to do or believe that thing.

[35] Harrison's analysis of normative reasons differs structurally from typical accounts. Primarily, he denies that a fact or proposition or some objective natural feature could be a reason for an action. Because he takes there to be good evidence that reasons are favourings, he objects to more typical accounts; "[t]he idea of a fact or proposition or objective natural feature literally *favouring* something is just plain incoherent. I mean *how on earth* can a fact *favour* anything (2018, p.37)"?

(2018, p.18)

According to Harrison, when he is aware of, say, some moral obligation bearing on his deliberation, he is aware of it as a command being made to him. And that may very well be how things appear to him, and others. I doubt, however, that that captures *all* of the phenomenology of normative reasons. Let me speak about my experience, which Harrison would grant I am an authority about. It is not just that I feel commanded when I become aware of a moral obligation (if indeed I do feel commanded), it is that I feel commanded by something that is *authoritative* over me. After all, I do not feel like the command is just some loud, annoying voice that happens to be in my head. Its power to move me is not just that it berates me and psychologically coerces me into action. Rather I sense that the command has some *legitimate claim* over me. Same with the way normative reasons might plead with me, request things of me, and so on. These are authoritative pleas and requests.

So if Harrison's phenomenology of normative reasons is anything like mine, the evidence that he appeals to for his analysis of normative concepts in fact supports the analysis that normative reasons are *authoritative favourings.* Now there is a normative concept of authority and a non-normative concept. *A* has non-normative authority over *B* when, just as a matter of fact, *B* defers to *A,* submits to *A,* and so on. But this is clearly not the kind of authority that the phenomenology reveals normative reasons to have. Again, the favourings of normative reason, if we grant that they appear as favourings, do not simply psychologically coerce us into de facto submission to its demands. Rather, the favourings have a legitimate claim on us. It is the *normative* concept of authority that we must appeal to in understanding the phenomenology. And the resulting analysis of normative reasons would be that they are *normatively* authoritative favourings. But now we have had to appeal to a normative concept to provide an analysis of normative concepts. This is circular, and so fails to truly be an elucidation of

normative concepts at all.

It seems that Harrison himself would accept that circularity in analysis entails that one has failed to actually provide an analysis. After all, this is the exact criticism he makes of analysing normative concepts in terms of *oughts*:

> Some suggest that my judgement that I have reason to do/believe something is equivalent to judging that I *ought* to do or believe that thing. … But it is wholly uninformative. This is because the term *ought* is ambiguous, and has normative and non-normative meanings. … But if we have to invoke the concept of a normative reason in order to distinguish this [normative meaning of] ought from others, then we cannot use *ought* or *oughtness* to explicate the concept of a normative reason or normativity. We have gone in a circle and are now helping ourselves to that which we were supposed to be explicating (2018, pp.14-5).

Harrison's analysis, if he goes for the one actually supported by the phenomenological evidence, would be similarly circular. So he has failed to show that normative concepts have an analysis that leaves an open role for theistic properties to satisfy and natural properties to fail to satisfy.

My objection is not to be confused with another objection to do with authority that Harrison responds to. Here is how he frames this other objection:

> Some may worry that if normative reasons are the attitudes a mente divina is adopting towards us doing and believing things, then they will lack any real authority … It can be expressed as a question: "why should I do as this god favours me doing and believing? …" (2018, p.161).

Harrison considers what this objection amounts to if authority is understood to be normative authority. He claims that it simply results from a misunderstanding of his thesis: "… the questioner is asking to be shown that there is normative reason to comply with what the mente

divina is favouring them doing and believing. This question is confused. By hypothesis, normative reasons *just are* this mente divina's biddings (2018, p.161)."

That objection may very well be confused, but it is not my objection. I am not claiming that even if moral obligations are divine commands, we would need some normative reason to obey them. I am challenging the analysis Harrison gives of moral and normative concepts. My objection strikes at Harrison's argument at its first premise – that normative reasons can be analysed as favourings simpliciter – it strikes before he even gets around to identifying those favourings with any divine mind. I conclude, then, that Harrison's argument fails.

One might, however, try and save his argument in a more modest form. Suppose Harrison abandoned analysing normative concepts *in general* as favourings but focused only on moral concepts like the morally obligatory. He could claim that the analysis of MORAL OBLIGATION reveals it to be a normatively authoritative command (I shall omit "normatively" from here on, but the reader should remember that it is the normative concept of authority that I shall have in mind). This is not obviously circular. The concept of moral obligation is not the same concept as that of normativity in general so normativity can be appealed to elucidate moral obligations.

What we need, however, is some reason to analyse moral obligations in this way. Perhaps the reason is that we still need to do justice to the phenomenology of moral obligation. When we become aware of our obligations they feel like commands bearing down on us. This analysis respects that phenomenology. But note that on the proposed analysis, we are supposing that NORMATIVITY is a constituent concept of MORAL OBLIGATION. But unless NORMATIVITY can itself have the right sort of conceptual connections to other concepts, then, according to the Conceptual Constraint, normativity will be irreducible. And so moral obligations will only

come about if there are commands and some distinct, irreducible normative fact. The natural picture that emerges is that these commands are invested with normative authority by the irreducible normative fact.

Let me explain why this ends up undercutting the motivation for this analysis in the first place. If there is some distinct normative fact that invests divine commands with their normative force, then the divine commands themselves are not the best explanation of our appearing to be commanded in our moral phenomenology. Consider a comparison with a human case. As Wolterstorff explains (2008, pp.268-271) when a military officer issues a command to a private, there is some standing normative fact that explains why those commands are obligatory. The normative fact invests the commands with their authority.

Now suppose that war is imminent, and I am a private terrified of being deployed. To my horror, the commanding officer gives the orders to deploy. I deliberate about what I will do in response. I want very much to disobey, drop out of the armed forces, and remain at home. Yet I become aware of my civic duty to remain in my post and be deployed. As I wrestle with what to do, one voice in my head, representing my civic duty, seems to be commanding me to deploy. But this is not the voice of my military superior. I am not simply replaying his orders in my mind. In fact, we can imagine that he gave his orders in a manner quite contrary to the commanding voice presently in my head. He may have given them in the form of a declarative sentence rather than a command, "you will be deployed" rather than "deploy!" and he may have said it quite gently. By contrast, I feel my civic duty to deploy intensely screaming imperatives at me whenever I consider dropping out.

It seems to me that the feeling of being commanded comes from my appreciation of the

background normative fact, rather than the officer's command. In Wolterstorff's language, I am aware of the standing obligation I have to obey the officer's orders. This is the background normative fact that invests the officer's orders with authority. And it is the standing duty that is screaming imperatives at me. Put it this way, it is not so much the officer's command that is weighty in my head, it is my awareness of my civic duty telling me to *obey the officer's command*. In a sense, the civic duty is commanding me to obey my officer's command. The officer's actual command is not the focus of the *being-commanded* phenomenology. This is certainly how the phenomenology seems to me.

Structurally, the same would be true if divine commands invested with normative force by some prior normative fact. If God were to command me to ϕ, what would be salient in my phenomenology would be the sense of being commanded, by the normative fact, to do what God commands. This is what would be different between me deliberating about whether to obey God's commands, and whether to obey the command of a peer. I might disregard my peer's command without any internal voice trying to persuade me otherwise. But in the God case, it is this sense of a second-order command to do what has been commanded that bears down on me. The actual commands of God are not what I am directly aware of when I feel the pressure of this second-order command. But then it seems that there is no motivation to analyse moral obligations as involving the literal commands of any concrete being such as God. Those commands would not explain the phenomenology.

One might object that we are still left with a phenomenology involving a sense of being commanded. How are we to respect that fact without appealing to any actual commands? I do not know. Perhaps the phenomenology cannot be taken seriously as revealing the actual nature of normative reasons. After all, as I have just argued, attempts to take it seriously in that way

run into major problems. They end up with circular analyses or end up failing to explain the phenomenology after all. It is now the burden of the moral arguer to show that these problems can be bypassed. Until they do, we should be cautious about what the phenomenology of moral deliberation reveals.

§5.4 ROLE-SATISFACTION: ADAMS

Next I turn to the arguments of Robert Adams, who most moral arguers defer to. I will focus on his arguments concerning moral obligation rather than moral value simply because everything I say about moral obligation can be applied, with small modifications, to moral value. Unlike Harrison above, Adams does not claim that commands or favourings are part of the analysis of MORALLY OBLIGATORY. Rather, he gives an analysis of moral concepts in terms of a broader functional role and then argues that it is more plausible that divine commands would satisfy that role than anything naturalistically acceptable. So, commands do not enter the argument at the analysis stage, only at the stage of considering what could satisfy the open role provided by that analysis. His analysis is therefore not vulnerable to the charges I levied against Harrison.

Adams first develops his analysis in his paper *Divine Command Metaethics Modified Again* (1979b, pp.73-5), and then develops it further in *Finite and Infinite Goods* (1999, pp.231-248). The functional role that Adams lays out in his analysis of MORALLY OBLIGATORY has several elements. The morally obligatory is that property which is:

(i) Objective

(ii) Instantiated in the majority of actions that we pre-theoretically consider

morally obligatory

(iii) Involved in regulating our epistemic judgments about what is morally obligatory

(iv) Able to rationally motivate individual and communal actions and attitudes

The kind of objectivity Adams is concerned with in (i) is the kind moral realists are interested in, the kind explicated in §2.2. What is involved in (ii) is what Adams calls "extensional adequacy." The thought is that prior to doing our meta-ethics, we hold to various judgments about which actions are morally obligatory (I will often focus on moral wrongness in what follows). According to Adams, whichever property is the reduction base of moral wrongness it cannot turn out that it is distributed over actions completely differently than how we pre-theoretically believe that wrongness is distributed. A candidate reduction base that did that would involve such a revision to our understanding of moral wrongness that it would seem more a change in topic than a successful reduction. Adams does not say much about what the epistemic regulation in (iii) involves except that it may require the property to have some reliable causal influence on how we form judgments that some action is wrong (1999, p.247).

Before we consider (iv), let us take stock of how natural properties could fare in satisfying (i) through (iii). Clearly, a great many natural properties are objective, so there should not be much problem in them satisfying that condition. As for their extensional adequacy, I will address that in the following chapter, as other moral arguers object to natural reductionism on those grounds. I will argue that there is no dialectically effective way that moral arguers can show theistic properties superior to natural properties with respect to extensional adequacy. As for their role in epistemic regulation, it seems to me that if natural properties can satisfy extensional adequacy, then they will satisfy this also. Suppose for instance, that the natural property

maximising happiness is roughly distributed among all and only those actions that we pre-theoretically take to be obligatory. It would then seem highly plausible that maximising happiness has a powerful causal role in how we form judgments about what is obligatory. So, the only condition left that can eek out an advantage for theistic properties over natural properties is (iv), which I will call the *rationality requirement.*

One might wonder, though, whether I have omitted a further condition that Adams includes in his analysis. Adams is known for defending his voluntarism in the context of a broader social account of obligations. This might suggest to some that Adams includes a specifically social criterion within his analysis of obligations. And I agree that he does, but the picture is complicated. There are various features of obligations that Adams considers to be, in some important sense, social. Some of these features he takes to be given by conceptual analysis of MORAL OBLIGATION, but others are not. What conceptual analysis reveals is that:

> If I have an obligation … it can only be in a personal relationship or in a social system of relationships. If an action is wrong, likewise, there must be a person or persons, distinct from the agent, who may appropriately have an adverse reaction to it (1999, p.233).

What Adams seems to mean here is, firstly, that obligations concern our relationships to other agents. By contrast, "the beauty of a scene or the badness of pain can be understood in abstraction from any social setting (1999, p.233)." The point here is that agents have no opportunity to fulfil or violate their obligations absent contact with others. But this is already accounted for in the extensional adequacy criterion. If, pre-theoretically, wrongness is distributed over actions that involve other agents, a theory of wrongness must reflect that. Secondly, Adams is saying that wrongness must be able to motivate certain attitudes on the part of those other agents in the social context. Again, this is already accounted for in the

rational requirement that I unpack below.

Adams does also argue that obligations are *created* by social relationships. Moral obligations are created by the requirements God puts on us in our relationship with him. But no such social criterion is included in the analysis of MORAL OBLIGATION. This can be seen in the distinction Adams makes between the two tasks of spelling out the semantic role of obligation concepts and finding the best candidate for satisfying that role: "… the semantics of obligation leaves metaphysical questions open, regarding the reality and nature of obligation. We can understand the role of obligation and still ask, 'Is there really something that is suited to fill this role? And if so, what is the best candidate?' The best candidate, I will argue, is to be understood, roughly, in terms of social relations and, ultimately, in terms of divine commands… (1999, p.234)." Adams appeals to the creative power of social relationships at the stage of finding the best candidate, not at the stage of conceptual analysis.

There is, however, a connection between the social features of obligations that are part of the analysis of MORAL OBLIGATION and those that are not. Adams holds that social relationships create obligations *because* that makes sense of the rational requirement. It would make sense of why other agents in the social setting can appropriately have adverse reactions to wrongness. If obligations have a social nature, it would make sense of why wrongness has implications for social attitudes and reactions. However, I argue below that it is trivially easy for a large number of properties, even natural properties, to satisfy the rationality requirement. At least, one cannot argue against their satisfying it without objecting to their extensional adequacy. There is, then, no reason to prefer a social account, or any other account friendly to voluntarism, on the

grounds that it best satisfies the rationality requirement. [36] And the extensional adequacy criterion turns out to do most of the work in Adams' argument.

Let us consider, then, what Adams says about the rationality requirement. According to Adams, it is a conceptual truth that morally wrong actions are ones that we rationally ought to take seriously:

> It is part of the semantically indicated role of moral obligation that it is something one should take seriously and care about… Part of taking moral obligation seriously is our response to violations of it - that is, to wrong actions, in the relevant sense of 'wrong'. If an act is morally wrong, then in the absence of sufficient excuse, it is appropriate for the agent to be blamed, by others and by himself (1999, p.235).

For Adams, taking wrong actions seriously involves having certain attitudes in response to them. If an agent performs a wrong action, the agent has reason for self-blame. And onlookers have reason to cast blame at the agent. But Adams is clear that the appropriate response to wrong actions goes far beyond blame. If we are wronged, then we have reason to feel anger. There are also actions that are appropriate in relation to wrong-doing. We have reason to publicly discourage such action, perhaps punish it. We have reason to apologise for any wrong actions we are responsible for. In sum, whatever property is morally wrongness, it stands in rational relationship to attitudes and actions like blame and punishment. We can call these attitudes and actions *reactive stances* to wrongness.

There are, however, several different ways to understand the relationship between the concept, MORAL OBLIGATION, and the reasons we have for reactive stances. And on any of these views,

[36] Adams also has other arguments for preferring a social account but they do not directly concern the features he includes within the analysis of moral obligations. Since I am focusing on the Conceptual Gap Argument, I will leave these other more minor arguments to the side. At any rate, Murphy (2011, pp.124-132) has already given a detailed critique of Adams' reasons for preferring a social account of obligations.

it is difficult to show that a theistic property better satisfies this element of the wrongness role than any natural property. Consider first that on one view, the reasons we have to adopt reactive stances simply reduce to the conceptual connection between moral obligation concepts and reactive stances. Our reasons for reactive stances are nothing over and above these conceptual connections. So long as a property can be represented under the concept MORAL WRONGNESS, that property stands in the desired rational relationships – it does so just by virtue of being represented under MORAL WRONGNESS. If this is the picture we are working with, then any natural property that satisfies the other criteria in the wrongness role will trivially satisfy this criterion.

To see this, note that properties can be represented under different concepts. This is precisely one of the lessons that can be drawn from examples of reduction like water reducing to H_2O. The property *being water* is identical to the property *being* H_2O. But the *concepts* of water and H_2O are different, even though they refer to the same property. It is sometimes said that properties have differing *modes of presentation.* Each concept that a property can be referred to under is a different mode of presentation for that property. If we hypothesise that moral wrongness is identical to some natural property, we hypothesise that one and the same property has two distinct concepts – two modes of presentation. This property can be represented under a natural concept but also under MORAL WRONGNESS. But then by hypothesis, this natural property has the required rational connections to reactive stances. It has them just by being presented under MORAL WRONGNESS and enjoying the conceptual connections of that concept – connections which constitute the required rational connections.

Suppose, then, that we take an alternative view and deny that the rational connections reduce to the conceptual connections. Suppose, instead, that the reasons to adopt reactive stances are

independent, substantive facts. A property can only be presented under the concept MORAL WRONGNESS if there are substantive normative facts linking that property with reactive stances. That's fine as it stands, but to fill out this view, we need to ask whether the normative facts connect to the property *wrongness* under any particular conceptual presentation of wrongness. Consider that when I judge that Superman is strong, I refer to Superman under the concept SUPERMAN. Even though Superman is Clark Kent, that does not entail that I judge that Clark Kent is strong. I refer to Superman through a particular conceptual presentation of him. This is in contrast to, say, causal relationships. A rock does not smash a window through any conceptual presentation of windows. The question, then, is whether the connection that normative facts have to wrongness is more like the above case of reference, or more like causation.

Let us consider what follows if the connection is more like causation. Take Adams' view that wrongness is identical to being proscribed by God. This view could be true even if we were not aware of any normative facts connecting being proscribed by God and reactive stances. That is, it could be true even if, when we conceive of an action being proscribed by God, we do not intuit any normative relevance to reactive stances. So long as being proscribed by God is *in fact* identical to wrongness, then the normative facts *do* link being proscribed by God to the reactive stances. They just do not link being proscribed by God to the reactive stances through the *concept* PROSCRIBED BY GOD. Thus, when we conceive of an action presented under PROSCRIBED BY GOD, we are blind to the normative connections between the *property* of being proscribed by God and the reactive stances. The same would apply to any natural property posited as identical to wrongness. Such an identity could not be objected to on the grounds that when we conceive of an action bearing that natural property, we do not intuit any normative relevance to reactive stances. It would have to be objected to on some other grounds,

like extensional adequacy. So once again, natural properties trivially satisfy the requirement of motivating reactive stances.

Next, we can consider what follows if the connection between the normative facts and wrongness is sensitive to conceptual presentation. It won't help Adams if the normative facts connect to wrongness under the concept WRONGNESS. For once again, any natural properties posited as identical to wrongness will trivially satisfy the requirement. If a natural property is identity to wrongness, then it can be presented under WRONGNESS. As such, it gets rational connections to reactive stances for free. What Adams must require is that the normative facts connect to wrongness under its non-moral conceptual presentation. That is, wrongness is identical to proscribed by God only on the condition that the normative facts connect to wrongness under the concept PROSCRIBED BY GOD. Likewise, wrongness is identical to a natural property only on the condition that the normative facts connect to wrongness under the natural concept. Thus moral arguers can object that natural properties *qua natural properties* do not seem to motivate reactive stances.

But consider those theories in first order normative ethics that posit natural properties as what make actions right or wrong. Utilitarianism, for instance, claims that failing to maximise happiness is what makes actions wrong. These theories have the plausibility they do because there are at least some actions that bear these properties and satisfy the rational requirements. On many occasions, when an action fails to maximise happiness, we have reason to blame the agent, feel angry and so on. Now one might object that utilitarianism, and any other theory that posits natural properties as those that make actions wrong, are implausible. But the primary way to show that is to argue that these theories make predictions that conflict with our pre-theoretical judgments about how rightness and wrongness are distributed over actions. This,

however, is just to argue that the natural properties they focus on are not extensionally adequate. Alternatively, Adams might argue that it is not these natural properties that motivate reactive stances. They merely make actions wrong, and it is the wrongness itself that motivates reactive stances. But it is not clear why one could not say the same thing about theistic properties like proscribed by God.

One might object to natural properties providing normative reason to adopt reactive stances on other grounds, namely, that there is no conceptual connection between the relevant natural concepts and normative concepts. It may be charged that there is no analytic connection between maximizing happiness and having a normative reason to ϕ, and we want analytic connections. But we need some argument as to why conceptual connections are required here. Even if there are conceptual constraints on reduction, it does not follow that there are conceptual constraints on normative reasons. This would be an unmotivated move. It seems, then, that there is no view of normative reasons that allows Adams to show that theistic properties have the relevant advantage.

One move left for Adams is to change his proposal so that he is concerned with motivating reasons rather than normative reasons. On that reading, Adams is claiming that an analysis of MORAL WRONGNESS is *whatever property that, among other things, psychologically motivates us to ϕ when we recognise that the property is present.* This is no help. We are paradigmatically motivated by our intentional mental attitudes. Desires, in particular, are involved in motivation. But the problem is that our desires only refer to some property under a particular concept. When I desire that an action that maximises happiness be performed, the propositional content of that desire [that an action that maximises happiness be performed] only refers to the property of maximising happiness under the concept MAXIMISING HAPPINESS. But a property can be

represented under different concepts.

Suppose, then, that moral obligatoriness is identical to some natural property. It follows that one and the same property has two distinct concepts – two modes of presentation. This property can be represented under a moral concept and under a natural concept. The moral concept is MORAL OBLIGATION and the natural concept is, say, MAXIMISING HAPPINESS. Now one can desire that an action that is morally obligatory be performed and one can desire that an action that maximises happiness be performed. But even though these two desires refer to actions that have the same property, their psychological, motivational effects can be quite different. Audi makes this same point, only he does so while considering reductive voluntarism rather than natural reductionism: "… desires with different conceptual content differ even if their different contents - say to do one's duty toward one's children and to fulfil God's will in this matter - express the same property (2011, p.161)."

Now here is the upshot of all this. Given that we are interpreting Adams' motivational condition psychologically, rather than normatively, the analysis he gives of MORAL WRONGNESS turns out to be *the property, whatever it is, that, among other things, has a mode of presentation under which it motivates opposition.* But *any* property that one posits as being the reduction base of moral wrongness will satisfy this. For any property, natural or theistic, that one posits as being the reduction base of moral wrongness will, by hypothesis, have a moral mode of presentation and a non-moral mode of presentation. And the moral mode of presentation will be a mode under which the property motivates opposition. If I represent an action as being a morally wrong action, I will be motivated to oppose its performance. But since any property can trivially satisfy this condition, there is no room here for theistic properties to have any advantage over natural properties. I conclude, then, that on either the motivational or

psychological reading, Adams' argument fails. Adams does not present an account of the moral wrongness role that sustains the Moral-Natural Gap.

§5.5 A MODEST ARGUMENT

The arguments presented by theistic moral arguers have been shown not to work. But in the previous chapter, I noted that Bedke (2012) – not himself a theistic moral arguer – has a conceptual gap argument against natural reductionism. This argument does not depend on offering any analysis, even in part, of moral or normative concepts. It simply notes that whatever the application conditions of moral concepts are, natural properties do not satisfy those conditions. After all, imagine an idealised cognitive agent able to process an unlimited amount of information, and give this agent an exhaustive list of all the truths about natural facts and properties, and knowledge of the application conditions of moral concepts. Would this agent be able to tell, just from this information, what actions moral concepts apply to? Bedke holds that most of us would answer negatively. The implication of this is that either moral concepts lack open roles or if they have them, they are not satisfied by naturalistic ontology.

This argument could be put to service by moral arguers. Perhaps just as one does not need to give an analysis of moral concepts to see that natural facts do not suffice for the application of those concepts, one does not need an analysis to see that theistic facts *do* suffice for the application of those concepts. After all, as was conceded in the previous chapter, there is much more intuitive resistance to natural reductionism than there is to theistic reductionism. What Laskowski calls the sense of incredibility in ethics is stronger for natural reductionism than for reductive voluntarism. Perhaps, then, the moral arguer does not need to provide a specific account of the conceptual connections between moral properties and theistic properties. That

there is some sort of connection is plain, which is enough to show that theistic reductionism should be preferred to natural reductionism.

But this is not enough. That is because even if moral arguers want to refrain from giving the details of the conceptual connections involved, we can unearth those details for them. And when we do, we see that they are irrelevant to reduction. It seems to me that when people hold that there is a conceptual connection between theistic propositions and moral propositions, this is because they take it to be an analytic truth that God is morally good. In other words, they agree with Craig that God's existence entails the existence of moral value. But as I argued above, any analytic connection between GOD and MORAL VALUE fails to support reductive voluntarism over natural reductionism.

Further, if there *is* a connection between GOD and MORAL VALUE, it is a compound-to-constituent connection. This connection would not allow theistic ontology to pass Bedke's test for reduction. The conceptual test that Bedke endorses relies on asking whether an ideal agent given a set of truths *L,* and the application conditions for moral concepts, would have enough information to know whether those concepts apply. There are, however, rules on what concepts the truths in *L* are allowed to use. They are not allowed to use the moral concepts themselves (Bedke 2012, p.115).[37] Now if we take the *L* truths to be the full set of theistic truths, *L* violates this rule. If GOD is a compound concept that is constituted in part by MORALLY GOOD, then any truth in *L* that uses GOD uses a moral concept. No doubt a lot of truths in *L* use GOD! Therefore, moral arguers cannot appropriate Bedke's argument.

[37] As Bedke puts it, "A set of minimal L truths is a set of propositions (i) formulated without *use* of H concepts, though *mention* of H concepts is allowed so long as (ii) explicit information about what H concepts refer to is excluded, and (iii) otherwise, the L truths are complete and maximally discriminating of the way things are L-wise (2012, p.115)."

Of course, one might wonder *why* the *L* truths are forbidden from using moral concepts. The answer is that Bedke designs his test such that passing it is sufficient for reduction. But if the *L* truths can use the moral concepts themselves, passing the test ceases to suffice for reduction. For instance, if GOD has MORALLY GOOD as a constituent part, then a world where GOD has application is also a world where MORALLY GOOD has application. The theistic truths that use GOD do give enough information to tell an ideal agent that the moral concept applies. But they give no specifics as to what ontology the concept applies in virtue of. They leave open whether God is morally good because moral goodness reduces to God or because, say, God meets some independent Platonic standard of goodness. But if a non-reductive option is left on the table, passing the test does not suffice for reduction.

Of course, moral arguers may want to protest that some other conceptual connection is responsible for the lack of intuitive resistance to reductive voluntarism, not a compound-to-constituent connection between GOD has MORALLY GOOD. But they need to show us what that connection is. This involves getting back into the details of how theistic concepts connect to moral concepts. Until those details are provided, they cannot rely on Bedke's argument.

Overall, then, there is presently no successful version of the Conceptual Gap Argument for preferring reductive voluntarism to natural reductionism. No account of what the conceptual connection *c* amounts to allows moral arguers to satisfy all the desired criteria. What this means is that the major argument for preferring reductive voluntarism to natural reductionism fails. That is, the major argument for P2 fails. There is, however, a second prominent argument for preferring reductive voluntarism to natural reductionism. Although less influential, what I call the *divergence argument* is still utilized by many moral arguers. It concerns the extensional adequacy of natural properties. In the next chapter, I address this argument.

6

The Divergence Argument

Recall that theistic moral arguers aim to motivate P2 [it is plausible that theism accommodates objective moral properties and more plausible that it does so than that naturalism does so]. One way they have tried to do that is by arguing that reductive voluntarism is more plausible than its naturalistic competitors. That requires it having an advantage over some *specific* competitors, like natural reductionism. In the previous chapter, I showed that the major argument for preferring reductive voluntarism to natural reductionism, the Conceptual Gap argument, fails. But there is another prominent argument for preferring reductive voluntarism, what I call the Divergence Argument.

In this chapter I argue that the Divergence Argument cannot be motivated in a way that is dialectically effective against atheists. For motivating it requires arguing against all naturalistically acceptable first-order normative theories. This significantly reduces the dialectical value of the argument. But more importantly, moral arguers need to adopt a naturalistic normative theory to defend voluntarism against an important atheist objection. They need to defend against what I call the Counterfactual Argument. I argue that without appealing to a naturalistic normative theory, voluntarists will not be able to give atheists a reason to reject the Counterfactual Argument.

§6.1 INTERPRETING THE ARGUMENT

Here is Moreland and Craig expressing the divergence argument, which they take to be the second major objection to natural reductionism:

> Second, every attempted reduction of a moral property to a natural one has failed because there are cases where an act is right even if it does not have the natural property, and an act can have the natural property and not be right. For example, suppose one reduces the moral property of rightness in "*X* is right" to "*X* is what is approved by most people." This reduction is inadequate. For one thing, the majority can be wrong. What most people approve of can be morally wrong. If most people approved of torturing babies, then according to this version of ethical naturalism, this act would be right. But even though it was approved by most people, it would still be wrong. On the other hand, some acts can be right even if they are not approved of (or even thought of, for that matter) by most people (2003, p.401).

The argument includes an implicit constraint on reduction:

> **Tracking Constraint:** a moral property *m* reduces to a property *f* only if necessarily all and only actions with *m* also have *f.*

This is the demand that the moral property supervenes on the non-moral property. There can be no divergence between the moral property and the non-moral property. The argument also includes the claim that natural properties violate that constraint:

> **Moral-Natural Divergence:** it is not the case that necessarily all and only actions with *m* also have some natural property *f.*

Or at least, the argument claims that no natural property *f* that has so far been proposed as a candidate reduction base of the moral property has satisfied the tracking constraint. But this detail can be ignored for the sake of simplicity, doing so will not affect my response.

The way that Moreland & Craig defend Moral-Natural Divergence is to argue that hypothesising that *m* tracks some natural property *f* leads to counter-intuitive implications. We

have pre-theoretical intuitions about which actions count as morally wrong, or morally good etc. If hypothesising that *m* tracks *f* goes against these intuitions, Moreland and Craig think that provides evidence that *m* and *f* diverge. This seems to be the reasoning in their toy example where obligatoriness is hypothesised to reduce to whatever most people approve of. This hypothesis leads to the conclusion that if most people approved of torturing babies, then doing so would be morally obligatory. But Moreland & Craig clearly expect us to have strong intuitions against the claim that torturing babies is morally obligatory in that scenario – intuitions that provide evidence that obligatoriness and majority approval diverge.

Clearly, the Tracking Constraint and Moral-Natural Divergence entail the conclusion that:

> It is not the case that *m* reduces to some natural property *f.*

There is no problem of logical form in the argument. Of course, in order for this argument to not be turned against theistic reductionism, Moreland & Craig, and other moral arguers that follow them, will need to endorse:

> **Moral-Theistic Tracking:** it *is* the case that all and only actions with *m* also have some theistic property *f.*

And Moreland & Craig claim that where *m* is moral wrongness, the tracking theistic property is *proscribed by God.* And where *m* is moral goodness, Moreland & Craig claim that the tracking theistic property is *resembling God.* This is the meta-ethical, reductive voluntarism that we have been concerned with in the last couple chapters. For simplicity, however, I shall

focus on moral rightness and wrongness; everything I say could be given minor tweaks to apply to moral goodness and badness.

We caught the first glimpse of something like the divergence argument in the previous chapter when I discussed Adams' role-satisfaction argument for preferring reductive voluntarism to natural reductionism. According to Adams, an analysis of MORAL WRONGNESS reveals that whichever property is the reduction base of moral wrongness, it cannot turn out that it is distributed over actions completely differently than how we pre-theoretically believe that wrongness is distributed. A candidate reduction base that did that would involve such a revision to our understanding of moral obligation that it would seem more a change in topic than a successful reduction. What Adams called the criterion of *extensional adequacy* is a kind of tracking constraint on reduction.

Adams' argument, however, can be given a strong reading and a modest reading. On the strong reading, he claims that natural properties completely fail to meet the criterion of extensional adequacy. This is a stronger claim that what Craig & Moreland defend. If natural properties failed to satisfy extensional adequacy, then any natural property would be distributed in such a way that it would not only diverge from how it seems to us, pre-theoretically, that obligatoriness is distributed, but diverge *sharply*. And the divergence would be so sharp that anyone aware of that divergence who still maintained that obligatoriness reduces to the natural property would not have a proper grasp of the concept of moral obligatoriness. Craig & Moreland do not go that far. While they claim that positing a natural property as what moral obligatoriness tracks leads to counter-intuitive implications, they do not claim that the implications are so counter-intuitive as to involve switching the subject.

One might be sceptical that this stronger argumentative burden could be met. Someone who is a meta-ethical utilitarian holds that moral obligatoriness reduces to maximising happiness, a natural property. This has the same implications for which actions are morally obligatory that ordinary, normative utilitarianism has. One might think that these implications are so counter-intuitive as to render both of these utilitarianisms implausible but not *so* implausible that they are unrecognisable as moral theories. Thorough-going utilitarians who embrace its counter-intuitive implications do not seem to be simply talking about something other than the morally obligatory. As Wielenberg says:

> Consider, for example, the disagreement between a Kantian who maintains that what makes morally wrong actions wrong is that they involve treating a rational being as a mere means and a Utilitarian who maintains that what makes morally wrong actions wrong is that they fail to maximize utility. It is too cruel (and, more relevantly, not plausible) to suggest that at least one of these characters is so confused that her position is conceptually incoherent – that one of them doesn't even grasp the concept of moral wrongness (2014, p.19).

But even if one has greater hope for the stronger reading of Adams' argument than that, it clearly only succeeds if the modest one does. And I argue in this chapter that the modest one does not succeed.

On the modest reading of Adams' argument, it is identical to Moreland & Craig's. The claim is not that natural properties completely fail to satisfy extensional adequacy. Rather, they do not satisfy it as well as theistic properties. Positing natural properties as what moral obligatoriness tracks has counter-intuitive implications, even if not *radically* counter-intuitive implications. And positing theistic properties as what moral obligatoriness tracks has fewer counter-intuitive implications. Let us, then, consider Adams as joining Moreland & Craig in their commitment to the divergence argument laid out above. Now as I argued in the previous chapter, Adams' entire role-satisfaction argument for reductive voluntarism turns out to depend

on issues of extensional adequacy. So, in arguing that the Divergence Argument fails, I will also be completing my case against Adams' overall argument.

One problem with the divergence argument as stated above is that its conclusion is not incompatible with natural reductionism. The argument concludes that it is not the case that a moral property *m* reduces to some natural property *f.* But a natural reductionist need not claim that *m* reduces to some *single* natural property *f.* The problem is that the argument launches from a Tracking Constraint on reduction that only covers cases of reduction via *identity.* But as I argued in §2.3, relationships of *constitution* and *realisation* also count as reductive relationships. Accounts that posit these relationships need not claim that the moral property reduces to just one natural property. Constitution accounts can claim that two or more properties are jointly the reduction base of the moral property. And realisation accounts claim that one or more properties are the reduction base in one situation, and one or more properties are the reduction base in another situation.

This problem, however, can be fixed. The tracking constraint can be modified to cover these other reductive accounts, as follows:

> **Tracking Constraint*:** a moral property *m* reduces to property *f* (or properties *fs*) only if there is some relevant set of properties *s* such that necessarily all and only actions with *m* instantiate all the properties in *s.*

Let me explain how this covers all three reductive relationships. An identity account will posit one single property *f* that is the reduction base of *m.* And the set *s* need only include *f* itself. This captures the constraint on identity that was present in Moreland & Craig's original

Tracking Constraint. A constitution account may posit two or more properties *fs* that are jointly the reduction base of *m.* And, again, the set *s* need only include the *fs* themselves.

A realisation account will posit either one property that is the reduction base, or multiple properties. But with realisation, no reduction base property need appear in the set *s.* After all, *s* requires that its members be present whenever *m* is, but realisation allows that different instances of *m* reduce to different properties. So not all of those reduction base properties will be present in every instance of *m.* But realisation *does* require that every reduction base property fulfils the same functional role. That is, all the reduction base properties must share the same *contextual* features. For instance, in philosophy of mind, a realisation account of pain might claim that pain reduces to c-fibres firing in one case, and to a silicon chip state in another. But what makes them both instances of pain is that they both fulfil the pain role – something like being a reaction to bodily damage, the cause of wincing, and so on. So realisation accounts are committed to there being some properties that occur in all and only cases of *m* – they are committed to there being some properties in set *s* – just not the reduction base properties themselves.

Therefore, all reductive accounts are committed to some form of supervenience claim. And as I said in §3.1, a claim of supervenience is just the claim that something tracks something else. Hence there is indeed a tracking constraint on reduction. With the modified Tracking Constraint in place, the rest of the divergence argument can be modified as follows:

> **Moral-Natural Divergence*:** it is not the case that there is some relevant set of properties *s,* featuring only natural properties, such that necessarily all and only actions with *m* instantiate all the properties in *s.* Therefore,

It is not the case that a moral property *m* reduces to any natural property or properties.

Moral-Theistic Tracking*: it *is* the case that there is some relevant set of properties *s,* featuring only theistic properties, such that necessarily all and only actions with *m* instantiate all the properties in *s.*

This gets the desired conclusion, one that is the denial of natural reductionism. So with the divergence argument patched up, it is time to consider the problems with it that cannot be so easily fixed. (Note that from here on, I will drop the * when using the names of the modified premises).

§6.2 NORMATIVE COMMITMENTS

In this section, I argue that although the divergence argument attempts to establish a meta-ethical conclusion (that reductive voluntarism is to be preferred over natural reductionism) it takes on substantial commitments in first-order normative ethics. The result is that these normative commitments significantly reduce the dialectical value of the Divergence Argument.

Recall from §2.1 that normative ethics attempts to figure out which properties *make* actions right or wrong. We can call any such properties *moral-makers.* More specifically, we can call properties that make an action right, *right-makers* and those that make an action wrong, *wrong-makers.* But normative ethics actually goes further than identifying moral-makers. It attempts to identify the most *fundamental* moral-makers. For sometimes a property makes an action

morally wrong only by virtue of being an instance of, or bringing about an instance of, some other property. For example, perhaps being an instance of stealing makes an action morally wrong but only because stealing is an instance of failing to maximise happiness. In which case, failing to maximise happiness is a more fundamental wrong-maker than being an instance of stealing.

The important thing to note is that normative theories have supervenience commitments like meta-ethical, reductive theories do. They claim any action that bears a moral property will also bear one of the fundamental moral-maker properties. This is what makes normative theories projectible. When all the facts about the distribution of the moral-makers are known, we then know how the moral properties themselves are distributed across actions. But it is not just that both normative theories and reductive theories commit to some supervenience claim or other. I will argue that for any reductive theory, one could construct a counterpart normative theory with the *same* supervenience claim. And vice versa; for any normative theory, one could construct a counterpart reductive theory with the same supervenience claim.

This is not an argument that someone who endorses a normative theory is *committed* to its reductive counterpart, or vice versa. That would depend on the account we give of how the making relationship relates to the identity relationship. The point here is just that one kind of theory is constructible from the other kind, whether or not one *entails* the other. The upshot is that to object to the supervenience claim of every version of natural reductionism is to object to the supervenience claim of every naturalistically acceptable normative theory. (A naturalistic normative theory should be understood as one that posits natural properties as fundamental moral-makers.) The Divergence Argument therefore commits one to the view that every naturalistically acceptable normative theory is false.

It is easiest to see how counterpart theories can be constructed when we examine those reductive theories that claim there is a single reduction base for the moral property, and those normative theories that claim there is a single fundamental moral maker. Take normative utilitarianism, which claims that maximising happiness is the single fundamental right-maker. There is a counterpart reductive theory, which claims that maximising happiness is *identical* to moral rightness. Likewise, take normative divine command, which claims that obeying God's commands is the fundamental right-maker. There is a counterpart reductive theory that claims that obeying God's commands is *identical* to moral rightness. Both the normative theory and the reductive theory claim that moral rightness supervenes on maximising happiness, or on obeying God's commands. It is just that one theory claims that the relevant property *makes* an action morally right, and the other claims that the relevant property is *identical* to moral rightness.

Clearly, then, running the Divergence Argument commits one to holding that normative utilitarianism is false. It commits one to the view that maximising happiness is not our single, most fundamental moral obligation. But it also commits one to the view that *any* normative theory that posits a natural property as the single, fundamental moral-maker is false. So, it also commits one to saying that acting in accord with a Kantian categorical imperative is not the single fundamental moral-maker. Indeed, some of the most defended normative theories are those that precisely posit a single, fundamental, natural moral-maker. For the Divergence Argument to claim that all of these theories are false is for it to take on some very substantial normative commitments.

What about normative theories that do not posit a *single* fundamental moral-maker, but multiple? Suppose that a normative theory claims that for every obligatory action, the same *two* properties jointly make that action obligatory. Perhaps, for instance, it is the combination of an action's causing harm *and* being against one's telos that makes an action wrong. Well, there is a counterpart reductive theory that claims that these two properties jointly *constitute* moral rightness. And both the normative and the reductive theory would claim that moral wrongness supervenes on these two properties.

The most complicated case is where a normative theory claims that every obligatory action is made obligatory by one or another property in a disjunctive set of properties. According to such a theory it is a disjunction of properties that are the fundamental moral makers.[38] Again, there is a counterpart reductive theory. It claims that moral obligatoriness is *realised* by different properties in different contexts. And both the normative and the reductive theory would claim that moral wrongness supervenes on this same disjunction of properties. The complication is that this is not enough to show that the normative theory and the reductive theory share *all* of the same supervenience commitments. For the reductive theory seems to have one further supervenience commitment. While realisation allows a disjunctive range of properties to be reduction bases, it still requires that these properties have contextual features in common. It requires that moral obligatoriness supervene on these contextual features. But it would seem that a normative theory is not committed to there being any shared contextual features.

We can handle this problem by noting that any disjunctive normative theory with some plausibility will posit moral-makers that *do* have contextual features in common. Indeed, these

[38] Such a theory seems to have been defended by Ross (1930), who denied that there was one, *single* fundamental duty.

will be features that Adams' own analysis of MORAL OBLIGATION included. According to Adams, whatever plays the role of the morally right is objective and regulates our judgments about what is obligatory. But any plausible normative theory will posit objective properties as the moral-makers (remember that we are assuming the truth of moral realism). And so long as the moral property supervenes on the disjunction of properties, it is very plausible that those properties will regulate our judgment about which actions are morally obligatory. After all, if the supervenience claim is plausible, that's because we do generally judge that morally obligatory actions have one of the properties in the disjunctive range. And a simple explanation of that would be that those properties regulate our judgments in some way.

Of course, Adams also included other features in his analysis. But I showed in the previous chapter that those features either are trivially satisfied by any properties posited as the reduction base. So, to rule out realisation accounts constructed as counterparts to disjunctive normative theories, moral arguers will need to give us an amended analysis of MORAL OBLIGATION. Until they do that, we are within our rights in treating all plausible disjunctive normative theories that are naturalistically acceptable as having counterpart reductive theories.

There do not seem to be any other options for how a normative theory could be structured. And so, it seems that to attack the supervenience claim of every natural reductionist theory is to attack the supervenience claim of every naturalistically acceptable normative theory. The Moral-Natural Divergence premise turns out to carry huge normative commitments. This seriously diminishes the dialectical value of using the Divergence Argument to support the theistic moral argument. The theistic moral argument this thesis is concerned with tries to motivate meta-ethical conclusions. If it carries first-order normative commitments also, moral arguers will either need to restrict the audience of their argument to those who already reject

naturalistic normative ethics or will need to shoulder the burden of arguing against all naturalistic normative ethics.

If they take the first option, the audience for the argument will almost vanish. Recall from §2.3 that moral arguers tend to take theism and naturalism as the only live options for philosophers today. So insofar as moral arguers are dialoguing with philosophers who hold live options, their audience is restricted to naturalists. And they are dialoguing with those naturalists who accept moral realism. But there are surely very, very few moral realist naturalists who do not think that there is any acceptable normative theory compatible with naturalism. To be this kind of naturalist, you would have to think that there are real, objective moral norms that are such that no actual agent's actions are relevant to those norms. Perhaps this naturalist thinks that one fundamental moral norm is to obey God, and that there is no God. So there is a moral norm that gets no purchase on us. While being this sort of naturalist is theoretically possible, it would be surprising if there actually are any such naturalists.[39] Going the route of restricting the audience of the moral argument to such naturalists renders the argument irrelevant.

The second option is therefore *prima facie* more appealing. Indeed, one might think that moral arguers would be very happy to shoulder the burden of arguing against all naturalistic normative theories. For one might think that moral arguers *already* take themselves to be shouldering that burden. That is, they already see normative theories like utilitarianism and Kantianism to be competitors to their divine command theories that need to be defeated. But this is not correct. Most moral arguers seem, at least implicitly, to set up the metaphysical moral

[39] This is not to say that no naturalists think that if there were a God, then we would have a moral obligation to obey God. It is just that they likely think this obligation derivative of a broader obligation to obey those with a certain relationship to us – like being our benefactor.

argument for theism to be independent of any controversial normative premises. It is set up to take a position in a meta-ethical debate, not a first-order normative one.

Consider that the moral argument is explicitly said to be a *meta-ethical* argument that utilises meta-ethical divine command theories. Further, the competitors to these divine command theories that moral arguers criticise are meta-ethical theories, not normative theories. Consider the naturalistic competitors that C Stephen Evans addresses in his book *God and Moral Obligation.* He addresses error theory, expressivism, constructivism, natural reductionism and natural non-reductionism (2013, pp.118-154). And consider the naturalistic competitors that Baggett and Walls attack (at least with respect to moral obligation): evolutionary accounts of morality, natural reductionism and natural non-reductionism (2016, pp.145-178).[40] Neither consider, say, ordinary utilitarianism to be a competitor to meta-ethical divine command theory. Implicitly, moral arguers concern themselves with what moral obligations *are* (divine commands) and remain neutral on what our obligations *tell us to do.*

Of course, one might think that moral arguers are simply mistaken about the normative commitments of reductive voluntarism. They may believe that reductive voluntarism is neutral on first-order theories like utilitarianism, but it is not. Thus, if moral arguers changed tact and starting to argue against naturalistic normative theories, they will simply be doing what they were already (unknowingly) committed to doing.

This line of thought, however, is mistaken; normative theories concern moral-makers, and there are two views of the making relationship that allow reductive voluntarism to be compatible

[40] Both Evans and Baggett & Walls use the more standard names for natural reductionism and natural non-reductionism: 'ethical natural' and 'ethical non-naturalism'.

with naturalistic normative theories. I will now spend some time unpacking these views. This will vindicate my claim that moral arguers would be taking on *new* commitments in arguing against naturalistic normative ethics. But also, there are problems moral arguers face should they choose to adopt these new commitments. These problems only come into view when the options for understanding the making relationship are laid out.

§6.3 THE MAKING RELATIONSHIP

When I discuss the making relationship in this section, I will be discussing it in the context of fundamental moral-makers. These are the kind of moral-makers that normative ethics is principally concerned with. There are four main accounts of the (fundamental) making relationship: making as reduction, making as causation, making as what motivates God to command, and making as what brings some action under the content of a moral law. I argue that the latter two allow reductive voluntarism to be combined with a naturalistic normative theory. Thus, so long as reductive voluntarists endorse one of these accounts, they are not committed to rejecting naturalistic normative theories.

In fact, I argue that these two accounts *commit* voluntarists to a naturalistic normative theory. On the other hand, the other two accounts commit voluntarists to theistic properties being the only fundamental moral-makers. It turns out, then, that voluntarists are either committed to a wholly theistic normative theory or a wholly naturalistic normative theory. There is no room for a blended theory that posits theistic properties as the fundamental moral-makers for some actions, and natural properties as the fundamental moral-makers for others. At least, there is no room for a blended theory in accounting for all those actions typically recognised as wrong by atheists. This will become important later when I argue that even if moral arguers want to

take on the commitment of defeating all naturalistic normative theories, they still run into problems.

Let us start, then, by considering the view that the making relationship is a reductive relationship.[41] It is not hard to see why this view precludes the compatibility of reductive voluntarism and naturalistic normative ethics. On this view, a property that makes an action wrong is one that moral wrongness itself reduces to. And, of course, reductive voluntarism claims that wrongness reduces to (and only to) *proscribed by God.* The reductive view of making therefore implies that *proscribed by God* is the only fundamental wrong-maker. This completely excludes natural properties from being fundamental moral-makers, and thus excludes naturalistic normative ethics.

Next, consider the causal view, which can be attributed to Parfit. He holds that the making relationship actually excludes reduction: "when some property of an act makes this act right, this relation holds between two quite different properties (2013, p.301)." The moral-maker must be ontologically distinct from the moral property itself. Parfit seems to have in mind what I argued in §3.2; that robust explanation between *x* and *y* precludes a reductive relationship. A property cannot explain itself. If Parfit understands making as something robustly explanatory like causation, that makes sense of why he takes it to preclude reduction. So if moral arguers adopt this account, natural properties are only fundamental moral-makers if they do something like cause God to command what he commands. But no natural property causes that. They do not have coercive causal power over God's will. Hence, the causal view also excludes naturalistic normative theories.

[41] It might be thought odd that a reductive relationship could be described as making. But as Copp (2011, p.45) points out, there is a reductive use of "because", e.g. Bill knows that *p* because *p* is true and Bill justifiably believes that *p.* It is no more cumbersome to say "what *makes* Bill know that *p* is that *p* is true and Bill justifiably believes that *p.*"

Let us now consider the two views that, I argue, do *not* rule out the combination of reductive voluntarism with naturalistic ethics. Natural properties may not be able to cause God to command what he does but they can feature in the *reasons* God commands what he does. Perhaps God has desires to promote certain natural properties and issuing commands achieves that. Or perhaps the goodness of certain natural properties gives him a normative reason to issue commands that promote them. Audi proposes an account where the making relationship is precisely this conferral of reasons (although he uses the language of grounding or the *in the virtue of* relationship rather than making). It is an account that he explicitly formulates to combine reductive divine command theory with natural properties as fundamental moral-makers:

> Suppose first that, with some versions of divine command theory, we take the property of being obligatory to be the same property as that of being divinely commanded. Instead of stopping there, however, we might take '*both'* properties (i.e. the property expressed by the theological phrase 'divinely commanded' and the property expressed by the non-theological phrase 'being obligatory') to be grounded - thus consequential, in a strong sense - on non-moral, "natural" properties belonging to the type of obligatory act in question. For instance, acts of loyalty to one's family or religious community might be both obligatory in virtue of one's special relation to them - involving, say, one's promise of giving them support - and *also* commanded by God for that very range of reasons. The divine commandedness of an act, which on this view is the same property as its obligatoriness, is thus in a sense *embedded* in its non-moral grounds. The act is divinely commanded on the same grounds in virtue of which it is obligatory (2011, p.140-1).

According to Audi, God commanding for certain reasons is just the same as something being obligatory in virtue of certain properties. Since natural properties can feature in God's reasons, they can be fundamental moral-makers.

In fact, on this view of making, for all those actions that atheists will recognise as wrong, natural properties will plausibly be fundamental wrong-makers. For any wrong actions like rape or torture, there seem to be natural properties that gives God reason for proscribing those actions. If there are normative reasons to promote things like happiness, love, truth, beauty and so on, then those are reasons for God to proscribe rape, torture, and so on. Perhaps those normative reasons are just God's reasons to promote his preferences. As moral arguers maintain, God "is by nature loving, generous, just, faithful, kind, and so forth… (Kurtz & Craig, 2009)." These are character attributes which plausibly entail that God has preferences for things like happiness and love etc. So as long as God has reasons to promote his preferences, he has reasons to proscribe the sorts of actions atheists recognise as wrong.

I have restricted the above to actions atheists recognise as wrong because it is not impossible that God command us to do things for reasons that have nothing to do with natural properties. Perhaps there are some things God commands only for reasons to do with theistic properties. God might command something for, say, *his glory*. But if any actions that we consider to be obligatory are so on the basis of theistic properties only, they are likely to be those only believed by theists to be wrong. Many theistic religions feature obligations that strike most atheists as having no point on a naturalistic view of the world. These are perhaps obligations to offer worship or sacrifice. If God has commanded anything for only divine-centric reasons, these are the likely candidates.

Enoch provides material for another account of making that allows compatibility between reductive voluntarism and naturalistic normative ethics. Enoch draws attention to the explanatory role that the *content* of a norm has, using the following example:

> If someone arrives at our jurisdiction, notices that people have drinking-status properties, that they also have age properties, and that the former supervene on the latter, and proceeds to ask for an explanation of this (perhaps surprising) fact, all we need to do is direct her attention to *the relevant legal norm.* We can tell her things like "Oh, the law here says that so-and-so", or "you must understand, the relevant legal norms here view gender as irrelevant to drinking-status, and determine only the age as a relevant factor", and so on. What explains the supervenience of legal drinking-status on age is simply the content of the relevant legal norms (2011, pp.143-4).

The example concerns the supervenience of some normative status – in this case legal drinking-status – on a natural property, age. The normative status supervenes on the natural property because of a legal norm. The legal norm says something like "if and only if a citizen is eighteen years or older can they legally purchase alcohol." The natural property referred to by the norm can be thought of as a triggering condition for the normative status. When a person bears the property *being eighteen or older* that person meets the condition specified in the norm for the legal-drinking status to be applied. The natural property *makes* a normative status apply in the sense that it brings the person (or in other cases, the action) under the scope of the norm.

On reductive voluntarism there are moral norms with content. These moral norms are just God's commands. So suppose God issues the command "maximise happiness!" This is a norm that applies the status *morally obligatory* to any act that satisfies the triggering condition of maximising happiness. Maximising happiness can make an action obligatory in the sense that it brings an act under the scope of the moral norm. Indeed, so long as God issues a command to perform an action with a certain natural property, whatever that property is, it can be a fundamental moral-maker. This view allows reductive voluntarism to be combined with naturalistic normative ethics.

Wielenberg, however, objects to interpreting the making relationship in this way. Picking up on Enoch's example, he claims that

> the fact that (a) drinking-status properties supervene (within the jurisdiction) upon age-status properties in a particular way, and the fact that (b) the relevant laws within the jurisdiction say that one can purchase alcohol if and only if one is over a certain age, seem to be the same fact – or at least, these facts are too closely related for one to constitute an adequate explanation of the other (2014, p.22).

But this is a puzzling objection. The fact that property *x* supervenes on property *y* is not the same fact as there being a *law* that connects *x* and *y.* The first fact allows that the supervenience of *x* on *y* is purely coincidental. The latter suggests that the supervenience is not merely coincidental. There is some law that *governs x* such that it supervenes on *y.* Of course, Wielenberg might hold to some view in which law-talk reduces to supervenience talk. But we can simply stipulate that we are talking about law in a stronger, governing sense. We need not be offering an analysis of ordinary talk about laws. Rather, we are positing a theory that attempts to explain what *making* amounts to.

Of course, if we posit laws, there is a further question as to why those laws exist. And in cases of human legal systems, that further question is the more pressing one. It is obvious that the property *illegal* supervenes on natural properties because of a system of law. In most ordinary contexts, if we asked, "why is *x* illegal?" we would indeed not be satisfied by an answer merely citing the existence of a law according to which *x* is illegal. We would want to know why that law exists in the first place. But just because citing the law would not satisfy the question we are asking in that context does not mean that it would not satisfy another question asked in a different context. After all, it is false to say that the existence of the law plays no explanatory role at all. Again, appealing to a law shows that the supervene of *illegal* on natural properties is not wholly coincidental.

Further, in the moral case, it is *not* obvious that moral properties supervene on natural properties because of some *moral law* – some thing with content. And further still, reductive divine command theorists can answer the question about where any such laws come from. Any such laws are divine commands that are caused by God's willing to issue them. So there is no problem in theists appealing to this account of the making relationship. Note also that I did not appeal to this view of making to provide an ultimate explanation of moral supervenience. Rather, only as an account of the making relationship that normative ethics is concerned with. One need not expect the making relationship to get to explanatory rock-bottom. Suppose the normative utilitarian shows us that maximising utility is what ultimately has moral significance. This could be because the only moral law is to maximise utility. We would not expect the utilitarian *qua* utilitarian to explain why there are laws in the first place.

Leaving objections to this account of making aside, I will now argue that on this account, for all those actions that atheists will recognise as wrong, natural properties will plausibly be fundamental wrong-makers. This might be surprising. God can clearly issue commands that refer to the theistic property of some action. And even though atheists will not recognise the theistic property as the wrong-maker, there may be natural properties that are co-extensive with the theistic property. These natural properties will be what atheists (mistakenly) appeal to as the wrong-makers. As an example, God could command "do not harm creatures that bear the divine image!" If all those creatures that bear the divine image are all those creatures that, say, have the capacity for reason, then there could be a theistic wrong-maker for a class of actions recognised as wrong by atheists.

Things, however, are not as simple as this. It is widely recognised that divine command theories like Adams' are subject to an epistemic constraint. That is, God only places an obligation on a

person if God ensures that his command meets some threshold of epistemic accessibility for that person. According to Adams, God must communicate his commands through some *sign,* i.e. through moral intuitions, or through social requirements, or through injunctions given by some prophet or holy text etc. And "the sign must be such that the intended audience could understand it as conveying the intended command (1999, p.256)." Now if God commands something like "do not harm image-bearers!" and this command applies to everyone, then God must have communicated that to everyone. But he does not seem to have communicated that to reasonable non-believers. They are not aware of any obligation to avoid harming divine image-bearers; they do not even believe any such beings exist. The same could be said for any command that references theistic properties.

This is a variant of the Promulgation Objection to voluntarism. The standard version objects that *none* of God's commands meet the epistemic requirement with respect to reasonable non-believers (Morriston, 2009; Wielenberg, 2014; Danaher, 2019).[42] Thus voluntarism is said to implausibly imply that these people have no moral obligations. By contrast, the version I raised above does not deny that reasonable non-believers have obligations given by God's commands. It merely denies that any of these commands refer to the theistic properties of actions. Thus, there is no command like "do not harm image-bearers!" behind the widely recognised obligation not to harm others. At best, there are only commands that refer to theistic properties behind obligations that are less universally recognised. Perhaps only those recognised by theists.

Of course, voluntarists have offered responses to promulgation objections. Typically, they argue that the epistemic requirement is softer than objectors suppose. It is often claimed, for

[42] Wielenberg (2018a) also has a version that focuses on psychopaths rather than reasonable non-believers.

instance, that God does not need to communicate the divine source of the command. God just needs to communicate that the command has authority of some sort (Peoples, 2011; Evans, 2013). Thus the requirement is said to be soft enough to let reasonable non-believers through. Critics have claimed in response that voluntarists need to go further and claim that the command need not be communicated in the *form* of a command (Morriston, 2009; Danaher, 2019). For it is not just that atheists do not believe in God, they do not believe that their moral obligations are the commands of any agent at all. Note, however, that voluntarists will have to soften the requirement even further to disarm my objection. They will say that God does not even have to communicate the *content* of the command correctly. Thus, moral arguers will potentially be in a position where they claim that neither the source, nor the form, nor the content of the command need be communicated.

But the whole point of the epistemic requirement is to retain an analogy between divine commands and the way human commands work. The military officer who does not communicate his command to his private cannot be said to place an obligation on the private. The more the epistemic requirement is softened in the divine case, the more strained this analogy becomes. It is hard to imagine that a military officer could generate an obligation on his private to do *x* if his command is to do *y.* Surely, if anything counts as the minimum that must be communicated, it is the content of the command. If voluntarists soften the epistemic requirement so much that the content is inessential, it seems to me that the analogy between the divine and human cases is so stretched as to be broken. The upshot is that on Enoch's view of the making relationship, voluntarism will be committed to natural properties being the fundamental moral-makers of all those actions atheists recognise as wrong.

Here, then, is where we have got to. Moral arguers want to motivate the Moral-Natural Divergence premise in the Divergence Argument. To do that, they need to defeat all naturalistic normative theories. This requires taking on substantial commitments in first-order normative ethics. Most moral arguers would consider these to be *new* commitments. They are not commitments they already had just by virtue of endorsing reductive voluntarism. I have argued that this is correct. Normative ethics is concerned with the moral-making relationship. And two accounts of the that relationship allow reductive voluntarism to be combined with naturalistic normative theories. There is, therefore, nothing inherent in voluntarism that requires taking a stand against naturalistic normative theories. It is only when moral arguers employ the Divergence Argument that rejecting those theories becomes necessary.

Taking on these new normative commitments would significantly increase the argumentative burden that moral arguers must carry. But some moral argues may be willing to shoulder that burden. However, I argue in the next section that if they take on that burden, they will be left without a way to defend the Moral-Theistic Tracking claim from a powerful objection. This is because in order to reject all naturalistic normative theories, moral arguers must reject two views of the making relationship. As argued above, understanding making in terms of God's reasons or the content of moral laws commits voluntarism to natural properties as fundamental moral-makers. Rejecting natural properties as fundamental moral-makers therefore requires rejecting these two views of making. But these are the only two views of making that give moral arguers the resources to respond to the Counterfactual Argument against voluntarism. I unpack this argument below.

§6.4 THE COUNTERFACTUAL ARGUMENT

Recall that the tracking constraint on reduction holds that moral properties must track the relevant properties *of necessity.* The tracking must hold *across worlds.* Consider again how Moreland & Craig dismiss their toy example of natural reductionism, where obligatoriness is said to reduce to whatever most people approve of. They say, "if most people approved of torturing babies, then according to this version of ethical naturalism, this act would be right. But even though it was approved by most people, it would still be wrong (2003, p.401)." Thankfully, in the actual world, most people do not approve of torturing babies. It is only in a merely possible world that most people are so depraved. What Moreland and Craig are saying is that this 'majority-approval' reductive view is committed to certain counterfactuals. One of them is that in a world where most people approve of torturing babies, it would be right to torture them. Their objection is that this counterfactual is false. We have a strong intuition that torturing babies would be wrong in that world.

The difficulty for reductive voluntarism is that it can be objected to on the same grounds. Voluntarism seems to predict, for instance, that in a world where there is no God, and someone is raped, that act of rape is not wrong. But many people have the intuition that, even in an atheistic world, rape *is* wrong. This is a challenge to the claim of Moral-Theistic Tracking. Critics of theistic moral arguments have objected to voluntarism on this basis. Admittedly, critics tend to consider atheist counterfactuals with respect to moral goodness and badness, rather than rightness and wrongness. Here is Morriston offering such an argument:

> Could Craig get away with saying that love and generosity and justice and faithfulness and the rest are good-making only because there is a God who is loving and generous and faithful (etc.) to the maximum possible degree? This alternative seems incredible to me. It implies that if there were no God who perfectly exemplified them, these properties would count for nothing. A person could be as fair-minded and loving and generous and faithful as you please and still fail to be morally better than a cruel and malicious person. Why would a discerning and consistent atheist have to think a thing like that (2011, p.22)?

It seems that when Morriston considers a counterfactual world in which God does not exist, he has the strong intuition that love, generosity and faithfulness are still good. Hence he describes the alternative claim – that they are not good in that world – as incredible. He holds, then, that natural properties like *being loving* have good-making powers that do not depend in any way on theistic properties.

But there is no strong reason why atheist counterfactuals cannot be used to support conclusions about moral obligations. This is implicitly acknowledged even by Adams, who anticipates objections to his theistic theory of wrongness on such counterfactual grounds (1979a, p.147). Thus I will be happy to extend the arguments of these critics from moral value to moral wrongness. Let us call these arguments *counterfactual arguments* against voluntarism. Moral arguers have offered several responses to these arguments. Below I show that only one response succeeds, but it depends on an account of the making relationship that moral arguers cannot adopt if they want to defeat all naturalistic normative theories.

I will start by critiquing those responses that fail. Craig offers the least sympathetic response to the counterfactual argument. He seems to simply deny having any intuitions to the effect that torture or rape would still be wrong in an atheist world. This is expressed when he says things like "Why, given atheism, think that inflicting harm on other people would have any moral dimension at all? … Of course, I agree that it *is* wrong, since I am a theist. But I can't see any reason to think that it *would* be wrong if atheism were true (2008, p.177, emphasis in original)." So on Craig's view, there just is no problem here. He does not have the relevant intuitions. And those that do must be, in Craig's view, mislead by some mistaken assumption.

It seems to me, however, that when Craig conceives of an atheistic world and asks whether any inflicting harm would be wrong, he is asking whether there are *conceptual connections* between INFLICTING HARM and MORAL WRONGNESS. As I argued in §4.3, Craig's major objection to natural reductionism is that natural concepts do not have the right sorts of conceptual connections to moral concepts. It would not be surprising if conceptual connections are what he is looking for when considering counterfactuals. Since he, like many other people, see no conceptual connections between INFLICTING HARM and MORAL WRONGNESS he concludes that inflicting harm would not be wrong in an atheist world.

This focus on conceptual connections, which Craig does not always make explicit, has led to a lot of miscommunication between him and his critics. Consider how Wielenberg responds to Craig's claim that in the absence of God, there is nothing special about human beings that makes it wrong to harm them. Wielenberg points to various features of human beings that would still distinguish them from morally irrelevant entities even in an atheist world; "Human beings can reason, suffer, fall in love, set goals for themselves, and so on. God or no God, human beings obviously differ when it comes to their intrinsic properties from dogs and mere lumps of slime (2014, p.51)." For Wielenberg these properties have obvious moral relevance. But for Craig, citing these properties misses the point. Because Craig does not see any conceptual connection between these properties and moral properties, his concerns have not been addressed.

Plausibly, Craig's mistake is to think that the only kind of relationship there could be between natural properties and moral properties is a reductive one. This ignores the possibility that natural properties could *make* actions wrong rather than be the reduction base of wrongness. Even if there is a conceptual constraint on reduction of the kind discussed in §5.2, it does not

follow that there are conceptual constraints on making. Indeed, first-order normative theories concerned with making are not typically presented as offering conceptual connections. A normative utilitarian who thinks that maximising happiness is the sole, fundamental right-maker may not be the least inclined to think that there is a *conceptual* connection between MAXIMISE HAPPINESS and MORAL RIGHTNESS.

If Craig is mistakenly focussed too narrowly on conceptual connections and reduction, this would make sense of the disconnect between him and even other moral arguers with respect to counterfactual arguments. Other moral arguers think that there is a problem to take seriously here. For instance, when Wainwright discusses Adams' reductive voluntarism, he notes that it might entail that in atheist worlds, moral concepts have no referent. But rather than finding this entirely unproblematic he says, "whether this consequence is acceptable, however, seems doubtful (2005, p.95)." And the fact that other moral arguers have produced the below responses to the counterfactual argument suggests they think it cannot be dismissed so breezily. In light of Craig being the outlier here, it is more plausible that his professed absence of the relevant intuitions results from some error.

Adams has a more promising reply to the counterfactual argument. He is upfront about being committed to the claim that nothing is morally wrong in an atheistic world. Adams thinks that moral terms are rigid designators. So if, as he believes, our moral terms refer to theistic properties in the actual world, they refer to theistic properties in every world. Therefore, our moral terms have no referents in an atheist world. Nonetheless, he thinks that this implication can be made more palatable. He argues that the equivalent moral terms used by *inhabitants of the atheist world* would successfully refer. As Adams says:

> Even in a world without God, however, the best remaining alternative to divine command metaethics might be correct in the following way. In such a world there could be people very like us who would say truly, "Kindness is right," and "Cruelty is wrong." They would be speaking about kindness and cruelty, but not about rightness and wrongness. That is, they would not be speaking about the properties that are rightness and wrongness, though they might be speaking about properties (perhaps natural properties) that they would be calling 'rightness' and 'wrongness'. But they would be using the words 'right' and 'wrong' with the same meaning as we actually do. For the meaning of the words, I assume, leaves open some questions about the identity of the properties they designate (1979a, p. 147).

Adams is saying that the situation is similar to how our term 'water' refers to H_2O across worlds. Strictly speaking there is no water on twin earth if the translucent blue liquid that falls from the sky, comes out of taps etc is XYZ and not H_2O. Nonetheless, inhabitants of twin earth refer to XYZ with *their* term 'water'. There is a kind of meaning shared by both our term and their term but due to differing contexts of utterance, they have different referents. Both XYZ and H_2O are capable of fulfilling the role laid out by this shared meaning.

Likewise, Adams can claim that both theistic properties and natural properties are, in different worlds, capable of fulfilling the role laid out by the meaning of moral concepts. So when we have intuitions that there would still be moral properties in an atheist world, what those intuitions are picking up on is the fact that there would still be properties capable of playing the relevant moral role in an atheistic world. Thus, these intuitions get something right. It is just that they to fail to show that *our* moral terms do not refer to theistic properties.

The biggest problem for Adams' approach is his assumption that our intuitions do not discriminate between something being the referent of *our* terms and something being the referent of *other speakers'* equivalent terms. But the twin earth thought experiment shows precisely that our intuitions *do* discriminate in this way. Grant that in the actual world WATER refers to H_2O. Now consider a counterfactual world where on twin earth, XYZ rather than H_2O

plays the water role and ask, is there water in this world? The canonical view is that our intuitions say *no,* there is no water. They say that *even though* something plays the water role, and even though twin earther's equivalent concept refers to XYZ. This is precisely how we reached the conclusion that WATER is a rigid designator in the first place. Our intuitions said that on twin earth, there is no water, and so we devised a semantic theory to reflect that.

If MORAL WRONGNESS is a rigid designator as Adams claims, then when our intuitions say that there is still moral wrongness in an atheist world, that is evidence that our moral terms refer to something in that world. And whatever it is in that world that our terms refer to, that is what they refer to in every world. But since there are no theistic properties in an atheist world, theistic properties cannot be what our moral terms refer to in that world, or any. Thus, Adams fails to defeat the counterfactual argument.

One might think that the solution is to amend the account of moral semantics that Adams gives. Perhaps reductive voluntarists should simply drop the claim that moral concepts are *rigid* designators. So, rather than referring to the same property in every world, moral concepts can refer to different properties in different worlds. This allows moral arguers to say that our moral terms refer to theistic properties in the actual world, but those same terms refer to natural properties in atheistic worlds. When we intuit that harm would be wrong in an atheist world, those intuition genuinely tell us about the application conditions of *our* concepts. Yet, this is not evidence against our concepts referring to theistic properties in the actual world.

Unfortunately, this modification leads to a problem with multiple realisability. If the semantics of moral concepts work like this, then they work like how functionalists in the philosophy of mind say mental concepts do. Yes, this allows concepts like PAIN to have different referents in

different worlds. It could refer to c-fibres firing in the actual world and ectoplasmic states in some other world. But it also allows for different referents in the *same* world. One of the advantages functionalism boasts is that it does not commit one to saying that pain in humans is realised by the same physical states as pain in, say, octopuses. After all, the brains of these creatures are rather different. So long as these different physical states play the same functional role, all of them realise pain.

Moral arguers want MORAL WRONGNESS to refer to a theistic property in the actual world. But they want it to *only* refer to a theistic property in the actual world. This will be hard to maintain if they rely on functionalist semantics to escape the counterfactual argument. The world the argument asks us to conceive is one where the natural states of affairs are the same as in the actual world – all that is different is that God does not exist. The response to the argument we are considering is that some natural properties fulfil the role of moral wrongness in that world. Indeed, our concept of moral wrongness refers to natural properties in that world. But those natural properties have the same distribution, and causal and epistemic effects in that world as in the actual world. Again, the natural states of affairs are the same across both worlds. So if natural properties fulfil the moral wrongness role in the atheist world, they should also fulfil it in the actual world. Thus, even if MORAL WRONGNESS refers to theistic properties, it will *also* refer to natural properties.

Of course, Adams is known for claiming that MORAL WRONGNESS refers to the property that *best* fulfils the wrongness role: "what we can discover *a priori,* by conceptual analysis, about the nature of ethical wrongness is that wrongness is the property of actions that best fills a certain role (1979b, p.66)." Perhaps, then, multiple realisability is avoided because the role of the *best* can only be occupied by a single property. One can claim that theistic properties are

the best role-satisfiers in the actual world and therefore they are the only referents of MORAL WRONGNESS in the actual world. Yet some of the natural properties in the actual world satisfy the role well enough such that in a world where God does not exist, they count as the best role-satisfiers. So the Counterfactual Argument is still defeated.

It is, however, ambiguous as to what is meant by the *best* role-satisfier. On one reading, being the best is included within the conceptual analysis of MORAL WRONGNESS itself. Within the semantic content of the concept we find the information that wrongness is the property that is best – has the highest degree of – objectivity, extensional adequacy, epistemic regulation, and reason-giving force. On another reading, being the best is merely a requirement of an epistemic principle. The principle is that when two theories conflict, one should prefer the theory with the greatest theoretical virtues.

The first reading leaves a lot unclear. Suppose that one property has a higher degree of extensional adequacy, and the other has a higher degree of reason-giving force. Which counts as the best? There seems no principled way to answer this question. Further, this reading seems to make moral wrongness too dependent on contextual features. A property that realises wrongness in one world can have all the same intrinsic features, and the same causal profile, in some other world and yet not realise wrongness in that world because elsewhere in the world, some other property has those features to a higher degree. Surely wrongness should depend on features of properties that are more metaphysically fundamental than how they rank next to other properties. As for the second reading, it guides our choice of theory when two theories conflict. But on the functionalist semantics we are considering, reductive voluntarism and natural reductionism do not conflict. Both turn out to be true. So either way, appeal to the 'best' role-satisfier fails to save this response to the counterfactual argument.

Moving on, Murphy also has a reply to the counterfactual argument;

> If you are a theist … the "nearest" world in which there is no God is outrageously remote. It is an impossible world, a deeply, deeply impossible world. It is of the essence of every possible creaturely substance that it is a creature. It is of the essence of God that all things distinct from God depend on God. When I try to carry out this thought experiment seriously, as a theist, I go blank (2018, p.354).

According to Murphy, counterfactual thinking is of no use here. The theist who seriously tries to conceive of an atheist world where there are creatures who harm each other will realise that such a world is too impossible to get any cognitive grip on. The theist should "go blank" when trying. Without being able to get a clearer conception of such a world, the theist cannot have any moral intuitions triggered by it. Thus, the counterfactual argument cannot give the theist a reason to deny Moral-Theistic Tracking.

There is a difference, however, in theists being able to defend *themselves* against the counterfactual argument, and theists being able to offer *atheists* a reason to discount the argument. Perhaps those who are already theists can utilise Murphy's response to retain their rational justification for believing reductive voluntarism. But moral arguers want to give non-theists a reason to endorse reductive voluntarism. Reflective non-theists will find the counterfactual argument a barrier to endorsing it. They have intuitions that harm would still be wrong in a world without God. What moral arguers need to do is give *non-theists* a reason to distrust these intuitions.

Now Murphy considers how things look from the perspective of a non-theist, indeed an atheist:

> If you are an atheist, and you ask "what value would creatures have without

> God?", the "nearest" world is the one we live in. So just ask: "what value do they have?" (2018, p.354)

The point here is that atheists do not need to consult their intuitions about some *counterfactual* atheist world. They believe the *actual* world is an atheist world. So, as Murphy suggests, atheists just need to consult their intuitions about whether things like rape and torture are wrong in the actual world. And that's correct but it does not solve the problem. All it shows is that the difficulty facing moral arguers need not be framed in terms of non-theists' intuitions about counterfactual worlds.

Think of it like this. Everyone's moral intuitions are sensitive to the presence or absence of certain features. When coercion is present in sex, we have the intuition that something wrong has happened. But when coercion is absent, all else being equal, we lack the intuition that something wrong has happened. This suggests that the wrongness of certain sex acts depends, at least in part, on the presence of coercion. Atheists do not believe that any action has a theistic property like *contravening a divine command.* Yet their intuitions tell them that plenty of actions are wrong. This gives them some reason to believe that the wrongness of actions does not depend on God, contrary to reductive voluntarism. These are the intuitions moral arguers need to undermine. They need to offer atheists some reason to believe that their intuitions are selectively insensitive to the dependence of wrongness on theistic properties. Murphy provides no way of doing that.

Here is where Baggett & Walls might offer a helping hand. They argue that the intuitions contemporary non-theists have about which actions are wrong are historically contingent in a way that should give non-theists pause:

> For much of the history of the world, belief in intrinsic human dignity and value was *not* common. There were reasons this conviction gradually grew and replaced the older, "common-sense" views, and among them was the revolutionary teaching that human beings were made by God, in his image, for a reason and purpose. The sanguine treatment of this historical contingency as easily eliminable without seriously jeopardizing such sacred convictions is presumptuous… (2016, p.142).

Their argument is not developed any further than this and so it remains ambiguous and admits of a few interpretations. There are at least a couple interpretations that are relevant to our interests. Both start with the following empirical claim. The kind of actions contemporary non-theists believe are wrong reflect intuitions about intrinsic human dignity. These intuitions may seem to arise spontaneously and naturally in many non-theists but that was not the case for many of their cultural predecessors. They had to consciously reason to belief in human dignity. And they arrived at that belief by inferring it from their belief that humans had a theistic property: *made in the image of God*. So even though these theological beliefs have been abandoned by many, it remains the case that we today have the intuitions about human dignity that we do because of these beliefs.

Where the two interpretations of the argument differ is in the epistemological or normative principles they add to this empirical claim. But before we look at those principles, we should note that the empirical claim itself is hard to establish. The claim is that the theological belief in the *imago Dei* came first and then eventually gave rise to naturally felt intuitions about human dignity. But for all that can be shown empirically, the reverse is true. Perhaps humans started to intuit the value of human life, and they then explained how humans could have such value by linking human beings to God as his image bearers. Maybe the intuitions came first, and the theological beliefs followed.

Adopting this picture would not require denying that the theological beliefs had an important

role in spreading these intuitions and embedding them in a culture. Even this, however, may be partly explained by a certain propensity among humans to naturally have intuitions about human value. One can imagine that belief in the *imago Dei* was persuasive to many because it made sense of an intuition they already had, even if that intuition was dim and pushed to the periphery of awareness. It is one thing, then, to grant theological beliefs an important role in explaining how widespread and deeply felt intuitions about human value are today. It is another to say that the theological beliefs came first and so are responsible for the origin of the intuitions. This particular empirical claim would seem hard to verify based on the minimal evidence we have of the relevant period of history.

Let us, however, put the empirical concerns aside and grant Baggett & Walls their account of history. The first argument that builds on their historical claims goes like this. Atheists, even if not all non-theists, should grant that inferring moral beliefs from theological beliefs is an unreliable belief-forming method. After all, atheists believe all positive theological beliefs are false. Thus they are likely to be a poor guide to reality, moral or otherwise. However, history shows that belief in intrinsic human dignity arose from this belief-forming method. And ultimately, so do atheists' intuitions about human dignity. Thus, atheists who become aware of this history should therefore consider their intuitions about human dignity to be unreliable. They should cease to form beliefs on their basis.

The problem with this argument in our context is that moral arguers do not want to undermine atheists' justification for believing in human dignity and all the standard moral beliefs that come with it. They want to leave those beliefs in place but argue that a fully naturalistic normative ethics cannot account for them. It is on that basis that they want to argue for theism. All moral arguers want to do is give non-theists a reason not to trust the insensitivity of their

moral intuitions with respect to theistic properties. This version of Baggett & Walls' argument goes too far beyond that. It undermines intuitions about human dignity in general. It may be an interesting argument, perhaps ultimately successful in its own right, but it does not aid the Divergence Argument, or the metaphysical moral argument for theism.

The second argument claims that the original inferential grounds of a moral belief reflects the relevant moral-maker. So, if belief that humans have moral value originally arose because it was inferred from the belief that humans are made in the image of God, this is strong evidence that being made in the image of God *makes* humans have moral value. Further, non-theists should not take the insensitivity of their moral intuitions with respect to theistic properties to be strong counter-evidence against this. That is because there is a psychological and cultural mechanism that, over time, leads our intuitions to be insensitive to the moral-maker revealed by the original inference.

The claim is that, originally, our predecessors guided their judgments as to which people had intrinsic moral value by consulting their judgments about which people were made in the image of God (are slaves, foreigners?) But overtime, as they got more and more familiar with the people that counted as the image of God, it became easier and easier to identify which people had intrinsic moral value. In fact, it became so easy that one no longer had to consciously guide one's judgment about moral value by one's judgment about the *imago Dei*. One just had an instilled habit of regarding *all* people as having intrinsic moral value. Over time, as the epistemic role of judgments about the *imago Dei* continued to recede, people started to think that theistic properties were irrelevant to people's moral value. That is, as theistic properties became less epistemically important, they came to be seen as less normatively important. Hence, we arrive at today, where the moral intuitions of many non-theists are not sensitive to

theistic properties.

The overall picture, then, is that the more time has passed, the more our moral intuitions have drifted away from sensitivity to what really makes humans have intrinsic value. It was clear at the beginning that being made in the image of God was the fundamental moral-maker. But our epistemic access to that truth has degraded over time. The problem with this argument, however, is that implies an implausible epistemology of normative ethics. It implies that discerning fundamental moral makers is primarily an empirical, historical task. We need to uncover the original inferences that produced our moral beliefs. Most normative ethicists would reject this methodology. Besides, many of the relevant historical events will be inaccessible to us, resulting in moral scepticism. Both versions of the argument therefore fail.

Indeed, all the available responses from moral arguers to the counterfactual argument fail. At least, all those that allow for theistic normative theories do. There is one response that does not fail. I unpack that response now.

There is a simple reply to the counterfactual argument. Divine command theorists just need to claim that our intuitions about atheist counterfactuals are only sensitive to *first-order normative truths,* truths about wrong-makers. They are not sensitive to meta-ethical matters, like what moral properties reduce to. Thus voluntarists can claim that the fundamental moral-makers are natural properties, but the reduction base of moral properties themselves are theistic properties. But this sort of response only works on a view of the making relationship that allows voluntarism to be combined with natural moral-makers. These are the views of the making relationship that are excluded by attempts to motivate Moral-Natural Divergence.

We can illustrate how this sort of response can work on Enoch's law-content view of the making relationships. The claim would be that our intuitions are only sensitive to the content of what is commanded, not the divine nature of the commands themselves. So suppose God's only command has the content *maximise happiness*. Our intuitions would work like this: when we conceive of a situation, actual or possible, what triggers our moral intuitions would just be whether happiness is being maximised or not. Our moral intuitions simply will not be sensitive to whether God exists or has issued any commands. For our moral intuitions to be sensitive to that, God's commands would need to reference his existence or commands in their content. The content would have to be something like *maximise happiness only if I (God) exist.*

Harrison actually makes use of this sort of reply:

> … we must once more remember that with divine psychologism our rational intuitions are messages from a god, and so they need to be interpreted as such. Bearing this in mind, imagine you find a note from your partner reading "I want you to be happy even if I do not exist." What does this mean? Is it good evidence that there will exist a favouring of you being happy, even if our partner does not exist? No, of course not. The note is good evidence that your partner wants you to be happy for your sake, or for someone else's sake, rather than for her own sake. That is, it is telling you something about the nature of your partner's desire for your happiness, not the conditions of its existence (2018, p.126).

So long as God has issued commands that, for the most part, refer only to natural properties in their content, the divine command theorist can explain our intuitions about the moral-making powers of natural properties in atheist worlds as being sensitive only to this content. Of course, this requires the divine command theorist claiming that God communicates his commands through our intuitions, and that God would have reason to make our intuitions sensitive only to their content. But divine command theorists have independent reasons to adopt these commitments.

Firstly, as we saw above, divine command theorists *already* claim that God communicates his commands, in part, through our moral intuitions. They are committed to that because they are committed to some communication constraint on God's commands. It seems clear that if God is reliant on special revelation, he has not met that threshold for most people. So divine command theorists claim that God has communicated to everybody by means of moral intuition. Again, as I noted above, it is controversial as to whether the appeal to moral intuition suffices to remove all worries about whether God has met the communication constraint. But the point here is that if divine command theorists have a chance of arguing that the communication requirement is met, an appeal to moral intuition will have to be made.

Secondly, if God commands *x,* what God wants us to do is *x.* God wants our intuitions to track the actions we are obligated to do. So long as God meets the relevant communication requirement, there is no reason for God to also ensure that our moral intuitions track the metaphysical conditions of these moral obligations. That knowledge is likely to be a lot less practically important than knowledge of first-order normative truths. One might object that God would not keep the divine nature of moral obligations hidden. For God would not want to hide his existence for human beings. But this objection takes one into the territory of the hiddenness argument against theism. If that argument works, then theism as a whole is false, not just some specific theistic meta-ethical position. Overall, then, divine command theorists have good reason to deny that intuitions about atheist counterfactuals show moral properties to have no dependence on theistic properties at all.

One might think that this sort of reply could be adapted to work with a different view of the making relationship. Perhaps God only communicates the content of his commands, even though that content does not reveal the moral-makers. That way, the making relationship could

be held by a theistic property even if the content of God's commands refer to natural properties. The problem with this reply is that it generates too much scepticism about first-order normative ethics. God has not communicated the moral-makers on this view. So while we ordinarily take our intuitions to be good indicators of the moral-makers, we can no longer do that. This seems like too much of a price to pay.

To take stock, moral arguers want to motivate the Moral-Natural Divergence premise in the Divergence Argument. This is their second major argument for preferring reductive voluntarism to natural reductionism. But to motivate that premise, they need to show that all naturalistic normative theories are false. This requires them adopting certain views of the making relationship. But these views do not allow moral arguers the only successful response they have to the Counterfactual Argument. Either way, then, moral arguers cannot motivate the Moral-Natural Gap in a way that is dialectically effective against atheists. They either cannot refute all naturalistic normative ethics, or they cannot defend voluntarism against an important objection. Having defeated the two major arguments in favour of reductive voluntarism, I move in the next chapter to consider causal voluntarism.

7

The Explanatory Depth Argument

We have been looking at attempts by theistic moral arguers aim to motivate P2 [it is plausible that theism accommodates objective moral properties and more plausible that it does so than that naturalism does so]. So far, we have looked at attempts to motivate it by arguing that reductive voluntarism is more plausible than its naturalistic competitors. This requires showing that reductive voluntarism has an advantage over some *specific* competitors, like natural reductionism. In the past couple chapters I have argued that the two main arguments for preferring reductive voluntarism to natural reductionism fail. Thus, moral arguers have yet to successfully motivate P2 by appeal to reductive voluntarism.

But that does not yet rule out a metaphysical moral argument for theism. For a non-reductive voluntarism might yet turn out to have an advantage over all its rivals. Philip Quinn's *causal* divine will theory denies that moral properties reduce to theistic properties. Now as before, for this theory to have an advantage over all its naturalistic rivals, it must at least have an advantage over some *specific* competitor, like non-reductive naturalism. The main argument for preferring causal voluntarism over non-reductive naturalism is that the former can offer a deeper, simpler explanatory stopping point for the existence of morality than the latter can. But I argue in this chapter that this argument fails. For causal voluntarism must posit at least one brute moral fact – plausibly one about God's moral authority. Non-reductive naturalism can also appeal to just one brute fact – one about the respect due to persons. And attempts to show that voluntarism's one brute fact has greater internal simplicity than brute realism's fact, fail.

§7.1 CAUSAL VOLUNTARISM

Quinn says, “it seems clear to me that being morally wrong and being contrary to the commands of a loving God are distinct properties (2013, p.82).” So there is no appeal to identity relationships in Quinn’s divine command theory. Rather, the canonical interpretation of Quinn is that he appeals to a causal relationship, as indicated in this passage;

> One intuitive picture that divine command theorists have found attractive presents God as an agent bringing about moral obligations and prohibitions by means of his legislative activity. According to this picture, God imposes or creates moral obligations and prohibitions by his commands or acts of will. Moreover, moral obligations and prohibitions are depicted as being directly and solely dependent on acts of the divine will. This way of thinking about the connection between God’s commands and moral obligations suggests a causal relation between divine commands and moral duty (Quinn 2006, pp.41-2).

However, things are a little more complicated; it seems that later on, he moved away from the language of causation. In a later paper, he describes how he once defended a causal view, but then says;

> My current view is that dependence of morality on God is best formulated in terms of a relation of bringing about, though care must be taken to distinguish this relation from various causal relations familiar from science and ordinary life. In particular, the divine bringing about in question will have the following marks: totality, exclusivity, activity, immediacy, and necessity. By totality, I mean that what does the bringing about is the total cause of what is brought about. By exclusivity, I mean that what does the bringing about is the sole cause of what is brought about. By activity, I mean that what does the bringing about does so in virtue of the exercise of some active power. By immediacy, I mean that what does the bringing about causes what is brought about immediately rather than by means of secondary causes or instruments. And by necessity, I mean that what does the bringing about necessitates what is brought about (2013, p.83).

This sort of bringing about is the traditional view that theists have about how God creates ordinary concrete objects like planets. God has some active power that he exercises which brings about effects, with no need for them to be mediated by laws, instruments, or helped by

other contributing powers. Further, nothing can impede God's direct creative powers. Necessarily, if God exercises his power to create planets, then there are planets. Traditionally, theists are happy to call this sort of bringing about a *causal* relationship. So Quinn has not described anything that, in a traditional theist's mind, would typically be distinguished from causation.

There are, however, views that exclude this sort of bringing about from counting as causation. For instance, according to some views, causation must be governed or mediated by laws. While theists are not normally friendly to such views, Quinn is consciously writing for a broad philosophical audience and so may not wish to take anything for granted that excludes such views. But since I have no objection to treating Quinn's bringing about relationship as a causal relationship, I will adopt the canonical interpretation of his theory. I will regard it as a causal theory. Furthermore, Wielenberg, the proponent of natural non-reductionism who is most engaged in dialogue with theistic moral arguers, also has no problem with treating it as a causal relationship. Indeed, he models a causal relationship in his own theory on the type of causation that theists claim God can distinctly exercise (2014, p.19).

Another thing to note about interpreting Quinn is that although he uses the language of God's *commands,* he claims that the language of God's *intentions* better captures his view. When God intends that I keep my promises, God causes me to be obligated to keep my promises. God's intentions cause there to be moral obligations, God's commands merely express God's intentions. By God's intentions, Quinn does not mean God's intentions to bring about my promise keeping, as if God intends to cause my promise keeping. Rather, Quinn is referring to God's *preference* that I keep my promise. Nonetheless, I will continue using Quinn's language of intentions and will.

Quinn is aware that focusing on intentions raises a problem for him. God has intentions for lots of things, and not all of these intentions plausibly cause moral obligations. Suppose that I break a promise. God intends for me to apologise. But if I have an unqualified obligation to apologise, it seems that I have an obligation to do something wrong so that I can then apologise for it. That's an absurd result. So Quinn needs to pick a particular type of intention to do the moral work. He says that God's *antecedent* intentions cause moral obligations. The following passage gives us a sense of the kind of preferences Quinn has in mind:

> God surely prefers that I keep [the promise], other things being equal. Assume God also antecedently intends that I keep my promise, which makes it obligatory for me to keep it. If I break my promise, I create a bad situation by violating an obligation. But assume that God, in addition, antecedently intends that I apologize if I break my promise, which makes it obligatory for me to apologize if I break it. If I am in the bad situation of having broken my promise, then my obligation is to apologize. If I fail to apologize, I violate a second obligation (2013 p.84).

So before I break a promise, or perform any action, God has antecedent intentions for me *conditional* on the various situations I might find myself in. As such, I do not have an absolute obligation to apologise. Rather, I have an obligation to apologise in the situation in which I have done wrong.

Quinn's own view, then, is that God's antecedent intentions cause (with totality, exclusivity, activity, immediacy, and necessity) moral obligations. It does not seem to me, however, that the kind of antecedent preferences that Quinn picks out could have that causal power. The problem is that those intentions do not include anything about moral obligations in their content. On Quinn's view, if God intends that I promote the well-being of my fellow creatures, then that intention puts an obligation on me. His view does not require that God prefer that I be *morally obligated* to promote the well-being of my fellow creatures. This, however, seems

to violate the following plausible principle:

> Necessarily, a divine intention causes *x* with totality, exclusivity, activity, immediacy, and necessity, only if God has an intention that *x.*

That is, if God's intention causes some effect, God must also have intended that effect. It seems impossible, for instance, that God could cause (with totality, exclusivity, activity, immediacy, and necessity) there to be a planet *with green grass on it* while only intending that there be a planet. For it would seem that God's power to cause green grass had been activated outside of his control. But it is impossible that God not have control over his own direct causal powers.[43] So God must intend more than that I just promote the well-being of my fellow creatures. He must intend that I be *obligated* to do so. Otherwise, impossibly, God's power to cause obligations seems activated outside of his control. Further, the natural view would be to hold that God's intentions that creatures be obligated are the very same intentions that actually *cause* there to be moral obligations.

I suspect that Quinn appealed to intentions without moral content to do justice to the intuition many theists have that God's will or intentions are obligating even when they do not have moral content. Many theists think that if God prefers that one give to charity, then giving to charity is obligatory, even though the content of God's preferences was not that one be *obligated* to do so. But Quinn can still vindicate this theistic intuition. God can first intend that it be obligatory to do whatever God antecedently intends, even if those intentions do not have moral content.

[43] Some Christians might hold that there is an actual counter-examples to this principle. The first is that, according to Luke 8:46, Jesus (who is identical to God the Son), seemed to have his healing powers activated without him intending that they be activated. There are, however, various interpretations of this strange incident consistent with my principle. Firstly, any strategy one has for showing Jesus' human limitations to be compatible with essential divine attributes will likely accommodate this incident. Secondly, one might hold that the healing power that emanated from Jesus was not Jesus' own but rather God the Father's who transmitted it through Jesus.

From then on, it will be obligatory to do whatever God antecedently intends, whether or not moral content is included. So my modification of Quinn's view does not lose anything theists would want to keep.

One final issue in interpreting Quinn concerns what exactly it is that God causes when God causes there to be moral obligations. There are at least two possibilities. The first possibility is that what God directly causes is the instantiation of moral properties. So when God intends that every rational creature be obligated to promote well-being, what that causes is the instantiation of *being obligated to promote well-being* on every rational creature. The second possibility is that what God directly causes is the existence of a moral law. This moral law then confers the relevant moral property, like *being obligated to promote well-being,* on every rational creature. This is like how a human government can pass laws which then confer a legal status on certain actions of its citizens. On the first option, God directly causes the moral properties. On the second option, God causes the moral properties indirectly by directly causing a moral law. This second option does not undermine Quinn's claim that God brings about moral obligations *immediately.* For on the second option, the moral law is identical to the obligation, and the moral law *is* brought about immediately.

Whichever interpretation is adopted does not substantially affect my assessment of causal voluntarism. But because it makes for less cumbersome exposition, I will adopt the second, law-based based interpretation. Besides, one strength of that interpretation is that it fits well the picture Quinn (2006, p.42) invokes of "God as an agent bringing about moral obligations and prohibitions by means of his legislative activity." Further, it fits well with the content interpretation of the making relationship. And as I argued in the previous chapter, that view of making gives voluntarists a response to the counterfactual argument against voluntarism.

§7.2 BRUTE REALISM

As already noted, moral arguers must show causal voluntarism to be more plausible than natural non-reductionism. According to natural non-reductionism, naturalism is true, but moral properties do not reduce to natural properties. This view, sometimes called Robust Realism, has recently increased in popularity. It is defended by Enoch (2011), Parfit (2013), and Wielenberg (2014), among others. They all have slightly different views, but if there is just one that causal voluntarism cannot beat, then P2 is unmotivated. I will focus on Wielenberg's natural non-reductionism. This is partly because Wielenberg is the only defender of the view who has significantly engaged moral arguers.

Like Quinn, Wielenberg claims that the instantiation of irreducible moral properties is *caused*. Certain natural properties cause moral wrongness, or moral rightness to come about. An instance of, say, *inflicting pain for fun* causes moral wrongness to be instantiated. And as I indicated above, Wielenberg holds that this causation looks much like the kind that theists claim that God can exercise;

> A paradigmatic example of the sort of robust causation I have in mind is the causal relation that many theists take to hold between a state of affairs being divinely willed and the obtaining of that state of affairs. The relation between p's being divinely willed and p's obtaining is a causal one of a particularly robust sort. It does not depend on the existence of a law of nature connecting acts of divine will with states of affairs (indeed, theists typically hold that whatever laws of nature hold are themselves the products of divine willing). Theists typically maintain that if God wills that p, this necessarily brings it about that p obtains. And there is no reason that acts of divine will cannot be simultaneous with their effects. Necessarily, if God wills at time t that p obtains at time t, then p obtains at time t. (2014, p.18-9)

So according to Wielenberg, natural properties causally instantiate moral properties directly and necessarily. This causation is not governed by laws. Indeed, there is no deeper explanation of these causal transactions at all;

> Why does being an instance of torturing someone just for fun entail moral wrongness? Because being an instance of torturing someone for fun *makes* an act wrong. But why does being an instance of torture for fun make an act wrong? Perhaps further explanation is available: for example, perhaps torturing just for fun never maximizes utility and failing to maximize utility makes an act wrong. But why does failing to maximize utility make an act wrong? Eventually we hit bottom; no further explanation is available (2014, p.24).

Wielenberg is unapologetic about this: "I don't see why possessing this sort of explanatory bottom is a problematic feature for a view to have (2014, p.24)." It is not clear, however, where exactly this explanatory bottom is placed in his account. Consider that we might ask why does *this* instance of torturing an innocent for fun and *that* instance of torturing an innocent for fun *both* cause wrongness to be instantiated? If there is no explanation of that – if that is where one reaches explanatory bottom – it could seem an intolerable coincidence that *every* instance has the same consequence. Especially if we hold that every instance across *every world* has the same consequence. The regularity of the causal effects of each instance seems to cry out for explanation.

This is where Wielenberg might appeal to abstract universals. Abstract states of affairs and properties are entities he invoked in the fullest statement of his account (Wielenberg, 2014). He has since questioned the role of these entities; "in *Robust Ethics* I freely help myself to various abstract entities – e.g. properties and states of affairs – it *may* be that those aspects of my view are dispensable (Craig & Wielenberg, 2020, p.40)." But here is where these entities can be put to work. Perhaps the property *torturing an innocent* is a universal with a particular nature. The nature of the universal is that its concrete instantiations cause moral wrongness. Thus the reason

why every instance has the same causal effect is that each instance shares that same nature. Of course, we might ask why this property has this nature. And Wielenberg might reply that there is no explanation of that. But reaching explanatory bottom at this stage might seem less extravagant. When we posit abstracta, there is just one brute fact appealed to, rather than many, as would be the case if there is a new brute fact whenever any particular instance has the same causal effect.

Positing a universal – a property – that has an essential relation (such as causation) to a moral property is just the sort of thing that Murphy proposes as a *moral law*, something that explains regularities across particular instances:

> The account of moral laws that I favor is modelled on what can be called the universals account of the laws of nature, the view formulated independently by David Armstrong (1983), Fred Dretske (1977), and Michael Tooley (1977). On their view, it is an error to think that a law of nature is a regularity plus some further feature; rather, a law of nature is a relation between universals the presence of which explains the regularity. Rather than being a relationship between the extensions of properties, it is a relationship between the properties themselves (2011, p.34).

So it does seem that there could be an important role for abstracta and moral laws to play in Wielenberg's account.[44] At any rate, interpreting Wielenberg in this way allows for a simple comparison between his account and causal voluntarism. For simplicity, we can refer to his account as *brute realism.* Both accounts posit moral laws of some sort. Causal voluntarism claims that moral laws are caused by the will of God. And brute realism claims that moral laws are uncaused, brute facts. What we need to see is whether invoking divine causation has

[44] One might think that some universalizability thesis does away for the need for abstracta. Perhaps, for instance, it is conceptually impossible for one instance of torture to be wrong and another right if the only different between them is that they are different instances. This may correctly capture the conditions of applying moral concepts but it does not tell us how anything in reality can answer to those concepts. We need a metaphysical account for that, and positing abstracta may help with that account.

advantages over leaving moral laws as brute. Note, then, that in this chapter I will shift my focus from moral properties to moral laws.

Of course, some moral arguers, notably Craig, take issue with an appeal to abstracta. He objects: "how can a physical object somehow reach out and causally connect to a transcendent, causally isolated, abstract object (2018, p.336)"? But as Wielenberg points out, he need not claim that natural properties have causal interactions with the abstract universal, *moral wrongness*. Rather, they have causal interactions with particular instantiations of moral wrongness (2018b, p.365-6). This is just like how any realist about abstracta might think of causal transactions between *x* and *y.* If the rock smashes the window, the realist does not say that the rock smashes the window universal. Perhaps, however, Craig's point is that whatever natural properties are causally interacting with, they are not interacting with physical properties. Even if they are not interacting with abstract objects, a particular instantiation of moral wrongness is no physical property. And it is still mysterious how natural properties could causally interact with such a thing.

In short, Craig would be raising an *interaction problem* for brute realism, akin to the interaction problem facing dualists in philosophy of mind. Dualists claim that mental properties are not physical properties, yet physical properties can causally interact with them. Many philosophers find something puzzling about this. What is strange is that Craig is himself a dualist about the mind. In fact, in the book he co-wrote with Moreland, they answer the interaction objection to dualism in the following way:

> This objection assumes that if we do not know *how A* causes *B,* then it is not reasonable to believe *that A* causes *B,* especially if *A* and *B* are different. But this assumption is not a good one. We often know that one thing causes another without having any idea of how causation takes place, even when the two items are

> different… Furthermore, it may even be that a "how" question regarding the interaction between mind and body cannot even arise. A question about how *A* causally interacts with *B* is a request for an intervening mechanism between *A* and *B* that can be described… But the interaction between mind and body may, and most likely is, direct and immediate. There *is* no intervening mechanism, and thus a "how" question describing that mechanism does not even arise. (2003, pp.243-4)

Wielenberg can make use of these same responses. We know that natural properties cause moral properties even if we do not know how. Perhaps we know this, or at least are justified in believing it, because we know that natural properties *make* actions have moral properties, and causation is the best account of the making relationship that we have. Further, the interaction between the natural and the moral is direct and immediate, so a "how" question about their interaction does not arise.

Craig also objects; "how do these physical situations "know" which abstract object to instantiate? What if instead of picking out moral *goodness,* some physical situation might pick out moral *badness* (2018, p.337)?" But again, if this is a problem, it also seems to be a problem for dualism in philosophy of mind. How do physical states "know" which mental state to instantiate? The answer the dualist gives is that some physical states just *do* cause some mental states. There is no agency on the part of physical states required. And Wielenberg can say the same thing about natural properties and moral properties. So we can leave these objections to the use of abstracta in brute realism aside.[45]

§7.3 THE EXPLANATORY DEPTH ARGUMENT

[45] Of course, there are also objections to positing abstract objects in general, even outside the context of meta-ethics. But these would take us far afield. And any moral arguer that objects to Wielenberg on the grounds that abstract objects are inherently problematic will be taking on significant metaphysical commitments.

In interpreting brute realism, I defended it against several objections. But the main objection moral arguers have is that its appeal to brute facts is unsatisfying. This is where voluntarism may have an advantage. Granted, Quinn himself was rather shy about whether his causal voluntarism has advantages over naturalistic competitors. His project was a defensive one, devoted to showing that divine command theory has not been shown false or unreasonable (1990, p.345-6). But other voluntarists *do* go on the offensive. While these voluntarists are typically reductive, not causal voluntarists, one can anticipate how, if they were forced to switch to the causal version, they would still view it as preferable to brute realism.

Here is what Baggett & Walls say about Wielenberg's appeal to brute facts:

> It is a dubious insistence that says, since explanations have come to an end, stopping where Wielenberg does is just as good a place as any. It is true that explanations have to come to an end, but the final explanation needs to be adequate to the task… when Wielenberg says deliberate cruelty makes something wrong, what he is saying may well be true, but it doesn't advance the discussion or offer illumination by way of explanation (2016, pp.142-3).

And Craig says:

> If our approach to meta-ethical theory is to be serious metaphysics rather than just a "shopping list" approach, whereby one simply helps oneself to the supervenient moral properties or principles needed to do the job, then some sort of explanation is required for why moral properties supervene on certain natural states or why such principles are true (2008, p.177-8).

The objection is that Wielenberg introduces brute facts too prematurely. We should expect explanation to go further and deeper than brute realism permits. Here is where these moral arguers might think that even causal voluntarism is superior to brute realism. Quinn's account goes further in causal explanation than Wielenberg's. After all, as already noted, both posit moral laws, but Quinn's provides a causal explanation of those laws rather than leaving them

as brute facts. There is an opening to argue that causal voluntarism's explanation goes deep enough.

It is, however, hard to specify how deep explanation should go, which facts cry out for explanation more than others, and so on. The argument for preferring causal voluntarism to brute realism will be on safer ground if it is framed more modestly. It just needs the following principle:

> **Explanatory Preference**: All else being equal, an account that causally explains *x* is rationally preferable to an account that does not explain *x* at all.

Rather than taking a stand on where exactly explanation should end, the Explanatory Preference principle merely states unless there is some over-riding factor, an explanation should be preferred to the positing of brute facts. This is an extremely plausible principle. Suppose one offers you a good explanation as to why apples fall downwards. It would surely be irrational to prefer to say that their falling downwards is a brute fact, unless one had some good overriding reasons to reject the explanation.

Of course, there are lots of ways one could challenge the claim that all else *is* equal when it comes to causal voluntarism. One might argue that Quinn's view merely pushes the brute fact back a step. After all, many theists think that God's existence is a brute fact. It may not be clear that there is any advantage to positing God as the brute facts rather than the moral laws themselves. But there *is* a potential advantage to ending explanation at God rather than the moral laws. There are perhaps numerous moral laws, but God is just one entity. Simplicity is a

theoretical virtue and the fewer entities a theory posits at explanatory bottom, the better.

Granted, this requires that there *are* multiple fundamental moral laws. It assumes that we do not have a single fundamental moral obligation to, say, maximise happiness or act in conformity to the categorical imperative. If we do have a single fundamental obligation, then there does not seem to be any advantage in positing a single being, God, behind it. If anything, positing such being seems a gratuitous increase in complexity. So causal voluntarists can only argue for an advantage in simplicity if they argue against certain theories in first-order normative ethics. Of course, they might be willing to do that.[46] So, let us assume for the sake of argument that we have multiple fundamental moral obligations.

There are other objections. One might think that the ontological costs of positing God outweigh the gains in explanation. Or perhaps some version of the Euthyphro objection succeeds. But my objection to the argument will not take these routes. In the remainder of the chapter, I argue that causal voluntarism can be interpreted in two different ways. On the first interpretation, God causes moral obligations without the help of any ontologically distinct normative facts. I will argue that this is impossible. So the second interpretation is the only viable one, where God causes moral obligations *with* the help of ontologically distinct normative facts. However, I argue that this leaves causal voluntarism without any advantage over brute realism. More explicitly, my argument is structured as follows:

(1) Either (a) God causes moral obligations without the help of any ontologically distinct normative facts or (b) God causes moral obligations *with* the help of

[46] This does not require them arguing against *all* naturalistic normative theories, just those that posit a single fundamental moral-maker. So this is not as burdensome a commitment as that carried by the Divergence Argument (§6.2).

ontologically distinct normative facts.

(2) It is impossible that God causes moral obligations without the help of *any* normative facts.

(3) It is impossible that God causes moral obligations with the help of normative facts that are reducible to God.

Therefore, God causes moral obligations with the help of ontologically distinct normative facts.

(4) If God causes moral obligations with the help of ontologically distinct normative facts, then causal voluntarism does not have any explanatory advantage over brute realism.

Therefore, causal voluntarism does not have any explanatory advantage over brute realism.

Premise (1) just states the two interpretations of Quinn's view and, indeed, is a logical truth on the assumption that God creates moral obligations at all. The first interpretation, given in option (a), allows either that God creates moral obligations without the help of any normative facts at all, or with the help only of normative facts that are *not* distinct from God. Such facts would have to be facts that reduce to God (where I understand this broadly to cover God, or his nature, or his actions, or his attitudes etc).[47] Premises (2) and (3) claim that neither are possible. Hence the second interpretation given in option (b) is all that remains. Premise (4) claims that this

[47] Recall I argued in §3.1 that the claim that there is no robust ontological distinct between *x* and *y* is the claim that *x* reduces to *y*.

interpretation leaves causal voluntarism without any advantages. I will now defend premises (2) and (3) before moving on to defending (4).

§7.4 NO NORMATIVE EX NIHILO

One way to interpret causal voluntarism has it that God does not cause moral obligations with the help of any ontologically distinct normative facts. This might be because God does not cause moral obligations with the help of any normative facts *at all.* This strong view is suggested by Quinn's claim that God causes moral obligations with totality, exclusivity, activity, immediacy, and necessity. It suggests that God brings moral obligations about merely through causal power, not through the help of some normative principle that, say, grants God the *right* to issue binding commands or intentions. This seems to be how Murphy (2019) interprets the language of totality, exclusivity etc. God's intentions are facts that explain moral obligations in a *merely causal* way, he says. We can call this the creation of moral obligations from *normative ex nihilo.* This is analogous to the claim that God causes the universe to come into being without the help of any prior physical substances – from *material ex nihilo.*

In fact, Quinn's view probably was not that God's intentions explain moral obligations in a merely causal way. As I show later, Quinn very likely thought that God caused moral obligations with the help of some prior normative fact about God's right to command. Indeed, there seems to be a consensus among theists that *necessarily* God does not bring about moral obligations through mere causation. That is, there is consensus around premise (2) of my argument. Murphy is the most explicit about this, and I discuss his objection to creation from normative ex nihilo below. But other theists implicitly accept (2).

For one, as we will see, many of them give accounts of what the normative fact *is* that gives God the right to issue obligations. This is often done in the context of responding to the prior obligation objection to voluntarism. The objection is that God's commands cannot generate *all* moral obligations because there must at least be a prior moral obligation on creatures to obey God's commands. If God could bring about moral obligations through mere causal force, one could respond: God doesn't *need* there to be any prior normative fact, let alone a moral obligation. But this response is not given.

Given this consensus, I will not labour heavily on arguing for premise (2). The argument I give for (2) is rather rough and intuitive, although I believe it stands behind theists' rejection of creation from normative ex nihilo. I do not expect that every reader will find the argument convincing. But as I am arguing against theistic moral arguers, who plausibly accept this argument, I will not address the concerns of those who may be less sympathetic to it. However, before turning to my argument, it is worth addressing another potential reason for denying creation from normative ex nihilo.

One might think it is trivially obvious that *God* cannot create from normative ex nihilo. For necessarily, God is perfectly good and perfectly rational. But goodness and rationality are normative properties. So these facts about God's properties are normative facts. And necessarily, God manifests his goodness and rationality in his creation of moral obligations. Therefore, one might conclude, it impossible that God not be helped by normative facts. This, however, is not a good reason to reject normative ex nihilo, and I doubt it is what theists have in mind. The argument is vulnerable to the charge that it misunderstands what creation from normative ex nihilo is. God's goodness and rationality do not help him in the sense at issue. God's goodness and rationality merely *guide* God in the exercise of his powers. They do not

constitute or directly enable those powers.

To illustrate the distinction, consider two people, one of whom is morally good and the other evil, and call them *G* and *E.* Both *G* and *E* are young and physically strong. One day, both *G* and *E* witness an old lady struggling to carry her shopping bags. *G's* goodness guides *G* to exercise her powers – her physical strength – to help the lady with her shopping. Whereas *E's* evilness guides *E* to avoid lending a hand. Suppose that *G* would never assist with the shopping if *G* wasn't good. In one sense, then, *G* cannot assist with the shopping without the "help" of some normative fact – the fact of *G's* goodness. Nonetheless, all the *power* to assist with the shopping resides in *G's* muscle mass, not in *G's* moral disposition. There is a sense in which *E* has the same power to assist with the shopping. *E* is just as physically strong as *G.* The normative facts about their moral dispositions merely affect how they choose to exercise those powers.

But the claim that God creates from normative ex nihilo is not the claim that God is not 'helped' in any way at all by normative facts. It is not the claim that God is not even guided by his goodness and rationality. It is that God's *power* to create moral obligations resides in his causal powers alone. Just like how one's power to lift shopping bags resides in one's muscle mass alone. This is what premise (2) denies. And the main reason that creation from normative ex nihilo is denied is that there seems to be some unintelligibility in the idea of mere causation having the power to bring about moral obligations. We can start to see this unintelligibility by considering how Wolterstorff criticises the idea of a voluntarism that posits creation from normative ex nihilo. His criticism starts by pointing out that human commands do not generate obligations from normative ex nihilo.

Wolterstorff asks us to consider that there are clearly some cases of humans issuing commands to other humans where no obligations are generated. (Or at the very least, there are some cases of humans issuing speech acts with the surface features of commands yet failing to generate obligations). One example is an imposter military officer issuing 'commands' to the troops. The troops don't seem to have an obligation to obey these commands. What we need, Wolterstorff says, is some explanation of why this is so, and his is that:

> the imposter officer lacked the *power,* the *potestas,* to issue military commands to this body of persons; he lacked the *capacity* to do so. Hence he also lacked the *authority* to do so, that is, the permission-right to do so and the claim-right to be free to do so. (2008, p.269).

What genuine officers have that imposters lack is a certain normative power. Imposters lack the *authority* or *permission-right* to generate obligations. These are normative notions. And Wolterstorff holds that necessarily, the genuine officer's possession of that authority comes with a corresponding *standing obligation* on the part of the troops. The two are corollaries of each other:

> Let me call the permission-right and the freedom claim-right of the legitimate officer to issue commands to his troops on military matters *standing* rights on the part of the officer; let me likewise call the implied obligation of the troops to obey such well-formed commands as the officer may issue a *standing* obligation on their part, and the implied right of the officer to his troops to obeying such well-formed commands as he may issue a *standing* right on his part. (2008, p.269)

In other words, there is a prior normative context that explains why obligations are generated by the commands of a legitimate officer. Antecedent to the troops having an obligation to obey a *specific* command of the officer (like, "deploy to the garrison!"), they have a general, standing obligation to obey *whatever* he commands (within certain limits). And this prior normative fact bestows or enables a power in the officer to bring obligations about. Of course, this power is

used in conjunction with the officer's merely causal power. The officer must have the causal power to utter sounds that are recognisable as commands. But this causal power is not enough. The prior normative context here is doing serious explanatory work in how it is that the officer can generate obligations. In the human context, commands do not generate obligations out of normative ex nihilo.

The objection that Wolterstorff raises from these observations against divine command theories that posit creation from normative ex nihilo is relatively modest. Since there is a disanalogy between the human case and the divine case (as understood by these theories), the motivation for these theories is undercut. Wolterstorff says, "The theory proposes to illuminate how God generates moral obligations by pointing to an analogous phenomenon in human affairs: we generate obligations in each other by, among other things, issuing commands (2008, p.272)." If the disanalogy breaks down, then the human case can no longer illuminate how God generates moral obligations and there is no reason to endorse a divine command theory.

It is not hard, however, to see why someone might go further and claim that there is something unintelligible about creation from normative ex nihilo. Try and conceive of a situation in which a military officer successfully generates an obligation but without the help of any antecedent normative context. So the officer does not have any *authority* or *right* to generate obligations, and the troops do not have any *standing obligations* to follow the officer. Merely by being able to utter a sentence like "deploy to the garrison!" the commander brings forth a moral obligation. To my mind, there is something not fully intelligible about this. And I think many divine command theorists implicitly agree. It may not be an explicitly contradictory notion. But the bringing about of the obligation does not seem fully explicable just by appeal to the officer's ordinary causal powers.

Of course, God's causal powers are far more expansive than those of humans. So one might think that even if humans cannot create from normative ex nihilo, that doesn't mean God can't. After all, humans cannot create physical objects from material ex nihilo but that doesn't mean God can't. In response, Murphy cautions against relying too much on the uniqueness of God's causal powers:

> … it is hard to make any sense of the view that moral facts are made to obtain simply in virtue of an exercise of God's causal power… This is an entirely unfamiliar phenomenon: nowhere else do we encounter a merely causal connection between a nonnormative fact and a normative one. (The appeal to the very strangeness of divine causation itself is not sufficient to answer the objection. For there is an *extra* strangeness here: that the relationship between nonmoral and moral facts is in every case with which we are familiar a rational relationship, whereas on the causal formulation of theological voluntarism the relationship is merely causal. Creation ex nihilo does not constitute *carte blanche* to multiply strangenesses.) (2011, p.113)

It is, though, not entirely clear what Murphy's objection is. He starts off introducing the idea that creation of moral facts from normative ex nihilo is dubiously intelligible – "hard to make any sense of" – but then seems to justify that by pointing out its unfamiliarity. Unfamiliarity, however, is not sufficient for unintelligibility. Again, creation from material ex nihilo is not familiar, yet to most of us it is not unintelligible.

There might, however, be a better way to get a grip on the unintelligibility that Murphy speaks of. Consider a case of a physical object coming to be from *causal ex nihilo.* That is, a physical object comes into being without being *caused* to come into being, even indeterministically, by anything. Many of us have a strong metaphysical intuition that this is impossible.[48] This

[48] Not all of us do, and not all us that do believe it to be a veridical intuition. But recall that I am merely trying to explicate the line of reasoning that both myself and divine command theorists already agree on.

intuition is behind the premise in another argument for God that Craig endorses. The first premise of the Kalam Cosmological Argument is *everything that begins to exist has a cause.* According to Craig, this premise is "obviously true" (Craig & Sinclair, 2009, p.182). To those for whom it seems obvious, it isn't enough to postulate that God exists antecedent to a universe that begins to exist. The existence of the universe remains unintelligible unless God exercises some causal power in bringing it about.

There is, I think, a broader, more abstract principle at work behind this intuition. It is that some developments in the state of a world must be explicable, and that they are only explicable with the exercise of the right sort of powers. Such developments do not need to be understood as temporal changes. Even if God timelessly creates, the following is still true; there is a logically prior state in which there is no creation and then a logically posterior state in which there *is* creation. Now in the case of physical objects, the right sort of powers are *causal* powers. But suppose that in the case of moral obligations, the right sorts of powers are *normative* or *rational* powers. And suppose further that normative powers are not reducible to causal powers. It would follow that God's causal powers alone are not enough to bring about moral obligations.

This is precisely the line of thought that seems to motivate intuitive resistance to creation from normative ex nihilo. Murphy remarks, "it seems to me that the formulation of theological voluntarism most worth taking seriously is a formulation that takes the dependence of moral necessities on the divine will to be not a causal but a rational relationship (2011, p.133)". This is why he takes Adams' view more seriously than Quinn. He interprets Adams such that "the way that God's commands make performance morally necessary is not a causal power; it is an essentially normative power (2011, p.114)." Note the contrast between two kinds of power, causal and normative. And the language of an *essentially* normative power suggests that it is

not reducible to causal power. Indeed, there is good reason to think that many moral arguers will reject its reduction to causal power. The *concept* of a normative power seems very distinct from the *concept* of a causal power, and many moral arguers endorse a conceptual constraint on reduction, as I argued in §4.

This, then, is the thought that unites me and divine command theorists in rejecting God's creation of moral obligations from causal power alone. Bringing about moral obligations needs to be explicable through the exercise of an irreducibly *normative* power. There must, then, be some prior normative fact that helps God in his creation of moral obligations. Causation alone won't do. To be clear, the claim is not that any normative fact must be explained by some prior normative power. That would require an absurd infinite regress. The claim is that a normative fact that is *brought about* without a prior normative power is unintelligible. Compare: one need not make the claim that the universe existing without a prior causal power is unintelligible. One can simply find the claim that the universe *began to exist* without a prior causal power to be unintelligible.

Of course, it is important not to make too strong a claim about the normative irrelevance of causal powers. Causal powers can bring about changes in non-normative facts, and non-normative facts affect the normative situation. To quote Murphy again:

> A nonmoral fact can be part of what constitutes a reason to perform an action. (That you promised to φ can be cited in explaining why you have a reason to φ; your promising to φ constitutes, at least in part, the reasons that you have to φ.) It can be part of an enabling condition for that reason. (The existence of a social practice of promising can be cited in explaining why you have a reason to φ; the existence of that practice might explain why your promise has the reason-giving force that it has.) It can be cited as a defeater-defeater for a reason. (While the fact the promisee told you that you need not fulfil your promise to φ typically releases you from your promise to φ, the fact that you threatened to beat up the promisee if he or she did not tell you that you need not fulfil your promise invalidates that

> release, and can be cited in explaining why you have a reason to φ.) Pretty obviously what all of these relationships have in common is that they are rational relationships, not causal ones (2011, p.112).

What Murphy seems to be saying is that non-normative facts affect we what might call the *triggering conditions* of normative facts in a given situation. For instance, if there is a moral obligation to keep one's promises, then when one exercises one's causal powers in uttering a promise, that moral obligation is triggered. You now have an obligation to keep that particular promise you made. Likewise, when a society exercises its causal powers in bringing about an institution of promising in the first place, the moral obligation to keep one's promising now has the potential to be triggered. What causal powers cannot do, though, is bring about the moral obligation to keep promises itself. This is why premise (2) is correct.

However, one question that may have been on the reader's mind for a while now is how the above discussion bears on Wielenberg's view. After all, on his view natural properties cause moral properties, and there seems to be no other relationship in view than causation. So would Wielenberg's view also be ruled out? As the view stands, yes. But a modification will save it. The problem with Wielenberg's stated view is that it is not a thorough-going non-reductionist view. True, Wielenberg does not reduce moral *properties*. But he does reduce the *relationship* between non-moral properties and moral properties. Rather than leaving that relationship as an irreducibly normative one, he reduces it to a causal relationship. This is what he should correct. He should endorse an irreducibly normative relationship between them. Divine command theorists should not fault such a move since they seem to agree that there is the need for such an irreducible relationship.

§7.5 IRREDUCIBLE RIGHTS

God needs the help of some normative fact to cause there to be moral obligations. That was the conclusion above. In this section, I argue that this prior normative fact cannot be reduced to God, his nature, actions, or attitudes. Hence, premise (3) is correct. Once it is shown that causal voluntarists must posit an irreducible prior normative fact, I will then be able to argue that their view has no advantage over Wielenberg's brute realism.

One might think, however, that there is a more direct and aggressive attack on causal voluntarism available. One might think that the lesson to be drawn from the rejection of creation from normative ex nihilo is not merely that there needs to be some prior normative fact or other. The lesson is that there needs to be some prior *moral obligation.* If that's right, then causal divine command theory doesn't merely fail to have an advantage over brute realism, it is positively false. For the theory claims that *all* moral obligations are caused by God. If there needs to be a prior moral obligation to enable God to cause obligations, then God does not cause all of them. This is the Prior Obligation objection.

I do not think that the Prior Obligation objection succeeds. It takes an important lesson from the impossibility of creation from normative ex nihilo but makes an unjustified inference from that lesson. While it is true that not just any prior normative fact will do, there is no good reason to believe that acceptable prior normative facts are limited to moral obligations. Hence why I do not rely on this objection to defeat causal divine command theory. It is worth explicating what this lesson is and where the Prior Obligation objection goes wrong. In doing so we can uncover what limits there are on acceptable prior normative facts and thus see whether any of them are reducible to God.

Moral obligations can only be brought about with the right kind of power – *normative* power.

But normative power comes in different degrees. Different normative facts provide different strengths of reason for action, e.g. many of us think that a prudential reason to ϕ is not as strong a moral reason to ϕ. And there can be unintelligibility not only when one conceives of a moral obligation coming to be through causal power but also through normative power of *insufficient strength.* Consider again an analogy with physical objects and causal power. It is inexplicable that a billiard ball should move across the table without any prior causal force impinging upon it. But the following is also inexplicable, although less glaringly so.

Suppose that one billiard ball very slowly makes its way across the table and gently taps another billiard ball. But then upon being tapped, this second ball races across the table and smashes into a third ball with such force that it is disintegrated. It is inexplicable as to where the massive causal force of the second ball came from. It was only invested with the tiny causal force of the first, slow-moving ball. If the first ball had also been moving very fast, the movement and impact of the second ball would make more sense. Likewise with normative force. Without the prior normative fact being sufficiently forceful, the emergence of a new, much stronger normative fact coming to be does not fully make sense. That is, we should endorse:

> **Prior Normative Force Principle:** normative facts can only come to be with the help of prior normative facts with equivalent normative force.

Granted it is not entirely clear how to 'measure' comparative normative force. But a rough test for seeing which of two kinds of normative fact is stronger, *A* or *B,* is to see which kind would prevail if they prescribed conflicting actions. So suppose that you have a social obligation to ϕ and a moral obligation to ~ϕ. Intuitively, the moral obligation prevails, and easily so. Moral obligations provide extremely strong reasons for action, arguably the very strongest. Social

obligations provide much weaker reasons. As such, we should expect some unintelligibility in moral obligations coming to be with the help of merely social obligations. And when we conceive of the standing obligation that invests the officer's commands with normative force to be merely an obligation of social etiquette, there *does* seem to be something unintelligible about the officer's commands producing moral obligations. Of course, perhaps there is a moral obligation to do what is socially obligated. But this just solves the problem by conceding the point; the prior normative fact must have the right force.

Now it is easy to see why somebody might infer from the Prior Normative Force Principle that when it comes to God's creation of moral obligation, the prior normative fact must specifically be a moral obligation. For one might think that the only normative fact with normative force equivalent to a moral obligation is another moral obligation. But this is, at the very least, not a strict entailment of the Prior Normative Force Principle. Consider that different physical objects or events can have equivalent causal force. My throwing a tennis ball and my throwing a rock may both have enough normative force to smash a window. There is, then, the logical possibility that there are normative facts distinct from moral obligations that nonetheless carry force equivalent to moral obligations. And as we shall see below, normative facts about *moral rights* are a promising candidate.

But first, although I think the inference from the Prior Normative Force Principle is the main source of the intuition that a prior moral obligation must be present, there are other arguments for that conclusion – arguments that also fail. Wolterstorff, supposing that the divine command theorist will want to hold that the prior normative fact is an obligation of some type, even if not a moral one, argues that this obligation *must* be a moral one. He notes that if God has commanded a particular action, then the following obligations are in play:

> (2) the standing general obligation to obey God by performing such actions as God may command me to perform, and
> (3) the generated specific obligation to obey God by performing the action A that God commanded me to perform. (2008, pp.278-9)

The divine command theorist will want to deny that these obligations are both moral. But Wolterstorff's reason for thinking that they must be both moral is that there is no way to violate the general obligation without violating the specific one, or some other specific one. On divine command theory, all of the specific obligations are *moral* obligations. And so, all violations of the general obligation involve violations of a moral obligation. But Wolterstorff holds that it is therefore impossible for the general obligation to not also be a moral obligation. It seems that Wolterstorff is committed to the following principle:

> **Violation Principle** If the only way to violate an obligation *o* is to violate a moral obligation, then *o* is a moral obligation.

Now, Wolterstorff notes that the Violation Principle does not commit one to the claim that *o* is *only* a moral obligation. It could be *both* a moral obligation and a relational, or socio-linguistic obligation, or some other obligation. The point is that whatever other sort of obligation *o* is, it is at least moral.

However, the Violation Principle is subject to counter-examples. Many have thought that it is a moral obligation to obey all of your legal obligations. Clearly, the only way to violate this moral obligation is to violate a legal obligation. Therefore, according to the Violation Principle, this implies that the moral obligation to obey all of your legal obligations is itself also a legal obligation. But that does not seem right. If it were a legal obligation, it would have to be

recognised in the legal system. A legal system, however, need not include any second-order law according to which it is illegal to break first-order laws. And so, the moral obligation to obey the law of a legal system need not itself be a legal obligation, even though the only way to break that moral obligation is to break a legal obligation.

Wolterstorff has another related but potentially distinct argument for the Prior Obligation Principle. He says; "I do not see how my obligation to perform this instance of the type can be a *moral* obligation when my obligation to perform *any* instance of the type that becomes possible for me is *not* a moral obligation (2008, p.280)." This thought seems an expression of the following principle:

Type Principle If I have a general obligation to do acts of type *x*, and I am morally obligated to do this token of *x*, then my general obligation to do acts of type *x* is a moral obligation.

This principle can also be shown false. Suppose it is *not* the case that we have a moral obligation to obey all our legal obligations. And suppose that our legal obligations do, this time, include a second-order law to obey first-order laws. I therefore have a general obligation to do acts of the type *obey legal obligations.* Now suppose that one of my legal obligations is to not steal from others. And suppose that it is also a moral obligation not to steal from others. In this situation, I am morally obligated to do a token of the type *obey legal obligations.* Yet by stipulation of the case, my general obligation to do acts of that type is *not* a moral obligation.

So far, then, we have seen that there is no way to motivate the Prior Obligation objection from more basic principles. Divine command theory must respect the Prior Normative Force

Principle, but this principle does not demand – not as a matter of logical entailment anyway – that the prior normative fact be a moral obligation. To argue that it must be, one must show that other candidate normative facts do not have the requires strength. The next task, then, is to see what normative facts other than moral obligations divine command theorists can appeal to. A few have been proposed by divine command theorists as part of their responses to the Prior Obligation objection. I argue that the appeal to *moral rights* is promising. We can start, however, by weeding out a couple proposals that are not promising.

Evans' first reply to the standard Prior Obligation objection is that:

> The defender of a DCT might respond by arguing that these "prior obligations" to obey God are not actual obligations but just hypothetical ones: To say that we have a prior obligation to obey God is just another way of saying that we are obligated to obey God *if* God issues commands, but there is no actual obligation until a command is issued. The conditional proposition itself can be understood as simply spelling out the meaning of the claim that God has moral authority. (Evans, 2013, p.99)

According to Evans, in addition to the standard kind of moral obligation we focus on, there could be *hypothetical* moral obligations. Thus, God's commands bring about all actual moral obligations, but there are prior, hypothetical moral obligations present. It is not clear what exactly a hypothetical obligation is. Evans seems to say that hypothetical obligations are merely true conditionals. To have a prior hypothetical obligation to obey God is just for it to be true that *if God commands ϕ, then we are obligated to ϕ.*

The job of a prior obligation, though, is to make God's generation of moral obligations explicable. But mere conditionals do not offer any explanatory connection between the antecedent and the consequent. The conditional *if there is fire, then there is smoke* does not claim any explanatory connection between fire and smoke. Granted, in most contexts the

utterance of "if there is fire, then there is smoke" would imply an explanatory connection. But that is because of background assumptions assumed in those contexts. If Evans is relying on background assumptions in our context, these need spelling out. Absent such assumptions, the truth of the conditional *if God commands ϕ, then we are obligated to ϕ* does not entail any normative force that can explain why God's commands generate obligations.

Elsewhere, Evans (2013, p.64) gives a couple accounts of how principles about what is *morally good* could provide God's commands with the right normative force:

> Good and conscientious parents give their children a great deal: time, effort, material resources. Gratitude for such gifts is good, and gratitude appropriately expresses itself in respect and appreciation for the parents. However, God as Creator is the giver of every good that any humans have, including their very existence. It seems plausible then that the debt of gratitude humans owe to God vastly exceeds the gratitude they owe to parents or any human benefactor. In fact, the gratitude owed to God must always be far greater than that owed to any human person or even to all human persons to whom one owes gratitude. Since God is the creator and sustainer of those human benefactors, all the goods given by those human benefactors, including parents, are also gifts of God.
>
> A second possible way that divine authority could be justified is to appeal to the goodness of a relation to God. Aquinas, for example, holds that a relation with God is the highest good possible for a human person. The ultimate happiness of a human person is the knowledge of God that is described as the "beatific vision." Obviously if this is so, humans have very good reasons to do what is necessary to achieve such a relation.

So gratitude and relationship with God make it very good to, and give humans very good reasons to, obey God's commands. But the normative force of something being very morally good is still significantly short of the force of something being morally obligatory. Granted, it is much closer to that force than, say, the force of something socially obligatory. And so Evans' suggestions are more plausible than many others that could be offered. But they aren't plausible enough. There still seems to be something inexplicable about God's generation of moral

obligations against the background of only normative facts about moral goodness.[49]

Now to the more promising kind of prior normative fact that divine command theorists can appeal to. Rather than facts connecting properties of God with moral goodness, there could be facts connecting God's properties with *moral rights,* or *moral authority.* Evans suggests this very possibility:

> A third possible normative principle that might justify divine authority is the claim that God, by virtue of his creation of humans and the natural world, has a rightful claim to be the owner of that created world and everything in it, including human beings (2013, p.65).

Evans claims that God's property of being the creator of the world grants him some kind of moral right or authority over it. But one need not accept Evans' particular focus on God's property as creator as what grants God his moral authority. Perhaps other properties more plausibly grant moral authority.[50] The thing to focus on here is the invocation of moral authority itself. The language of moral authority and of moral rights is not identical to the language of moral obligations. After all, 'authority' and 'rights' are different words from 'obligations.' A divine command theorist might hope, then, that normative facts about authority or rights (I will use these interchangeably from now on) do not reduce, even in part, to facts about moral obligations.

Yet moral rights seem to have greater normative force than moral goodness does. To say that one has the moral right to ϕ seems stronger than saying that it is morally good to ϕ. And saying

[49] The same criticism can be made of Hare's (2015, p.18) suggestion that God's goodness entails that God is to be loved, and that love entails obedience to God's commands. God's goodness might make it appropriate to obey God or very good to obey God, but this falls short of the force of obligation.
[50] Granted, Murphy (2002) argues that no divine property is such that necessarily, it grants God authority over all creatures. But I will not step into this debate – I will allow that some divine property can get the job done.

that someone has the moral right to not be treated ϕ-wise seems stronger than saying that it is morally bad to treat someone ϕ-wise. Rights seem to carry the strength of normative force that obligations do. There is, then, the promise for a kind of normative fact distinct from moral obligations yet with the force of moral obligations.

Now there *is* tension in claiming that some normative fact is distinct from a moral obligation but carries the force of moral obligation. While *x* carrying the force of moral obligation does not strictly entail that *x* reduces, even in part, to a moral obligation, it certainly makes such a reduction tempting. We want our ontology to be as little bloated by irreducible entities as possible, and if rights have the same force as obligations, then one might wonder what rights are doing in one's ontology. Wolterstorff, however, defends rights as doing work that moral obligations alone do not (2008, p.264-284). So I will be happy to grant the distinction between rights and obligations.

That appealing to rights is the best move divine command theorist can make may explain why this move is so widespread. Quinn himself seems to implicitly grant the need for prior authority. Like Evans, he suggests that God's role as creator might grant him authority, but also considers that God's power and his lovingness might (1978, p.48). Craig (2010) also thinks God's authority is key: "in the case of moral obligations, these arise as a result of imperatives issued by a competent authority. And in virtue of being the Good, God is uniquely qualified to issues such commands as expressions of His nature."

The remaining issue is whether facts about rights are reducible to God or not. I am not aware of any proposed theistic reductions of rights. Theists have proposed reducing moral goodness to God or his nature, and moral obligatoriness with his commands. But moral authority has not

been given reductive treatment. The most natural approach would be to take the relevant normative principle, like *creating the universe makes one a moral authority over it* and then interpret the making relation as a reductive relation. Thus, God's property of being a moral authority would reduce to his property of being creator of the universe (or whatever other property one's principle targets).

There is, however, a problem with a causal divine command theorist taking this approach. Strictly speaking, causal divine command theory is compatible with reductive views about moral properties distinct from the morally obligatory. Theoretically, one could claim that moral obligations are caused by God but say that moral goodness or moral rights reduces to God. One would need to be careful, however, that the reason why one reduces a moral fact to God does not undercut one's reasons for appealing to *causation* when it comes to moral obligations. For instance, if the making relation is a reductive one, and the causal divine command theorist endorses the normative principle *x being commanded by God makes x obligatory,* then this will entail that the morally obligatory reduces to being commanded by God. One needs a non-ad hoc reason to reduce one moral fact but not another. No such reason seems available. Causal divine command theorists should thus hold that normative facts about rights do not reduce to God. They should agree with premise (3).

§7.6 BRUTE RIGHTS

I argued above that voluntarists must posit an irreducible normative fact about rights. An irreducible fact is not the same as a brute fact. After all, on traditional monotheism, the physical universe is irreducible but it is still caused by God. I argue in this section that voluntarists must go further and posit a normative fact about rights that is both irreducible *and* brute. The

consequence of this, I argue, is that voluntarism has no explanatory advantage over brute realism.

One might think that God could be the cause of the very normative fact that grants him the right to issue obligations through commands. Perhaps God creates his own rights. The obvious problem with this is that if the normative fact about moral authority is caused by God, then God brings about a new normative fact. But the Prior Normative Force principle requires that there be some prior fact with the requisite normative force for that to happen. There would need to be some normative fact strong enough to grant God the power to create normative facts about moral authority. This normative fact would have to be neither a fact about rights nor moral obligations, since it is involved in the explanation of those.

Other candidates are not easily forthcoming. Further, even if one could identify another normative fact, this fact would either be brute or caused. If caused, then the regress continues. One has to find yet another normative fact with the right force. If brute, then one needs to explain why there is an advantage to positing *this* normative fact as brute rather than just positing the fact about rights as brute. The explanation would have to appeal to, say, the greater simplicity of this normative fact. But it is very hard to see what other normative fact voluntarists could appeal to that have the requisite strength and superior simplicity.

In addition, there needs to be some reason why it is that God can cause there to be some normative fact with the power to grant God moral authority, yet God does not possess the power to directly grant himself moral authority. Why must God go via the normative fact? In most cases, God himself has the powers that he bestows on created entities. For instance, God creates creatures that can move physical objects around. God has that power too. Now, I do not

claim that God must always directly possess the powers that he bestows on creatures. Perhaps, say, there are some phenomenal states, like feeling tempted, that creatures have the power to experience, but God does not – at least not directly. God might have to rely on created entities to experience those phenomenal states, by becoming incarnate in some physical body. What we need, though, is some principled reason why God does not have certain powers directly. The most obvious reason is that direct possession of certain powers entails imperfection in God. If God cannot experience temptation, the explanation is that experiencing temptation requires a certain weakness. But there does not seem to be any reason why an inability to directly grant himself moral authority would be an imperfection.

For the above reasons, it is implausible that God is the cause of the normative fact about moral authority. Divine command theorists should hold that this fact is brute. Indeed, in a recent defence of divine command theory, Choo explicitly takes this line:

> I admit that I think it is a brute fact that if a being with moral-authority-granting-properties commands us to *p,* then we have a moral obligation to *p.* This seems like an adequate explanatory stopping point. Even if I'm wrong and there is a further explanation, it seems right to think that the explanatory chain has to stop somewhere and so there will be a brute fact somewhere along the line. (2019, p.383)

At this stage, one might think that causal voluntarism as Quinn formulated it must be false. How can God cause there to be moral obligations with totality and exclusivity etc. when God must be helped by prior a brute, irreducible normative? It is certainly hard to see how God can, in any felicitous sense, be held to cause moral obligations with totality and exclusivity. But Quinn's view could just be modified to drop claims of totality and exclusivity. The claim that causal divine command theory has an explanatory advantage over natural non-reductionism never depended on those claims.

Granting an irreducible and brute normative fact about rights does not immediately foreclose the potential for explanatory advantage. As Choo explains, divine command theory "does not posit any brute moral obligations but offers a deeper explanation for all our moral obligations (2019, p.382)." That is, whereas Wielenberg posits multiple brute obligations, divine command theory need not posit any brute obligations at all. (This assumes that Wielenberg is committed to multiple fundamental obligations, but this is an assumption I am granting). Voluntarism can just posit one brute fact about moral *authority*, which, in conjunction with God's commands, explains all moral obligations. This promises a deeper and more elegant stopping point. It remains, then, a substantial task to show that causal divine command theory fails to have an explanatory advantage. It is the task I now turn to.

It is not clear that Wielenberg cannot modify his account to also make some fact about *rights* the one brute fact which then explains all moral obligations. The voluntarist's brute fact is:

> if a being with moral-authority-granting-properties commands us to ϕ, then we have a moral obligation to ϕ.

Wielenberg's brute fact could be:

> if a being has moral-right-granting-properties, then we have a moral obligation to treat that being *x*-wise.

Granted, although I have until now used talk of moral rights and moral authority interchangeably, the notion of moral authority in the voluntarist's brute fact differs from the

notion of moral right in Wielenberg's brute fact. The creatures that have rights in Wielenberg's brute fact cannot simply bring moral obligations into being by command – at least not without being in some very special circumstances. Nonetheless, rights and authority are plausibly connected notions. Wolterstorff treats this kind of authority as just the *right* to command; it is to have "the permission-right to do so and the claim-right to be free to do so (2008, p.269)". This is a different species of rights than, say, human rights, but they are all rights nonetheless. So if it is plausible to treat moral authority as not reducible, even in part, to moral obligations, it is plausible to give the same treatment to the moral rights of creatures.

Both views, then, can posit a single brute fact about rights that explains all moral obligations. If divine command theory can retain an advantage, it must be located elsewhere. One advantage it might be thought to retain is that the single brute fact about rights that it posits is *simpler* than the one Wielenberg's theory posits. Wielenberg's brute fact is that *if a being has moral-right-granting-properties, then we have a moral obligation to treat that being x-wise.* But what needs specifying is what exactly it is to treat a being *x*-wise. One might think that what is actually prescribed will be quite complicated. Indeed, it will end up merely specifying the multiple obligations that we have to the relevant beings. Obligations like *do not steal from them, do not harm them, do not violate their privacy,* and so on. Whereas, by contrast, the single brute fact that voluntarists posit is very simple. It just states that what we are obligated to do whatever God commands.

There is, however, reason to doubt that the brute fact about God's authority will really be as simple as that. For if it is, then God has the moral right to command anything whatsoever, no matter how horrendous. God would have the moral right to command, for instance, that a mother torture her baby, molest it, and then cut it up into pieces and eat it. But it is very hard

to believe that God has that right. My claim is not that what God has the moral right to command is as limited as what humans have the moral right to command. My claim is only that there are *some* limits to what God has the moral right to command. Thus, the brute fact that divine command theorists will have to posit will really be something like *if a being with moral-authority-granting-properties commands us to p,* ***and the command meets conditions c,*** *then we have a moral obligation to p.* And the divine command theorist needs to give us reason to believe that those conditions can be spelt out simply. Until they provide such reason, again, there seems to be no advantage for divine command theory.

One might object that it is not actually so implausible that God has a moral right to command anything whatsoever. If we find that thought violating our moral intuitions, that is because we are imagining that God actually *exercises* his right to command horrendous things. But, the objection continues, there is a difference between having a right and exercising it. God may have the right to command a mother to mutilate her baby, but God would never exercise that right because of his essential moral goodness. [51] Thus, the brute fact about God's authority is indeed very simple – it carries no restrictions. But so long as we keep in mind that God would never actually command some actions that he has the authority to command, nothing about this is morally offensive.

What should we make of this reply? I simply protest that I am not making the mistake the reply aims to correct. I find the idea that God has the moral authority to command that a mother mutilate her baby violates my moral intuitions, *even if* God would never exercise that right. There are just some things that even God's role as creator, or supreme benefactor, do not give

[51] While all moral arguers hold that God is good of necessity, not all theists agree. See Reichenbach (2014) for a recent defence of God's goodness as merely contingent.

him the right to command. And I do not think that my moral intuitions on this matter are idiosyncratic. Granted, I also do not think that my moral intuitions are universal. There surely are theists who are happy to assert that God has the moral right to command anything whatsoever. I do not claim to have an argument to persuade these theists otherwise. But one should remember the dialectical context. One is trying to offer a moral argument that gives an atheist a reason to become a theist. I submit that few atheists will be inclined to agree that God would have the moral authority to command anything whatsoever, no matter how horrible. In which case, one cannot rely on that claim in a moral argument for theism.

Now divine command theorists might concede that there will have to be conditions on what God has the right to command. But they might claim that these conditions are still simpler than those that an atheist like Wielenberg will need to rely on. Indeed, the condition could just be that the issuing of the command be compatible with moral perfection. The normative fact could be:

> If being b_1 with moral-authority-granting properties commands ϕ in context c, and a possible being b_2 with moral perfection and the same power as b_1 could have also commanded ϕ in c, then ϕ is obligatory.

What this move does is to take the appeal to moral perfection but build it *into* the principle about moral authority. So the claim is not that God has the right to command anything at all, and his moral perfection acts as an external constraint on his exercise of that right. Rather, internal to the right is the restriction that it only gives him authority to command in a way that is consistent with moral perfection.

Although this moral authority principle is more complicated than one that gives God carte blanch authority to command anything at all, it might still be thought simpler than the rights-principle that an atheist could appeal to. That is because acting consistently with moral perfection remains relatively simple. By contrast, the atheist's alternative still requires a complex shopping list of different stipulations. There may be nothing that unifies in a simple way all the kinds of treatment that persons have the right to receive. And so one may be left with an inelegant list like *do not steal from them, do not harm them, do not violate their privacy,* and so on. Again, the claim is that atheists are stuck with an inelegant multiplicity at explanatory rock-bottom.

Note, however, that the divine command theorist's move merely farms out any potential complexity in the principle about moral authority into principles about moral goodness. The modified moral authority principle invokes a normative notion, that of *moral perfection.* The moral authority principle does not actually make determinate what a being with authority-granting-properties has the right to command. It must be coupled with normative principles about what moral perfection demands. Now this is not yet to identify a problem. The voluntarist is not positing any new, further normative principles. Moral perfection is the maximum degree of moral goodness. So principles about moral perfection will just be within the wider family of principles about moral goodness. And very plausibly, theists and atheists are both already committed to such principles. The divine command theorist is just putting those principles to work in a moral authority principle.

The problem is that atheists like Wielenberg could also farm out the complexity in the principle about person's rights into other normative principles. They could claim, for instance, that the brute fact about rights that explains all moral obligations is:

> If a being has moral-right-granting-properties, then we have a moral obligation to respect that being's *dignity.*

I take it that *dignity* is a kind of value that merits a certain response: that of respect. And respect consists of a wide variety of behaviours and attitudes. There are, then, normative principles linking dignity with a wide variety of behaviours and attitudes. The atheist can simply put these principles to work in detailing the variety of behaviours and attitudes that persons have a right to receive. A right that makes it obligatory for us to offer those behaviours and attitudes. Note that just like principles about moral goodness, both theists and atheists are already committed to normative principles about dignity. So once again, causal divine command theory seems on a par with brute realism. No advantage has been gained.

There is, however, one further way that voluntarists can try to retain an advantage. Suppose that both moral goodness and respect are degreed properties. Someone can be more morally good than someone else and someone can offer more respect than someone else. Now as Swinburne (2004, p.55) argues, intuitively, there is something simpler about a property instantiated to the maximum degree than to some arbitrary finite degree. So the advantage could be this: the principle about God's moral authority appeals to a property instantiated to the maximum degree – moral *perfection.* By contrast, the principle about the dignity due to persons will need to appeal to some arbitrary finite amount of respect. Because of this contrast, the theist's principle is simpler.

But Wielenberg can claim that the amount of respect obligated may be finite but not *arbitrary.* The degree of respect that we are obligated to give to persons is the minimal positive amount.

That is, our degree of respect cannot be so low that it counts as *disrespect.* So long as we cross over into positive respect, we meet our obligations. Any further degree of respect is supererogatory. There is nothing arbitrary about this, like there would be if we are obligated to give 138 units of positive respect. Divine command theorists might object that it is controversial as to whether there *is* a minimum degree of positive respect. Perhaps positive respect can come in ever smaller amounts, getting closer and closer to 0 units of positive respect. But it is highly doubtful that if respect comes in degrees, the degrees are so fine-grained as to be infinitely divisible.

Further, it is controversial as to whether an omnipotent and omniscient being could have moral goodness to a *maximum* degree. Some critics of theism, like Rowe (2004) and Oppy (1995, p.166), have argued that for there to be such a maximally good being, there would have to be at least one possible world that is maximal in value. And there are good reasons to believe that there are no such worlds. As Swinburne (1994, p.135) says, any world we can conceive of can always be conceivably improved by featuring one more valuable state of affairs, such as the existence of one more conscious creature. Most theists argue in response that God need not create an unsurpassable world, but only a world that is *good enough.* But what degree of value is good enough? If it is some arbitrary finite amount, then theists are committed to a normative fact that is as inelegant and complex as they accuse the atheist's rights principle as being. Now perhaps instead, God need only a create a world that is minimally good – overall more good than bad. But if there are concerns about whether there is a minimal degree of positive respect, there are plausibly concerns about whether there is a minimal degree of positive value. So it is not at all clear that divine command theorists can boast an advantage here.

Indeed, there seems to be no way to show that causal voluntarism has an advantage over brute

realism. At least, no successful way of showing that has been produced by divine command theorists at present. And attempts to anticipate how they might adapt their arguments to maintain an advantage have no more success. I conclude, then, that the burden of proof is on the voluntarist to come up with some new way of showing that their theory has an advantage. Until they do that, we can conclude that premise (4) is true. And with that, the argument as a whole is sound.

Causal voluntarism has no advantage when it comes to explanatory depth. And so the major argument for preferring causal divine command theory fails. Moral arguers have not been able to motivate P2 by appeal to causal voluntarism.

8

Conclusion

Let us take stock of where we have got to. We have been assessing metaphysical moral arguments for theism, which focus on the existence and objectivity of moral properties. These arguments assume that a good independent case can be made for the existence and objectivity of moral properties. They then argue for the following:

> P2. It is plausible that theism accommodates objective moral properties and more plausible that it does so than that naturalism does so.

That is, moral arguments claim that theism offers better answers as to whether moral properties are reducible, what they might reduce to, whether their instantiation is explainable, and what might do the explaining, than naturalism. It is assumed, however, that theism and naturalism are the only two live options. Hence, these arguments claim that theism offers the very best answers to these questions. There is thus said to be a route from the objectivity of morality to theism.

Most moral arguers motivate P2 by appeal to restricted voluntarism. Restricted voluntarism claims that what is morally right and wrong depends on God's commands. And what is morally good or bad depends on God himself or his nature. Robert Adams claims that the dependence in view is reduction. Philip Quinn claims that it is causation, at least when it comes to the dependence of right and wrong on God's commands. These are the two main options that moral arguers appeal to for an account of how theism accommodates objective moral properties.

To establish P2, moral arguers must show that restricted voluntarism is more plausible than every naturalistic competitor. Obviously, this requires showing that restricted voluntarism is more plausible than some *particular* naturalistic competitors. I argued that there are particular competitors that restricted voluntarism has no advantage over. Or rather, I argued that moral arguers have not successfully shown voluntarism to have the advantage. Of course, I did not assess every possible argument for preferring voluntarism. Just the major arguments that are prominent in the writings of moral arguers.

I started by considering the Conceptual Gap Argument for preferring reductive voluntarism to natural reductionism – the view that moral properties reduce to natural properties. This argument claims that there is a conceptual constraint on reduction. Moral properties reduce to *x*-type properties only if moral concepts have a certain kind of conceptual connection *c* to *x*-type concepts. The claim is then that reductive voluntarism satisfies this constraint whereas no form of natural reductionism does. I argued that there is no account of the conceptual connection *c* that both makes it plausible that *c* is a constraint on reduction *and* allows that moral concepts have *c* with theistic concepts, but not natural concepts.

Next I considered the Divergence Argument for preferring reductive voluntarism to natural reductionism. This argument claims that there is a tracking constraint on reduction. A moral property *m* reduces to property *f* (or properties *fs*) only if there is some relevant set of properties *s* such that necessarily all and only actions with *m* instantiate all the properties in *s.* The claim is then that reductive voluntarism satisfies this constraint whereas no form of natural reductionism does. I argued that attacking natural reductionism on this score burdens moral arguers with significant commitments in first-order normative ethics. Further, if moral arguers

shoulder that burden, they will be unable to defend voluntarism against an important objection – the Counterfactual Argument.

Finally, I considered the Explanatory Depth argument for preferring causal voluntarism to brute realism – the view that moral properties do not reduce to any properties at all and that moral laws are without explanation. This argument claims that positing God's will as the causal explanation of moral laws is rationally preferable to positing moral laws as merely brute. It allows for a simpler explanatory rock bottom. I argued, however, that claims of superior simplicity cannot be sustained. For causal voluntarism must posit at least one brute moral fact – plausibly one about God's moral authority. But brute realism can also appeal to just one brute fact – one about the respect due to persons. And attempts to show that voluntarism's one brute fact has greater internal simplicity than brute realism's fact, fail.

Overall, then, the three major argument for preferring restricted voluntarism to some particular naturalistic competitors all fail. It has thus not been shown that voluntarism has an advantage over *all* naturalistic competitors. Moral arguers have yet to successfully motivate P2. As such, we do not yet have reason to believe that the existence and objectivity of moral properties points to the existence of God. This therefore completes my negative assessment of arguments from moral objectivity for theism.

Of course, my assessment of the present state of the arguments does not rule out other, more successful arguments for P2 being put forward. In particular, theists should devote more attention to the under-developed arguments I mentioned in §3.3. I noted that according to some theists, even if moral properties are brute and irreducible, this is more plausible on theism than on naturalism. This option for accommodating morality – appealing to brute and irreducible

moral properties – was the one remaining option not covered by voluntarism that still had promise for a theistic moral argument. To close, I will offer some suggestive thoughts as to how these arguments could be developed.

Let us first consider them in their undeveloped form. Evans says:

> Mavrodes points out that a world in which such facts are "brute facts" does not seem very much like a Russellian world. One would think that, in a Russellian world, basically a naturalistic world in which everything reduces to the basic particles of physics, moral facts would not be "deep." Rather, a world in which there are ultimate moral facts seems more like Plato's view of the world than Russell's. And the world as seen by Platonists, though perhaps not necessarily theistic, has usually seemed to be "congenial (especially congenial when compared to some other philosophical views) to a religious understanding of the world." In the Platonic view of reality, whatever moral obligations humans face reflect the basic character of reality itself, and that in turn suggests that the ultimate truth about that universe is not that it is simply "accidental collocations of atoms" (2013, p.115).

And Wainwright says:

> The general point is this: naturalists believe that the space-time world is all there is; the story of the world is ultimately a story of "atoms, and time and space and economics and politics." But "values and obligations cannot be deep in such a world. They have a grip only upon surface phenomena, probably only upon man. What is deep" in a world of this sort "must be such things as matter and energy, or perhaps natural law, chance, or chaos." In a Platonic or theistic world, on the other hand, the Good is "much more fundamental to reality than are the atoms," and those who align themselves with the Good align themselves "with what is deepest and most basic in existence." The existence of objective values, or of necessary connections between natural properties and the value properties which supervene upon them, are surely less surprising in a world of this sort than in one in which what is deepest is matter, energy, natural law, or chance. As a result, religious views like theism or Platonism can illuminate the existence of objective value in a way in which naturalism cannot. That they can provides a reason for embracing views of this sort (2005, p.67).

The puzzle is to understand how theism could have an advantage over naturalism even when God is not doing any explaining or serving as a reduction base. A clue lies in Wainwright's

statement that moral properties would be "less surprising" in a theistic world than in a naturalistic world. Language of what is "surprising" or "unexpected" often expresses a *probabilistic* claim. Perhaps, then, there is potential to show that in a purely epistemic sense, brute and irreducible moral properties are more probable on theism than on naturalism.

The easiest way to run such an argument would be to opt for a purely subjective interpretation of epistemic probability. So long as, just as a matter of psychological fact, we place a higher credence in the likelihood of brute and irreducible moral properties on theism than on naturalism, the existence of brute and irreducible moral properties turns out to confirm theism. Other arguments for God, such as the fine-tuning argument (Monton, 2006), have already been interpreted according to a subjective account of epistemic probability. This would just apply the same treatment to moral arguments.

Of course, not everyone is happy with subjective interpretations of epistemic probability. Here, then, is another option. Those features of reality that are brute and irreducible are the most foundational, deepest parts of reality. Everything else that exists must be explained by appeal to them. According to Evans, there is something strange about moral facts or properties being at this deepest level of reality if naturalism is true. What exactly is strange about it? We might get a clue from Rosen, not himself a moral arguer, who also thinks there is something strange here:

> the view entails that unlike other facts about large complex particulars like people and their actions, the basic normative facts about such things are not metaphysically grounded from below. If we could survey the complete metaphysical grounding graph of the universe we would find at the bottom, along with (say) the quantum state of the universe and the basic principles of mathematics, facts about the normative features of particular human actions. I find that hard to swallow (2017, p.138).

The key thought I take from this passage is that on this picture, naturalism is left with two deep, fundamental sorts of facts that are very different. And not just different by virtue of one sort being physical facts and the other moral. We can think of it like this: moral facts are, in a certain sense, anthropocentric. They concern the actions of human beings (and any other moral agents that might exist). By contrast, facts about microphysical laws are thoroughly unconcerned with human beings. They may have led to the emergence of human beings, but this was quite unanticipated by the laws themselves.

On theism, things are different. We have God and moral facts at the deepest level of reality, not microphysical laws and moral facts. And God, like moral facts, is in some sense concerned with human beings, or agents more generally. Indeed, perhaps in a very directly literal sense – God *cares* for created agents. And even prior to creating, plausibly God is by nature the sort of being who has desires to create agents and is himself an agent concerned with his own agency and how to exercise it. So there is a unity, on theism, between God and moral facts. Both are agent-centred features of fundamental reality. But on naturalism, there is a disunity between the two fundamental features of reality. Moral facts are agent-centred but physical facts are not.

Let us call the thesis that God and moral facts are the two fundamental features of reality the *M*+*G* thesis. And let us call the thesis that physical facts and moral facts are the two fundamental features of reality the *M*+*P* thesis. Potentially, one could argue that *M*+*G* is a simpler hypothesis than *M*+*P*. For at a very abstract level, it appeals to only one type of fundamental fact – agent-centred facts – rather than two. And if *M*+*G* is simpler, it arguably has greater intrinsic probability, all else being equal. Thus, if we have reason to believe that

moral facts are brute and irreducible, we have reason to refer *M*+*G* to *M*+*P*. That translates as a reason to prefer theism to naturalism (all else being equal).

This argument is only suggestive, but something along its lines is worth further research.

References

Adams, R., 1973. A Modified Divine Command Theory of Ethical Wrongness. In: R. Adams, ed. 1987. *The Virtue of Faith and Other Essays in Philosophical Theology.* New York: Oxford University.

Adams, R., 1979a. Moral Arguments for Theistic Belief. In: R. Adams, ed. 1987. *The Virtue of Faith and Other Essays in Philosophical Theology.* New York: Oxford University.

Adams, R., 1979b. Divine Command Metaethics Modified Again. *The Journal of Religious Ethics,* 7(1), pp.66-79.

Adams, R., 1999. *Finite and Infinite Goods: A Framework for Ethics.* Oxford: Oxford University.

Alston, W., 2002. What Euthyphro Should Have Said. In: W.L. Craig, ed. 2002. *Philosophy of Religion: A Reader and Guide*. Edinburgh: Edinburgh University. pp.283-298.

Audi, R., 2011. *Rationality and Religious Commitment.* Oxford: Oxford University.

Baggett, D., Walls, J., Copan, P. & Craig, W.L., 2011. Book Symposium: Theistic Foundations for Morality by David Baggett and Jerry Walls. In: Evangelical Theological Society, *63rd Annual Meeting: No Other Name.* San Francisco, California, 16-18 November, 2011.

Baggett, D. & Walls, J., 2011. *Good God: The Theistic Foundations of Morality.* Oxford: Oxford University.

Baggett, D. & Walls, J., 2016. *God and Cosmos: Moral Truth and Human Meaning.* New York: Oxford University.

Blackburn, S., 1993. *Essays in Quasi-Realism.* Oxford: Oxford University.

Blackburn, S., 1998. *Ruling Passions: A Theory of Practical Reasoning.* Oxford: Clarendon.

Bedke, M., 2012. Against Normative Naturalism. *Australasian Journal of Philosophy,* 90(1), pp.111-129.

Boyd, R., 1988. How to be a Moral Realist. In: P.K. Moser & J.D. Trout, eds. 1995. *Contemporary Materialism: A Reader.* Abingdon: Routledge. Ch.15.

Brandon-Mitchell, D. 2008. Naturalistic Analysis and the A Priori. In: D. Brandoon-Mitchell & R. Nola, eds. 2008. *Conceptual Analysis and Philosophical Naturalism.* Oxford: Oxford University. Cambridge, MA: MIT. Ch.2.

Brink, D.O., 1989. *Moral Realism and the Foundations of Ethics.* New York: Cambridge University.

Chalmers, D., 1996. *The Conscious Mind: In Search of a Fundamental Theory.* Oxford: Oxford University.

Chalmers, D., 2010. *The Character of Consciousness.* Oxford: Oxford University.

Choo, F., 2019. The Prior Obligations Objection to Theological Stateism. *Faith and Philosophy,* 36(3), pp.372-384.

Copp, D., 2012. Normativity and Reasons: Five Arguments from Parfit Against Normative Naturalism. In: S. Nuccetelli and G. Seay, eds. 2012. *Ethical Naturalism: Current Debates.* New York: Cambridge University.

Craig, W.L., 2008. *Reasonable Faith.* 3rd ed. Wheaton: Crossway.

Craig, W.L., 2009. This Most Gruesome of Guests. In: R.K. Garcia and N.L. King, eds. 2009. *Is Goodness without God Good Enough?* Plymouth: Lowman & Littlefield. Ch.9.

Craig, W.L. & Sinclair, J.D., 2009. The Kalam Cosmological Argument. In: W.L. Craig & J.P. Moreland, eds. 2009. *The Blackwell Companion to Natural Theology.* Oxford: Blackwell. Ch.3.

Craig, W.L., 2010. *Does Theistic Ethics Derive an "Ought" from an "Is"?* [online] Available at: <https://www.reasonablefaith.org/writings/question-answer/does-theistic-ethics-derive-an-ought-from-an-is/> [Accessed 2 October 2020].

Craig, W.L., 2011. *Sam Harris on Objective Moral Values and Duties.* [online] Available at: <https://www.reasonablefaith.org/writings/question-answer/sam-harris-on-objective-moral-values-and-duties/> [Accessed 2 October 2020].

Craig, W.L., 2018. Erik Wielenberg's Metaphysics of Morals. *Philosophia Christi,* 20(2), pp.333-338.

Craig, W.L. & Wielenberg, E., 2020. *A Debate on God and Morality: What is the Best Account of Objective Moral Values and Duties?* A.L. Johnson, ed. New York: Routledge.

Danaher, J., 2019. In Defence of the Epistemological Objection to Divine Command Theory. *Sophia,* 58(3), pp.381-400.

Davidson, D., 1970. Mental Events. In: D. Davidson, ed. 1980. *Essays on Actions and Events.* Oxford: Clarendon Press. Ch.11.

Enoch, D., 2011. *Taking Morality Seriously: A Defence of Robust Realism*. Oxford: Oxford University.

Evans, C.S., 2010. *Natural Signs and Knowledge of God.* Oxford: Oxford University.

Evans, C.S., 2013. *God and Moral Obligation.* Oxford: Oxford University.

Evans, S.C., 2018. *Moral Arguments for the Existence of God.* [online] Available at: <https://plato.stanford.edu/entries/moral-arguments-god/> [Accessed on 2 October].

Finlay, S., 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University.

Foot, P., 2001. *Natural Goodness.* Oxford: Oxford University.

Hare, J., 1997. *The Moral Gap.* New York: Oxford University.

Hare, J., 2015. *God's Command.* Oxford: Oxford University.

Harrison, G. K., 2018. *Normative Reasons and Theism*. Cham: Palgrave Macmillan.

Hume, D., 1739. *A Treatise of Human Nature*. D.F. Norton and M.J. Norton, eds. 2000. Oxford: Clarendon Press.

Jackson, F., 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis.* Oxford: Oxford University.

Kripke, S., 1980. *Naming and Necessity.* Oxford: Blackwell.

Kurtz, P. & Craig, W.L., 2009. The Kurtz/Craig Debate. In: R.K. Garcia & N.L. King, eds. 2009. *Is Goodness without God Good Enough?* Plymouth: Lowman & Littlefield. Ch.1.

Lambert, S., 2020. *Is Theism Compatible with Moral Error Theory?* Unpublished.

Laskowski, N., 2019. The Sense of Incredibility in Ethics. *Philosophical Studies,* 176(1), pp.93-115.

Leftow, B., Anderson, P.S. & Schellenberg, J.L., 2015. Whither Philosophy of Religion? *Religious Studies,* 51(Special Issue 3), pp.441-454.

Levine, J., 2001. *Purple Haze: The Puzzle of Consciousness.* New York: Oxford University.

Linville, M., 2009. The Moral Argument. In: W.L. Craig & J.P. Moreland, eds. 2009. *The Blackwell Companion to Natural Theology.* Oxford: Blackwell. Ch.3.

MacIntyre, A., 1969. Hume on 'is' and 'ought'. In: W.D. Hudson, ed. 1969. *The is-Ought Question: A Collection of Papers on the Central Problems in Moral Philosophy.* London: Macmillan. pp.35-50.

Mehta, N., 2017. Phenomenal, Normative, and Other Explanatory Gaps: A General Diagnosis. *Philosophy and Phenomenological Research,* 98(3), pp.567-591.

Miller, C., 2009. The Conditions of Moral Realism. *Journal of Philosophical Research,* 34, pp.123-155.

Monton, B., 2006. God, Fine-Tuning, and the Problem of Old Evidence. *The British Journal for the Philosophy of Science,* 57(2), pp.405-424.

Moreland, J. P. & Craig, W.L., 2003. *Philosophical Foundations for a Christian Worldview*. Downers Grove: InterVarsity.

Morriston, W., 2009. The Moral Obligations of Reasonable Non-Believers. *International Journal for Philosophy of Religion,* 65(1), pp.1-10.

Morriston, W., 2011. God and the Ontological Foundation of Morality. *Religious Studies*, 48(1), pp.15-34.

Moore, G. E., 1903. *Principia Ethica.* London: Cambridge University.

Murphy, M.C., 2002. *An Essay on Divine Authority.* New York: Cornell University.

Murphy, M.C., 2009. Theism, Atheism, and the Explanation of Moral Value. In: R.K. Garcia & N.L. King, eds. 2009. *Is Goodness without God Good Enough?* Plymouth: Lowman & Littlefield. Ch.6.

Murphy, M.C., 2011. *God and Moral Law: On the Theistic Explanation of Morality.* Oxford: Oxford University.

Murphy, M.C., 2012. Restricted Theological Voluntarism. *Philosophy Compass,* 7(10), pp.679-690.

Murphy, M.C., 2018. No Creaturely Intrinsic Value. *Philosophia Christi*, 20(2), pp.347-355.

Murphy, M.C., 2019. *Theological Voluntarism*. [online] Available at: <https://plato.stanford.edu/entries/voluntarism-theological/> [Accessed 7 October].

O'Connor, T., 2008. *Theism and Ultimate Explanation.* Oxford: Blackwell.

Oppy, G., 2018. *Naturalism and Religion: A Contemporary Philosophical Investigation.* Abingdon: Routledge.

Oppy, G., 1995. *Ontological Arguments and Belief in God.* New York: Cambridge University.

Parfit, D., 2013. *On What Matters: Volume Two.* S. Scheffler, ed. New York: Oxford University.

Peoples, G., 2011. The Epistemological Objection to Divine Command Ethics. *Philosophia Christi,* 13(2), pp.389-401.

Putnam. H., 1973. Meaning and Reference. *The Journal of Philosophy,* 70(19), pp. 699-711.

Quinn, P., 1978. *Divine Commands and Moral Requirements.* Oxford: Clarendon.

Quinn, P., 2006. Divine Command Ethics: A Causal Theory. In: C.B. Miler, ed. 2006. *Essays in the Philosophy of Religion*. pp.37-52.

Quinn, P., 1990. The Recent Revival of Divine Command Ethics. *Philosophy and Phenomenological Research,* 50(Supplement), pp.345-365.

Quinn, P., 2013. Divine Command Theory. In: H. LaFollette & I. Persson, eds. 2013. *The Blackwell Guide to Ethical Theory.* 2nd ed. Chichester: Blackwell. Ch.4.

Reichenbach, B., 2014. God and Goodness Revisited: A Case for Contingency. *Philosophical Christi,* 16(2), pp.319-338.

Ritchie, A., 2012. *From Morality to Metaphysics: the theistic implications of our ethical commitments.* Oxford: Oxford University.

Rosen, G., 2017. What is a Moral Law? In: R. Shafer-Landau, ed. 2017. *Oxford Studies in Metaethics 12.* Oxford: Oxford University.

Ross, W. D., 1930. *The Right and the Good.* Oxford: Oxford University.

Rowe, W. L., 2004. *Can God Be Free?* Oxford: Oxford University.

Street, S., 2006. A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies,* 127(1), pp.109-166.

Swinburne, R., 1994. *The Christian God.* Oxford: Clarendon.

Swinburne, R., 2004. *The Existence of God.* 2nd ed. Oxford: Oxford University.

The Bible: English Standard Version, 2002. London: HarperCollins.

Wainwright, W., 2005. *Religion and Morality.* Abingdon: Routledge.

Waghorn, N., 2015. Metz' Incoherence Objection: Some Epistemological Considerations. *Journal of Philosophy of Life,* 5(3), pp.150-168.

Wielenberg, E., 2009. In Defence of Non-Natural, Non-Theistic Moral Realism. *Faith and Philosophy,* 26(1), pp.23-41.

Wielenberg, E., 2014. *Robust Ethics.* Oxford: Oxford University.

Wielenberg, E., 2018a. Divine Command Theory and Psychopathy. *Religious Studies,* pp.1-16. http://dx.doi.org/10.1017/S0034412518000781.

Wielenberg, E., 2018b. Reply to Craig, Murphy, McNabb, and Johnson. *Philosophia Christi,* 20(2), pp.365-375.

Wolterstorff, N., 2008. *Justice: Rights and Wrongs.* Princeton: Princeton University.

Wynn, M., 1997. Trust-Relationships and the Moral Case for Religious Belief. *International Philosophical Quarterly,* 37(2), pp.179-188.

Yetter-Chappell, H. & Chappell, R., 2013. Mind-body meets metaethics: a moral concept strategy. *Philosophical Studies,* 165(3), pp.865-878.