

Objective Moral Values and Metaphysical Queerness

Tony Roy

Perhaps the most fundamental problem for a divine command theory is to say *how* commands, or objective facts more generally, constitute moral values. In the first chapter of his *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin Books, 1977), J. L. Mackie (whom we have already met) argues against the existence of objective moral values. He thinks objective moral values would be so “queer” (in the old-fashioned sense of “strange” or “bizarre”) that it is reasonable to conclude there aren’t any. In this paper, I develop one key strand of Mackie’s reasoning along with a reply. You will have to judge its success!

In asking whether moral values are *objective* or *subjective*, we ask whether, in a sense yet to be specified, moral values are *up to us* or not. If they are up to us, then moral values are subjective; if they are not up to us then moral values are objective. We will get more clear about the distinction in section one of this paper. Given the distinction, though, it is natural to wonder which is right: whether moral values are objective or subjective. The question is important, for answers go to the core of what *is* good, bad, right and wrong, and of *how* good, bad, right and wrong matter for the way we live.

Some think cultural differences show that moral values are subjective.¹ But the most forceful considerations against objectivism, and so for subjectivism, may be arguments which emphasize the “queerness” of objective moral values. Arguments which use queerness against objective moral values are embedded in much ordinary “educated” (anti-)moral thought; they occur forcefully in David Hume and, more recently, in the first chapter of J. L. Mackie’s, *Ethics: Inventing Right and Wrong*. Mackie develops the arguments based on queerness as part of a larger argument for “moral skepticism.” That larger argument goes something like this:

1. If there are “ordinary” moral truths, then there are objective moral values.
2. There are no objective moral values.

3. There are no “ordinary” moral truths.

The argument is valid in the sense that *if* the premises are true, then the conclusion is true as well. So its soundness depends on whether the premises are true. The first premise is a claim about “ordinary” moral practice. Of course, what is ordinary in one context may be extraordinary in another. So it is best to see the premise as a claim about certain, maybe traditional, approaches to moral value—*these* approaches require or presume objective values. The second premise is defended by the arguments from queerness. “If there were

¹ For discussion, see, e.g., J. Rachels, *The Elements of Moral Philosophy* (Philadelphia: Temple University Press, 1986), pp. 12-24.

objective values,” Mackie says, “they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe.”² This strangeness is a reason for thinking there are no objective moral values—like pink elephants and pixies, objective moral values would be so strange or bizarre, that we should think there aren’t any. Since the argument is valid, if its premises are true then its conclusion is true as well. We’ll develop motivations for accepting the premises in section two.

The discovery of a pink elephant, or a reasonable theory explaining how elephants could be pink, might cause us to revise a belief according to which there are no pink elephants. Similarly, the discovery of an objective moral value, or a theory explaining how objective values could be, might cause us to revise a belief according to which there are no objective moral values. In this way, Mackie’s reasoning is vulnerable to a positive explanation of supposedly queer features. In the third section, I attempt such a reply. I’ll develop an approach to objective moral values which, I claim, explains supposedly queer features. This account of objective moral values is particularly simple and perhaps mistaken. Still, if the simple approach explains queerness, there is room for more sophisticated (and correct) theories which do so as well. So, given the simple account, there is room for the suggestion that Mackie’s second premise is false. The first two sections lay groundwork for this positive theory.

I. The Question

As I say above, in asking whether moral values are objective or subjective, we ask whether, in a sense yet to be specified, moral values are *up to us* or not. If they are up to us, then moral values are *subjective*; if they are not up to us then moral values are *objective*. Mackie argues that there is a problem about the very idea of an objective moral value. To understand his argument, it is important that we understand the difference between objectivity and subjectivity in general, and this difference between objective and subjective moral values in particular. Unfortunately, the distinction is both slippery and controversial. In this section, I develop a version which should suffice for our purposes. I start with a general characterization, then move to some examples, and try to head off a few misunderstandings.

(A) General characterization. For a general characterization of the difference between objectivity and subjectivity, the notion of an *attitude* is required. For our purposes, an *attitude* is a thing or state with descriptive content or meaning—where, typically, this content or meaning is introduced by a ‘that’ clause. So, e.g., one may know *that* there are fish in the sea, believe *that* dogs fly, sense *that* someone loves you, experience *that* life is hard, demand *that* you get respect, command *that* your will be done, contract *that* you will pay some fee, hope *that* there is peace, desire *that* you receive a good grade, etc. Thus knowing, believing, sensing, experiencing, saying, demanding, commanding, contracting,

² J. L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin Books, 1977), p. 38.

hoping, desiring, etc., are all attitudes. In the first instance, it's *persons* who have attitudes. A person knows that there are fish in the sea, believes that dogs fly, etc. But, derivatively, attitudes may be embodied in groups, cultures, traditions, practices and the like. Americans believe *that* all persons have the right to life, liberty and the pursuit of happiness (this is so even though not every American believes it), it may be traditional *that* a first born child take up his or her father's profession, legislated *that* driving on the left side of the road is subject to punishment, etc. Notice: Not *every* mental state or sensation is an attitude. You may be happy that you understand the notion of an attitude. But perhaps there is also a simple sort of happiness which is not a happiness *that* anything. Thus the importance of content: An attitude's content gives it a "referential" or "representational" character—an "aboutness"—which differentiates it from other states. Your believing, saying, or being happy that there are fish in the sea includes or requires reference to the fish and sea. No such reference or "aboutness" characterizes "pure" happiness, a swim, or a tree.

Our primary distinction is between objective and subjective properties. Say a *thing* has a *property* under a certain *condition*. Thus we might say a ball is *round* insofar as every point on its surface is an equal distance from its center, and a skirt is *fashionable* insofar as it meets certain length, material, and cut characteristics. As we shall see, the former proper property is objective, and the latter subjective.

A property is *subjective* iff its condition is dependent on the attitudes of a person, group of persons, tradition, practice, or the like directed at ("about") the having of that condition. A property is *objective* iff it is not subjective.

Similarly, a relation is *subjective* iff its condition depends on attitudes about having that condition, and *objective* iff it does not. Derivatively, say a fact, sentence or claim is *subjective* (*objective*) iff it involves a subjective (*objective*) property or relation.³ Here is a simple test to distinguish the objective from the subjective: Imagine altering attitudes with respect to some condition; if altering the attitudes changes whether it is the condition under which a thing has a property, the property is subjective; if altering the attitudes leaves the condition under which a thing has the property unchanged, the property is objective. The test distinguishes objective from subjective properties just insofar as conditions that depend on attitudes change when attitudes change.

The question about whether moral values are objective or subjective, then, is the question about whether moral properties are objective or subjective. It is the question about whether conditions for being morally good, morally bad, morally right, or morally wrong do or do not depend on attitudes. In particular, on a *subjective* moral theory, every moral property is subjective, while according to an *objective* moral theory, at least some moral

³ It is sometimes convenient to reserve "subjective" for cases where the person, group of persons, tradition, practice, or the like is an individual, and to use "group-" or "culturally-relative" for cases where it is not. So, e.g., In Rachels, *The Elements of Moral Philosophy*, p. 16, the relativist sets up a direct opposition between "culturally relative" and "objective" moral views.

properties are objective. If a subjective theory is right, then moral values are up to us to the extent that the relevant attitudes are up to us. If an objective theory is right then, for at least some cases, the attitudes are irrelevant.

(B) Some examples. So far, our discussion has been fairly abstract, and therefore maybe difficult to follow. Some examples should help solidify the notions. First, some paradigm cases: A thing is *round* just in case all its points are equal distance from its center. Apply our test, and so imagine altering attitudes toward this condition: Suppose people don't like it when all the points of a thing are equal distance from its center, legislate against it, or whatever. It doesn't matter, in the end, a thing remains round just in case all its points are equal distance from its center. So *being round* is an objective property. In contrast, a skirt is *fashionable* just in case it meets certain length, material, and cut characteristics. Imagine altering attitudes toward these conditions: Suppose people think that to meet these conditions is to be ugly and disgusting. Then skirts with those characteristics aren't fashionable. So altering attitudes toward the characteristics alters the condition under which a thing is fashionable. So *being fashionable* is a subjective property.

Any subjective property depends on attitudes. But not every property that depends on attitudes is subjective (this is what makes the notion so slippery). Consider first, *knowing at least ten things*. Clearly this involves *knowing that* and so attitudes. But the property is not subjective. Subject to some debate, we may say that a person knows at least ten things so long as they have at least ten appropriately justified true beliefs. And revising attitudes toward this condition will not make it the case that a person with at least ten appropriately justified true beliefs doesn't know at least ten things. So *knowing at least ten things* is objective.

Another way to bring out this point is to consider the difference between (a guy's) *being admired* and *being admirable*. Say a guy is *admired* just in case certain positive attitudes are directed toward him by others (or, for that matter, by himself). That this is what is involved in being admired will not change if we alter attitudes toward the condition – in this case, having positive attitudes directed towards them. So *being admired* is objective. But *being admirable* is not. Perhaps some admire guys (along the lines Brad Pitt, or at least his image) with a ready smile a quick wit. Others may admire ones more macho, muscle bound covered in hair (along the lines of Sylvester Stallone, or at least his image). Debates along these lines may go round-and-round, ending with the conclusion “it's all subjective.” In this case, the conclusion is right. Suppose we say a guy is *admirable* just in case he has a ready smile and a quick wit. Clearly, altering attitudes toward these characteristics alters the condition under which a person is admirable. If one tends to puke at the sight of Brad Pitt, then the ready smile and quick wit won't do.

So subjective properties involve attitudes *at the level of the condition*. If the condition under which a property applies depends on attitudes about meeting the condition, then the property is subjective. Thus, again, subjective properties include *being rude*, as when talking with your mouth full (we can imagine a culture with attitudes according to

which talking with your mouth full is a sign of appreciation), and *being illegal*, as running a red light (in this case, the attitude involves legislation, and it is easy to imagine laws according to which red means go). Similarly, standards of taste and smell are subjective.

The status of moral values is not so obvious. A view like Ruth Benedict's, according to which "morality differs in every society, and is a convenient term for socially approved habits," would seem just to *be* a theory according to which moral facts are subjective.⁴ In contrast, G. E. Moore held that just as an object might be red or a parade long, so *being morally good*, *being morally bad*, *being morally right*, and *being morally wrong* are brute properties that an action or state of affairs might have.⁵ It is not immediately obvious how to identify such properties but, on his theory, actions or states of affairs have them "no matter what anyone may think"; so his is an objective theory of moral values. There are many other ways to have an objective or a subjective theory of moral values. Still, our question should be coming clear: we want to know whether moral values depend on attitudes like fashion or etiquette on the one hand, or are independent of attitudes like the shape of a ball or the earth on the other.

(C) Potential misunderstandings. The basic pieces of our distinction are in place. However, it may be helpful to distinguish what is typical (or, better, *stereotypical*) of the different views from what is definitional of them. Consider the following claims:

The objectivist insists that he is right and everyone else is wrong. Not necessarily. The condition for an objective property does not involve attitudes with respect to its condition so, in this sense, the truth or falsity of an objective claim does not depend on what we (or anyone) think. But do not mistake the claim that a sentence *has* a definite truth value for the claim that anyone knows what it is. Presumably, an objective claim *has* a single truth value. Therefore, if one person thinks the claim is true and another thinks it is false, one of them is mistaken. But the objectivist need not have any view about *which* of them is mistaken. Thus there is room for a tolerant and even modest objectivism: in cases of disagreement, the objectivist must hold that someone is right and someone is wrong, but need not insist that he is right and everyone else is wrong. But perhaps the point is just that, in cases of disagreement, the objectivist must hold that someone is right and someone is wrong, where the subjectivist might allow that *both* are right—one correct relative to certain attitudes and the other correct relative to others. This much is correct. Insofar as subjectivism allows that moral values may vary with attitudes, subjectivism does incorporate a sort of flexibility which objectivism does not. With this said, though, dogmatism is not the exclusive property of objectivism. It's just as easy to imagine a subjectivist screaming that good and bad are up to him as it is to imagine an objectivist screaming that they are not.

⁴ R. Benedict, "Anthropology and the Abnormal," *Journal of General Psychology* 10 (1934): 73.

⁵ G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903).

The objectivist holds that all moral truths are absolute—on an objective theory, if an act is morally wrong in one context, it is wrong in every context. Not necessarily. On an objective view, being morally good, morally bad, morally right or morally wrong doesn't depend on attitudes directed at the relevant conditions. But there are other sorts of dependency. So, e.g., the objectivist may accept a theory according to which killing a human being is morally wrong except on alternate Thursdays in Detroit or, more seriously, according to which killing a human being is morally wrong except when it would save multiple lives. Perhaps the objectivist thinks murder is always wrong in such-and-such contexts. But this makes the moral truth *universal*, not *absolute* in the above sense. Objectivism thus leaves room for a certain context-sensitivity. Of course, as above, it excludes the sensitivity to attitudes which is characteristic of subjectivism.

The subjectivist holds that no moral truths are universal—on a subjectivist view, the truth or falsity of a moral claim varies from one culture or context to another. Not necessarily. On a subjective view, the truth or falsity of a moral claim depends on attitudes. So truth or falsity may vary where attitudes vary (this is the point of our test for subjectivity). But nothing about subjectivism requires that the attitudes *do* vary—subjectivism sets up only the *potential* for variation. So, e.g., 'Ice cream tastes good' remains subjective even if everyone in the world likes ice cream. If the attitudes work out right, an objectivist and a subjectivist might *agree* about the moral evaluation of any, and even every, particular event. Of course, a subjectivist would hold that the moral facts are subjective and the objectivist that they are not. In principle, what separates objectivism from subjectivism is not the *content* of the moral principles accepted, but the *basis* for them: the subjectivist looks to attitudes where the objectivist does not.

The objectivist holds that attitudes are irrelevant to moral value. Not necessarily. First, consider 'An action is morally good if it is respecting those who respect you'. Attitudes are involved, but perhaps not attitudes *about* respecting those who respect you. On a subjective view, attitudes toward respecting those who respect you make it morally good or bad to do so. On an objective view, respecting those who respect you may be good or bad no matter what anyone thinks about it.⁶ Second, in typical cases, the objectivist may accept (purportedly) objective principles of the sort, 'It is good to obey the law' or 'It is right to obey your parents'. Then particular claims like 'driving on the left side of the road is good' or 'taking out the trash at night is morally right' depend on particular attitudes of a legislature or parents. The objectivist claims only that some moral principles do not depend on attitudes—not that all do not. In contrast, the subjectivist holds that every moral fact is attitude-dependent.

⁶ Similarly, though 'Broccoli tastes good' is subjective, 'Bob likes broccoli' is not. Of course, the latter involves Bob's attitude toward broccoli, but *holding the objects fixed*, attitudes about the objects—about Bob and broccoli—don't matter. This exposes an important point: underlying every subjective fact is the *objective* fact that the person, group of persons, tradition, practice, or the like have the relevant attitudes. One might say that a subjective fact about some object just *is* a broader objective fact including the object and whatever has the relevant attitudes.

II. Mackie's Argument

Mackie thinks ordinary moral claims require or assume objectivity—so that if ordinary moral claims are true, then there are objective moral values. But, at the same time, he thinks there are no objective moral values. He concludes that ordinary moral claims are not true. As above, his reasoning is valid: *necessarily*, if the premises are true, then the conclusion is true as well. So we need to know if the premises are true. In this section, I consider his premises in order, and develop reasons for thinking they are true. In the next section, we will ask whether there is a way to undercut them.

(A) Mackie's first premise: If there are "ordinary" moral truths, then there are objective moral values. Mackie thinks that ordinary moral claims require or assume objectivity. His own discussion mixes observations about what certain traditional moral thinkers have said, with observations about ordinary moral practice. I focus on considerations of the latter sort, and make some observations about moral *criticism*, moral *method*, and moral *authority*.⁷ In each case, the idea is that ordinary practice somehow sets moral values "apart" from human attitudes, and so treats the moral values as objective. Of course, again, what is ordinary in one context may be extraordinary in another. At best, then, the observations apply to certain, maybe traditional, approaches to morality. To avoid controversy about application, let's simply take "ordinary practice" as a *name* for any practice where moral criticism, moral method and moral authority are as follows:⁸

Moral criticism. The issue here is not criticism which arises when there is violation of accepted moral standards; rather, it's criticism of accepted moral standards. Ordinary practice accommodates such criticism. Suppose attitudes are arranged in some society so that, on a subjective view, it is morally appropriate for women to be beaten, raped and even killed. Perhaps everyone in this society, including women, accept that such treatment is morally appropriate. Can we make sense of the suggestion that this is morally wrong? If we can, then there would seem to be some sense in which the values stand "above" or "outside" the attitudes—and so a sense in which the values are objective. Notice: In asking whether the principles are wrong, we are not asking what the members of this society *think*—that is given; rather, we are asking whether there is a separate question about what is *wrong*. Also, their behavior may be wrong relative to the standards of some other society; but, as a person driving on the left side of the road in England doesn't do anything illegal (though they violate standards that would apply in another country), so we want to know whether members of that

⁷ Mackie, pp. 30-35. I don't follow Mackie directly. This discussion parallels, e.g., Rachels pp. 17-19. Rachels uses the relation between objectivity and ordinary moral claims, with the assumption that ordinary moral claims are correct, for the conclusion that moral values are objective. Mackie uses the relation between objectivity and ordinary moral claims, with a denial of objectivity, against ordinary moral claims. In this section, all I am after is the relation.

⁸ So "ordinary practice" may or may not be ordinary—as the Holy Roman Empire may not have been holy, Roman, or an empire, and a place may be called "Dartmouth" without being at the mouth of the Dart.

society do anything wrong (though they violate principles accepted in another society). With respect to this issue, consider your own reactions to the above example along with your reactions to, say, slavery, apartheid, or genocide. On the surface at least, a practice which includes criticism of accepted moral standards—whether our own or those of others—puts the values outside the attitudes.

Moral method. One might respond that when you criticize a moral principle, you don't have to be saying that the principle is mistaken; rather, such criticism need be no more than an attempt to persuade, and so to *make* principles morally wrong. But this brings us to our second point. Consider what people who are engaged in debate on some controversial moral topic, say, abortion, do. They do not say to themselves, "I propose to make abortion morally right (wrong), and so propose to engage in a propaganda campaign"; rather, advocates think they know what is morally right, and seek to bring others to (what they see as) the truth. Such persuasion may be an important part of a campaign to make abortion legal (illegal); it is, however, not conceived as an attempt to make abortion morally right (wrong), but as an attempt to make laws conform to moral reality. One may campaign for and create laws, but one doesn't ordinarily think about campaigning for or creating moral values. Or, again, think about the abortion debate, and consider the way argument works. If right and wrong were a matter of attitude, one would think that a survey or some self-evaluation could settle the overall question about whether abortion is morally right. But that is not how it goes. Somehow, the nature of the unborn's development, consequences for the mother, etc. are supposed to be decisive. Thus ordinary moral method makes right and wrong a matter for discovery, rather than creation, and a matter of external fact, rather than introspection.

At some level, these points might be accommodated by a subjectivist. She might hold that fundamental attitudes already have the consequence that abortion is morally right (wrong)—and all that remains is to help others to recognize consequences of their attitudes. Thus Mackie's point must be that ordinary moral method makes right and wrong a matter of discovery and external fact *to the exclusion of* attitudes. Return to a case where attitudes are arranged so that, on a subjectivist view, some X is sure to be morally right. The point about method is that, even so, matters of discovery and external fact might be sufficient for the conclusion that X is morally wrong. To the extent that this is so, ordinary practice puts moral values outside of attitudes, and so makes them independent of attitudes.

Moral authority. Finally, ordinary practice puts the *authority* of moral values above that associated with familiar attitudes. Imagine finding out that some action is opposed by your mother, is opposed by Democrats, isn't traditional, is illegal, and is immoral. The former may give you reasons for avoiding the action: avoiding the action will make others happy, make things better for you, etc. Still, the last seems to matter in a way that the others do not. The point here is not merely the point about method, that finding out about attitudes is not the same as finding out about moral values. Rather, it is that moral values *obligate* in a way that these attitudes do not. On Mackie's view, when a person makes a moral claim, what they say is "something that involves a call for action or for the refraining from action,

and one that is absolute, not contingent upon any desire or preference or policy or choice, his own or anyone else's."⁹ Again, think about the abortion debate and the authority a conclusion according to which abortion is morally right (wrong) is supposed to have over popular opinion, law, etc. Or think about the way we expect that orders to commit war crimes are to be refused. A practice on which moral values have this "absolute" authority seems to put the values above attitudes and so to treat them as objective. Of course, Mackie's view is that ordinary practice is therefore mistaken—and he tries to explain how people could be so systematically deceived. But the conclusion that ordinary practice is mistaken requires his second premise as well.

(B) Mackie's second premise: There are no objective moral values. This is where the argument from queerness comes in. Let me begin with an analogy. Consider what happens when a child concludes that there is no Santa Claus. Perhaps she discovers presents marked "from Santa" in her parent's bedroom closet, perhaps she sees parents filling her stocking, etc. There may be problems (*epistemological* problems) about the *evidence* for Santa—evidence which once seemed to show that Santa exists may fall away. But, further, the child may worry about how reindeer fly, how Santa visits so many houses, etc. That is, there may be problems (*metaphysical* problems) about the *nature* of Santa Claus—the child has beliefs about the nature of the world, and Santa doesn't fit in. If there is no positive evidence for a bizarre thesis, it is reasonable to conclude that the thesis is false, so the child concludes that Santa does not exist. Mackie thinks objective moral values are in the same boat: There is no positive evidence for them, and their supposed nature doesn't fit with what we know about the world; it is therefore reasonable to conclude that there are no objective moral values. In this case, the metaphysical and epistemological problems are linked. If it were part of the Santa story that both Santa and his presents were invisible, then the nature of Santa Claus would almost guarantee that there is no positive evidence for him; similarly, on Mackie's view, the supposed nature of objective moral values contributes to the lack of evidence for them. Thus the "root" problem would seem to be the metaphysical one, and it is important for us to understand how or why objective moral values are supposed to be "metaphysically queer"—to understand how or why there is a problem about their nature.

This problem for objective moral values is related to the third (moral authority) component of ordinary practice. Moral values are *prescriptive*—they prescribe what should and should not be done. Or, to put it another way, the morally good and the morally right have to-be-pursuedness and to-be-doneness somehow built into them, where the morally bad and the morally wrong have to-be-avoidedness and not-to-be-doneness built into them. But Mackie thinks "objective prescriptivity" is an oxymoron. It is one thing to imagine a person or a group of persons prescribing that such-and-such be done; this happens every day. But "prescribing that..." is an *attitude*. Thus, insofar as objective values are supposed to be divorced from attitudes, objective prescriptivity takes on an air of mystery: it seems to

⁹ Mackie, p. 33.

require an attitude that isn't an attitude. Of course, one might object that there is a sort of to-be-doneness that isn't attitude-based. But this seems mysterious as well.

To bring out this mystery, Mackie suggests reflecting on reasoning from David Hume's 1739 and 1740 *A Treatise of Human Nature*.¹⁰ Here is Hume's basic idea: Suppose you are offered the choice between \$1,000,000 and a punch in the nose. Does it follow, from the nature of the offer alone, which you will choose? It does not. Perhaps you have reasons for thinking that the money will lead to an empty, and therefore unhappy, life, while the punch will make you a better, and therefore happier, person. So you choose the punch. Suppose you somehow determine that taking the money will, in fact, lead to a happy life, while taking the punch will lead to a life of pain and despair. Does it follow which you will choose? It does not. For perhaps you are the sort of person who is into pain and despair! Or, again, suppose you are offered the choice between nuclear annihilation of the world and peace on earth. Does it follow, from this alone, which you will choose? No. The point is this: Objective facts, by themselves, are never a motive for action. Of course, what you do depends on the facts—but only with the addition of prior (subjective) components from the will. Given this, Hume's argument is simple: moral values are prescriptive—they are a motive for action—but objective facts are never a motive for action, so moral values are not objective.

It is natural to object against this argument that a *prescription* isn't the same thing as a *motive*; surely many people aren't motivated to do what is prescribed for them; so a problem about purely objective motives does not obviously transfer to a problem for purely objective prescriptions. But Hume thinks moral values are in very much the same boat as motives. Suppose I determine that pulling the trigger will cause a projectile to smash into such-and-such organism and so cause its biological functions to cease. Does it follow, from this alone, that pulling the trigger is wrong? It would follow, if we already had the fact that causing its biological functions to cease is wrong; but we haven't yet found *this* among the objective facts. So let's add some more objective facts: Suppose there is a lot of blood. Do the objective facts now add up to something moral? No. Now we need the fact that it is wrong to cause this bloody death, and we haven't yet discovered *this* among the objective facts. Or, again, suppose my pushing the button will result in nuclear annihilation of the world. Does it follow, from this alone, that it would be morally wrong to do so? It does not. The point is this: just as objective facts, by themselves, do not add up to motive, so objective facts, by themselves, do not add up to moral value. Of course, objective facts matter for decisions about right and wrong—but only with the addition of some (prior) moral values to tell us how they matter. If objective facts do not add up to moral values, then moral values are not objective.

As Mackie observes, Hume's reasoning does not actually remove the basic objection according to which there may be a sort of to-be-doneness that isn't attitude-based. It may be that objective facts do not add up to moral values only so long as one persists in leaving out

¹⁰ Mackie's discussion is on pp. 38-42; for Hume, see *Treatise* II.iii.3 and esp. III.i.1 in, e.g., the edition by Selby-Bigge, rev. Nidditch (Oxford: Clarendon Press, 1978).

the objective value facts! But now the force of the argument from queerness comes clear: Objective value facts cannot merely be facts about death, blood, and the like; by themselves, these are ethically (and motivationally) inert. But facts about death, blood, and the like are the ordinary facts with which we are familiar. Indeed, they may seem to be all the facts there are. At least, if there are objective value facts, they are extraordinary. But what could they be? And how do they generate prescriptivity? Like Santa Claus, objective values would seem to be entities utterly different from anything else in the world. (And, insofar as science and sensory experience seem directed just at ordinary facts, it's hardly clear what evidence one could have for objective values.) Mackie thinks it is therefore reasonable to conclude that there aren't any objective moral values.

Thus, Mackie's argument is complete. Ordinary moral claims presuppose that there are objective moral values; there aren't any objective moral values; so ordinary moral claims are false. Toward the end of his discussion, Mackie suggests that, "the objectivist may have recourse to the purpose of God: the true purpose of human life is fixed by what God intended (or, intends) men to do and to be"; and he concedes that, "if the requisite theological doctrine could be defended, a kind of objective ethical prescriptivity could be thus introduced."¹¹ On the face of it, it is surprising that Mackie would make this concession. On our definition, if God is a *person*, then an account of moral values which requires the intentions of God would be *subjective*. Still, we might have modified the definition to exclude God's attitudes from the subjective case. A moral theory based on God's attitudes would seem to give values a sort of independence from what *we* think that is consistent with objectivism. But, and this is the important point, why think that God's attitudes so much as matter for morality? Granted, God is a special sort of being. But, whatever God's nature, it is not obvious how the theist evades Hume's point. Suppose I clone some beings with the intent of having them work as slaves. Will it therefore follow that it is *morally right* for them to work as slaves? I think not. My being their "creator" is not enough. Perhaps God's commands come with a big carrot and a giant stick. But, again, how is this morally relevant? Might doesn't, by itself, make right. If facts about God are just more facts of the sort that don't add up to objective moral values, then Mackie's and Hume's argument sweeps away this approach to moral values as much as any other. In the next section, I'll say something about how God might make a difference.


III. An Objective Theory

There are, I claim, *ordinary* moral values, where these are "ordinary" in more than one sense: first, they are all that is presupposed by ordinary moral practice and, second, they are a familiar kind of thing to which arguments from queerness do not apply. I develop an approach to the good, and an approach to the right. According to the theory, the good is objective, but the right includes subjective elements. Taken alone, neither the objective nor the subjective elements are sufficient for ordinary moral practice; taken together, I claim,

¹¹ Mackie, p. 48, but cf., pp. 227-232.

they are. Given this, there is room for controversy about whether *ordinary* values should count as *objective* or *subjective* values. But it doesn't matter. On the one hand, if their subjective elements make ordinary values "non-objective," then Mackie's first premise is false: insofar as ordinary moral truths require just *ordinary* moral values, they don't require objective values. On the other hand, if their objective elements make ordinary values "objective," then Mackie's second premise is false: insofar as there are *ordinary* values, there are objective values. Either way, a premise is false, and the argument is unsound. I take the latter as my official account of what the theory accomplishes (thus the title of this section).

Before launching into the theory proper, let me begin with a preview of coming attractions. Consider a parent, child, rock, tree, and freeway as follows.



Suppose the parent commands the child, and the child behaves as in the following scenarios, with moral evaluations of the child's act as on the right.

	Command	Command for good	Action	Disobeys command	Wrong
1	don't go past rock	✓	child goes past rock	✓	✓
2	don't go past tree	✓	child goes past rock	–	–
3	don't go past tree	✓	child goes past tree	✓	✓
4	go play in freeway	✗	child doesn't go	✓	–
5	go play in freeway	✗	child runs into freeway	–	–
6	none	✗	child runs into freeway	–	–

In my experience most (but not all!) students agree with these evaluations. At this stage, it is not important whether you agree or disagree. If you agree, you may see the theory that follows as an explanation or account of your evaluations. If not, the account makes a case for evaluations as above. Presumably, it is not good for the child to run into the freeway and in the first cases, the parent issues the command for the child's own good. Observe, then that the child does something wrong when she violates a command which works "hand-in-hand" with the good, as in cases (1) and (3), but not in the other cases. So merely disobeying a command, or merely doing something bad for you isn't sufficient for wrong. With this much introduction, I begin with accounts of the good and the right, and then take up the relation between ordinary moral values and ordinary moral practice.

(A) The good. The account of the good depends on the notion of *satisfaction*. Let's begin with the simple observation that there are satisfaction facts. People are more or less satisfied when they leave home, get married, have children, etc. Satisfaction facts have

something to do with the fit between circumstances and desires; a person is more satisfied when desires are met, and less satisfied when they are not. Similarly, there are satisfaction facts for other creatures—not the satisfaction that you would have if, e.g., you were in a frog’s skin looking through its froggy eyes with your human thoughts; rather there are facts about “frog satisfaction”—whatever that may be. It may not be easy to put a number on satisfaction, but lack of numbers does not imply lack of facts. Presumably, there are also facts about satisfaction in non-actual circumstances. It is a fact that you would (wouldn’t) have been more satisfied if you had (hadn’t) gone to college right out of high school, if you had (hadn’t) married, etc. Again, the satisfaction that you would have had depends on the fit between desires and circumstances. But notice: The relevant fit is not the fit between *actual* desires and other circumstances; rather, it is the fit between circumstances and the desires you *would* have had in those circumstances. The actual desires of a drug addict are satisfied by the drug; but, in other circumstances, he or she might be satisfied—even more satisfied—without it.

Now for the hypothesis about the good. This is the part of the theory that is “particularly simple and perhaps mistaken.” The hypothesis is supposed to make sense of ordinary moral practice. For our purposes, any other (objective) proposal would do as well. Hypothesis:

Something is *good for* a creature iff it would increase that creature’s overall satisfaction. Something is *bad for* a creature iff it would decrease that creature’s overall satisfaction.

Perhaps this analogy will help: One might say that something is *profitable for* a person iff it would increase that person’s *net wealth*. Corresponding to objective facts about net wealth (in, say, 2010 dollars), are objective facts about profitability. It may not be easy to *discover* profitability facts; however this does not mean that there are no such facts to be discovered—either your net wealth would or wouldn’t have been greater if you had (hadn’t) gone to college out of high school, etc. Given this, there is room for a (not entirely precise) science of profitability facts—a science which might be called “economics.” Similarly, corresponding to facts about satisfaction, are facts about the good—about what would or would not increase a creature’s overall satisfaction. It may not be easy to discover these facts; however this does not mean that there are no such facts to be discovered—either your satisfaction would or wouldn’t have been greater if you had (hadn’t) gone to college right out of high school, etc. Given this, there is room for a (not entirely precise) science of facts about the good—a science which might be called “ethics.”

On this view, the good is objective. Perhaps this is surprising. As is the case for happiness, there may be a simple sort of satisfaction that does not involve attitudes; plausibly, however, even “simple” satisfaction depends on satisfactions *that...* Say this is right. Still, that something is good for an individual does not *itself* depend on attitudes directed at the condition. Your overall satisfaction may be increased by events you directly desire, but it may also be increased by events about which you have no attitudes at all (events

which take place while you are asleep, or in other parts of the world), and even by events which you regret (a spanking, or having the drugs to which you are addicted removed). Thus, that such-and-such would increase overall satisfaction (is good) does not itself require anything about attitudes with respect toward the condition. So it is objective. As above, it may be difficult to *discover* what is good for a creature—particularly if the creature isn't human. But this doesn't mean that there are no such facts to be discovered.

Already we can see how this approach to the good makes room for ordinary *moral criticism* and ordinary *moral method*. First, insofar as it is possible to be mistaken about the good, there is room for criticism of moral principles. Consider a culture according to which it is good for men to be strong and to suppress emotions, but good for women to be weak and to express them. Perhaps this culture is mistaken, perhaps not. The point here is only that such disagreement is not an empty “yes-it-is, no-it-isn't” debate of the sort that may arise about the taste of broccoli. Rather, it is a substantive disagreement about objective fact. We want to know whether this behavior promotes overall satisfaction or not. Second, with respect to moral method, the good's objective nature makes it a matter for discovery. It may not be easy to discover whether some behavior promotes overall satisfaction but, in an attempt to find out, we want to know about the behavior's consequences—as opposed to, say, propaganda campaigns mounted in favor of one side or the other.

But there is a problem—the very problem Mackie thinks all objective theories face—about the *moral authority* component of ordinary moral practice. Insofar as the good is objective, we run directly into the teeth of Hume's argument: It's not clear how facts about satisfaction can add up to facts about what we are obligated to do—unless there is some extra fact according to which we are *obligated* to promote satisfaction. For a response, I turn to the account of the right.

(B) The right. On the view being developed here, moral right is grounded in the moral good. But the right introduces or requires *responsibility* as well. So let's say a bit about responsibility. First, responsibility is a not-necessarily-moral social phenomenon which may bring with it praise, blame, reward, and/or punishment. So, e.g., you might be responsible to act as lookout as part of a gang robbery and be subject to praise, blame, reward and/or punishment for your performance. In ordinary cases, this responsibility isn't a moral one. Similarly, a shortstop is responsible for a certain part of the field—where this responsibility may bring with it praise, blame, reward and/or punishment. But, again, the responsibility isn't moral. Second, responsibilities can be given or taken—where there may be more or less giving and taking in different cases. So, e.g., a child may be given the responsibility of taking out the trash—though he doesn't want it. And a dictator might take responsibility for regulating industry—though his people never desire him to have it. Contract situations seem to be situations where giving and taking is more equal. Finally, the giving and taking of responsibility works only along certain, socially established, lines. A parent may give his or her child the responsibility of taking out the trash, but a neighbor cannot. This is not to say that in some other social organization (say a commune) a non-parent might do it; rather, it is to say that responsibilities can't be given or taken any old way.

Similarly, there may be giving and taking of responsibilities between a people and their leaders, where there are problems about one nation or culture giving responsibilities to another.

Insofar as they are a social phenomenon, responsibilities are subjective. Your responsibility to act under such-and-such conditions depends on attitudes with respect to those conditions—take away the attitudes and you take away the responsibility. Thus the account of the right mixes objective and subjective elements. Hypothesis:

An act is *morally wrong* iff it violates a responsibility with respect to the good of some creature or creatures. An act is *morally right* iff it is not morally wrong.

The subjective element comes from responsibility, the objective element from the good. There may be responsibilities “with respect to the good” which involve discerning the good—as a parent may take or be given a general responsibility for the good of his or her child, and there may be responsibilities “with respect to the good” which don’t involve directly discerning the good—as, for her own good, a child may be given the responsibility not to run into the street.¹²

The view is illustrated by the case of the rock, tree, and freeway. First, the view allows a certain flexibility under different responsibility contexts. This is illustrated by cases (1) and (2) where going past the rock is wrong in the one but not the other. More fundamentally, responsibility and the good are individually necessary, as illustrated by (4) where the child disobeys a command that isn’t for the good, and (5/6) where the child does a thing that isn’t good but doesn’t disobey any command. And the conditions are jointly sufficient as illustrated by cases (1) and (3) where they work together. Insofar as the objective good is required for right and wrong, not just anything can turn out to be right or wrong. And there is obligation because of the responsibility component.

(C) Ordinary practice. Presumably, we have responsibilities to the good. So there are ordinary moral values (as described above). It remains to show that ordinary moral values are all that is required by ordinary moral practice.

As above, the objective nature of the good makes room for ordinary *moral criticism* and *moral method*. Further, now we have room for prescriptivity and so *moral authority*. People are subject to praise, blame, reward and/or punishment insofar as they do or do not fulfill their responsibilities to the good. It’s worth noting that there is now room for specifically *moral* criticism of people for being mistaken about the good. So, e.g., a parent who accepts that running in the freeway builds character and so is for the best, probably makes a factual error about the good—where this might itself count as failing a moral responsibility.

¹² So, in the traditional vocabulary, we have room for both “deontological” and “consequentialist” aspects of the right.

But there are also limits to legitimate moral criticism. Different people may take or be given different responsibilities to the good, and there may be different, equally efficient, ways of assigning responsibilities. One parent may use a tree as his child's boundary and another a nearby rock—where, with respect to the good, one is as efficient as the other. Thus it may be that something is right for one person that is not right for another. Similarly, I am not responsible to distribute my goods equally to the children of the world precisely because my responsibilities to them (and I have some) are different from responsibilities to my own. This is a significant concession toward “relativistic” factors which may motivate subjectivism—but without sacrificing a capacity for moral criticism which ordinary moral practice requires.¹³ So far, then, the good and the right, on our account, may seem sufficient for ordinary moral practice.

But the subjectivity of moral prescriptivity may suggest reason to worry. Perhaps giving moral responsibilities at a high “degree” is sufficient to override other responsibilities and so to make sense of ordinary moral practice. But the subjectivity of prescriptivity spills into worries for moral criticism and method. For an extreme case, consider a culture that assigns no responsibilities whatsoever for the good, or assigns no responsibilities for the good of certain creatures (say, animals or the unborn), or assigns responsibilities but only at a minimal level. On our account, responsibility is required for moral wrong, and responsibilities can't be given in *any* old way—especially by one nation or culture to another. Thus, on our account, if nobody gives or takes the relevant responsibilities, *there is no wrong*. So far, we have been able to make sense of moral method and criticism, but only relative to the objective good. Ordinary method makes sense so long as responsibilities are understood, and what is at issue is the nature of the good. But it's not clear how investigation of the good is relevant to the responsibility component of right. And similarly, it's possible to criticize a culture for being mistaken about the good. But *moral* criticism seems problematic apart from responsibility.

Here is where God might help. If there is a God who is the creator of the universe then, as a parent to his or her children, or a government to its people, God might be in a position to assign responsibilities to governments or cultures. Insofar as God stands in a unique relation to us, such responsibilities might override others. And, if there is a God who assigns to cultures and governments at least general responsibilities for the good then, as one may criticize anyone who ignores responsibilities they are given, so one may criticize governments and cultures for failing to act on their God-given responsibilities. Thus, God's

¹³ And the concession has more moral “weight” than Rachels' related concessions on pp. 22-24. Compare also a sort of criticism that J. Chandler directs against theories which link moral values to the commands of a *loving* god (*American Philosophical Quarterly* 22 (1985): 231-239). He argues that, insofar as the moral values find their ground in god's loving character, god's commands become irrelevant to the moral values. But, as is the case here, the character (or the good) may impose a constraint that is necessary, but not sufficient, for the right.

attitudes might underwrite components of ordinary practice that our view seems to miss.¹⁴ It remains, however, that responsibility is attitude-dependent; if *neither* God nor people assign even general responsibilities to some good, there is no wrong.

Perhaps, though, the subjectivity of moral prescriptivity is a strength for our view, not a weakness. In typical cases, responsibility to the good is a given; so, on our account, in typical cases, ordinary moral criticism, moral method, and moral authority are vindicated. But some of the most intractable moral problems are ones (involving, say, animals or the unborn) where responsibility is in doubt. Once relevant responsibilities are in place, the question of whether, say, spanking is right or wrong is a question of its objective consequences. But objective investigations directed at animals or the unborn won't *themselves* assign responsibilities. With or without God, the assigning of responsibilities is attitude-dependent and so subjective. But this puts us in a position to *explain* why ordinary moral method doesn't seem decisive in problematic cases: the problematic cases are cases where responsibility is in doubt, and ordinary methods therefore don't apply.¹⁵ Maybe, then, this component of the view actually suggests *compatibility* with what's ordinary about "ordinary moral practice."

Insofar as our approach to the good and the right does account for ordinary moral practice, one of Mackie's premises must fail. If the overall approach counts as "subjective," his first premise is false; if the overall approach counts as "objective," his second premise is false. Either way, Mackie's argument is unsound. Here is a simple diagnosis of Mackie's mistake: Mackie treats moral values as monolithic entities, and argues that they cannot be both objective and prescriptive; it follows that ordinary practice is on the rocks. We distinguish the right from the good, and argue that there is room to locate objectivity in one place and prescriptivity in another; thus, plausibly, ordinary practice survives Mackie's objection.¹⁶

¹⁴ In Plato's dialog, *Euthyphro* 10a in e.g. the translation by Cooper in *The Collected Dialogues of Plato*, ed. Hamilton and Cairns (Princeton: Princeton University Press, 1973), Socrates famously suggests the question whether conduct is right/good because the gods command it, or whether the gods command it because it is right/good. On a religious version of our view, we have a "split" response: conduct is right because the gods command it, but the gods command it because it is good. So the good constrains the right, but a role remains for the commands. Of course, one might maintain that the gods themselves create satisfaction facts and are therefore responsible for both the good and the right. But, on our account, the responsibility for the one derives from attitudes in a way that responsibility for the other does not. In "Divine Commands and the Social Nature of Obligation," *Faith and Philosophy* 4 (1987): 262-275, R. Adams develops a closely related social notion of obligation in association with his rather different divine command theory.

¹⁵ Insofar as animals and the unborn don't enter into contracts, at least one part of the give-and-take which ordinarily clarifies who is responsible for what is missing. So we might expect controversy about responsibility in precisely these cases.

¹⁶ The paper has its origins in a long ago classroom debate on the Euthyphro dilemma. The idea was that the regular professor would play Socrates, I would be Euthyphro – and I would get steamrolled, as happens to Euthyphro in Plato's dialogue. But I didn't want to lose! So I gave some thought to "what Euthyphro should have said" – with results along lines from above. Though most wouldn't agree (and I'm

not sure I agree) this paper owes much, in inspiration and content, to discussion with Matthew Davidson, Meggan Coté, Susan Finsen, Richard Jensen, Chris Naticchia, and students in different introductory philosophy courses. I thank them and my wife, Rose Roy, for many helpful comments.