

Skepticism About Moral Knowledge¹

by

Geoffrey Sayre-McCord

UNC/Chapel Hill

Introduction

All the standard arguments for global (unrestricted) skepticism, needless to say, apply *mutatis mutandis* to moral skepticism. Whatever reason there is to doubt we know anything at all is reason to doubt that we have moral knowledge. Yet our moral views are distinctive in ways that invite skepticism even against the background assumption that global skepticism is false. One doesn't need to be a skeptic about everything to have serious doubts about peoples' claims to moral knowledge.

The grounds of moral skepticism are legion. They are found in the radical and apparently irreconcilable moral disagreements that are a familiar fact of life. They are found in the difficulty people often have in articulating, let alone explaining in any precise way, what they mean when they say that something is morally wrong or a moral duty. They are found in the fact that, as it seems, people might have all the same non-moral beliefs about (say) some course of action, about the alternatives, and about their effects, and yet still intelligibly disagree about whether so acting is morally permissible. They are found in the recognition that moral arguments are not something that might even conceivably be settled by empirical investigation. They are found in the striking extent to which an appeal to intuitions or gut feelings (that not everyone shares) are the only "argument" people have for the moral views they embrace. And (finally?) they are found in the ways in which peoples' moral views look to be more a reflection of private interest or cultural practice or perhaps even biological forces, than a recognition of independent moral facts.

These familiar and often frustrating features of our moral life offer significant reason to doubt that there is such a thing as moral *knowledge*, notwithstanding the strength of peoples' moral convictions and their readiness to sacrifice themselves and others in the name of those convictions. Not surprisingly, there is a long history of moral skeptics appealing, in different ways, to these features in support of their skepticism. Along the way, a remarkable variety of

¹ Copyright©Geoffrey Sayre-McCord

positions have emerged. Yet they all have in common the doubt they cast on moral knowledge.²

One good way to understand the variety of views that count as versions of moral skepticism is to relate them to a standard, and very influential, account of the nature of knowledge, moral and otherwise. According to this account, which traces its ancestry back to Socrates, knowledge is justified, true, belief.³ The guiding idea is that a person who claims to know something in particular will not count as knowing it if, as a matter of fact, she does not even believe it; and if she does believe it, she still will not count as knowing it, if what she believes is not actually true; and finally, even if she believes it, and it happens to be true, if her belief is not justified, if it is due to prejudice, or wishful thinking, or simply guessing, she will not properly be credited with knowing it.⁴ This account is regularly offered and defended as a general account of knowledge and is not supposed to be specific to moral knowledge. Nonetheless, specific versions of moral skepticism can be sorted according to whether (i) they doubt that people have moral *beliefs*; or, granting moral beliefs, (ii) they doubt that such beliefs are *true*; or, granting that people might have moral beliefs, and even granting that some of them might be true, (iii) they doubt that people ever have an appropriate *justification* for those beliefs.

Here I write of these skeptics *doubting* that we have moral beliefs, or that these beliefs are true, or that they are justified. The watchword of classical skepticism is doubt. But it is worth

² See Walter Sinnott-Armstrong's Moral Skepticism (New York: Oxford University Press, 2006) for an extended discussion.

³ This account is restricted to what is often called "knowing *that*" and is not meant to capture the conditions under which someone might count as knowing *how* to do something. This is worth noting since a good part of moral knowledge seems to be much more a matter of knowing how – how to respond, how to be sensitive, how to stand your ground – than knowing that – that something is wrong, or that it is important to pay attention to, or that it is worth fighting for. Still, claims about which bits of know how are morally important are themselves claims *that* something is morally important, and moral skepticism with regard to moral know how can reasonably be seen as an implication of moral skepticism about claims *that* knowing how to do certain things is morally important.

⁴ Not surprisingly, though this account is both standard and influential, it is far from uncontroversial. Most of the controversy has been about what should be added to the account, in the face of a broad range of examples that seem to show that not all cases of justified true belief are instances of knowledge. But some argue that the account is actually too restrictive. Most commonly this latter charge goes with holding that a person might know something even though she has no justification for it, as when (of instance) her beliefs reliably track the truth in the relevant domain, though she herself has no evidence of this. It is at least hard to imagine someone defending the view that knowledge does not require truth or belief, except if we are focusing on knowing how, rather than knowing that.

noting that many who are skeptical specifically about morality go beyond doubt to denial, arguing that there is reason to deny, and not merely doubt, that we have moral beliefs, or that these beliefs are true, or that they are justified. In the context of global skepticism, arguing in this way risks being self-defeating. But if the skepticism is focused on morality and plays out against a broader commitment to our having some non-moral knowledge, it is much less paradoxical to argue that, among the things we know is that we do not have moral knowledge.

In what follows, I will canvass, in order, the arguments people have offered for thinking that we don't have moral beliefs, that if we do, none of them are true, and that, even if some might be true, we are never justified in thinking of any of them (or, at least, any that are substantive) that they are true.

Do People Have Any Moral *Beliefs*?

A number of moral skeptics count as such not because they deny that some actions are morally wrong and others morally required, or that some institutions are just while others are not. Rather, their skepticism is born of their understanding of what we are doing in making such claims. According to them, we are not expressing moral beliefs, which they deny we have, but are instead expressing our emotions, or voicing our feelings, or prescribing behavior.

Initially, the suggestion that people do not have moral beliefs might seem difficult to take seriously. Of course, hypocrites are a familiar fact of life and there is really no denying that some people espouse moral views that they do not actually believe. Cynics may well argue that such hypocrisy is rampant and that everyone's motives and behavior belie their claims to believe in morality.⁵ Cynics aside, though, it seems clear that some people do sincerely believe that racism is morally wrong, or that justice demands equal treatment, or that people should keep their promises except if they find themselves in extraordinary circumstances. All of these claims seem to express beliefs that at least some people genuinely have. In fact, moral beliefs sometimes seem to be among a person's strongest commitments. Indeed, people regularly seem both to kill and to die for their moral convictions.

Still, there is a powerful strain of moral skepticism that has its roots, not in cynicism, but nonetheless in denying that people have moral beliefs. Unlike the cynics just mentioned, these

⁵ Cynics come in different varieties, of course, and some may be willing to grant that people genuinely believe as they claim even while arguing that believing is one thing, acting accordingly is another. On their view, the hypocrisy comes not with dishonesty about one's standards, but with a failure to live up to the standards one honestly believes apply. Cynics of this sort are skeptics about people's integrity but not, or at least not in virtue of this view of their fellow humans, skeptics about morality.

moral skeptics are perfectly willing to grant that people sometimes genuinely accept the positions they advocate. The skeptic's distinctive claim is that it is a mistake to see moral convictions, no matter how strongly held, as *beliefs*. Instead, they argue, a person's moral views are properly seen to be a reflection of their affective, and not their cognitive, attitudes. To hold a moral view, on this account is to have a distinctive attitude towards people, actions, institutions. To hold a moral view is a matter of looking at some things with things with approval and others with disapproval, or of being motivated to do some things and not others, or of hoping others will behave in one way rather than another. It is not a matter of having a belief with a distinctively moral content.

Of course these skeptics – who are usually called “non-cognitivists,” in recognition of their thought that there is no cognizing going on in making moral judgments – appreciate that we have the attitudes, motivations, and hopes we do in large part because of our beliefs. But the beliefs that play this crucial role are not, they argue, distinctively moral beliefs. Presumably, when we condemn someone for acting immorally we do so because of our beliefs about what she has done. Yet those beliefs, that she has broken a promise, or acted without considering the effects of her behavior on others, or whatever it is that leads us to disapprove of her, are non-moral beliefs about what she has done. They combine with our various attitudes (non-cognitivists argue) to cause our condemnation and, thus, our “judgment” that something wrong has been done. But the explanation of our moral view need make no appeal, they argue, to any beliefs that have a specifically moral content. And that is a good thing, they maintain, because there are no such beliefs.

According to non-cognitivists, once we distinguish a person's beliefs from her other (non-belief) attitudes, and, specifically, once we recognize the distinctive role of her preferences, of her desires and aversions, and of her approvals and disapprovals, none of which are themselves beliefs, we'll see (they argue) that a person's moral convictions are rightly located not among her beliefs but among these non-belief attitudes. A person without convictions, a person who has not taken a stand for or against things, a person who is indifferent to how the world is or might be (if such a person is conceivable) would be a person who so far lacks moral views (even if she might have all sorts of beliefs). While she might say that certain things are “right” and others “wrong” her indifference belies her sincerity. And this is because having a moral view is not a matter of what one believes but a matter of being engaged by how the world is or might be. Coming to grips with this, and recognizing the distinctive connection between ones moral convictions and one's responses to the world goes hand in hand, non-cognitivists argue, with recognizing that having a moral view is not a matter of forming a belief.

Indeed, an appeal to the distinctive connection – the apparently necessary connection – between a person’s moral commitments and what she is motivated to do is, I suspect, the single most influential argument for thinking that adopting a moral position is not a matter of forming moral beliefs.

The argument goes back at least to David Hume who emphasized the distinctively practical role of moral thought.⁶ The argument begins with the quite uncontroversial claim that there is a special connection between morality and action. And it follows up quickly with a much more controversial claim: that this special connection is best understood as being between the moral views a person accepts and her motivations.

The suggestion is that a person who sincerely holds that some thing is morally good, or that some course of action is morally right, is, *ipso facto*, positively motivated with respect to the thing or the course of action. And a person who sincerely holds that some thing is morally bad, or that some course of action is morally wrong, is, *ipso facto*, negatively motivated with respect to the thing or the course of actions. Of course, thinking this is compatible with recognizing that the motivations in question might be outweighed by other motivations, so that a person might on balance be motivated to do something she recognizes as wrong. But, the suggestion is, there must be a motivation there to be over-come, if she sincerely thinks the course of action is wrong. If, instead, she is utterly unmotivated to pursue, promote, or protect what she claims to be good or right (or to oppose and work against what she claims to be bad or wrong) her claims will have been shown not to express her genuine views.

With this view in place, the argument goes on to draw a contrast with beliefs all of which are (it seems) in themselves motivationally inert. Of course people are motivated to act as they do in light of their beliefs, but simply knowing which belief a person might have (knowing what it is she happens to believe) doesn’t by itself (the argument goes) tell us how the person will be motivated. To know how someone might be motivated we need to know not simply what she believes but we also need to know what preferences she has. This marks a sharp contrast, the

⁶ The *locus classicus* for this argument is in David Hume, *A Treatise of Human Nature*. But it played an especially important role in discussions and defenses of non-cognitivism starting with the work of C. L. Stevenson (who singled out the “magnetism” of the term ‘good’ as calling for special explanation) and A. J. Ayer (who argued that moral terms were designed to give voice to motivating states, not to report anything). See Stevenson’s “The Emotive Meaning of Ethical Terms” *Mind*, Vol. 46, (Jan., 1937), pp. 14-31 and Ayer’s *Language, Truth and Logic* (London; Gollancz, 1936). For the record, I think the standard interpretation of Hume has gotten him seriously wrong. See my “Hume on Practical Morality and Inert Reason,” in *Oxford Studies in Metaethics*, edited by Russ Shafer-Landau (Oxford University Press, 2007), pp. 299-320.

argument goes on, between a person's moral views, on the one hand, and her beliefs, on the other, since in knowing a person's moral views we do know that she will be motivated in certain ways. And this contrast shows, the argument concludes, that a person's moral views are not among her beliefs.

Of course there are various ways one might push against this argument. For instance, one might deny that in holding a moral view, one is, *ipso facto*, motivated at all one way or the other. Even if normally coming to hold that something is good or right has a motivational impact, there certainly seem to be cases where a person sincerely (and, one might want to add, accurately) recognizes that something is wrong and yet appears to suffer no hesitation in doing it. And, similarly, peoples seem regularly to lack any motivation whatsoever to do what they, apparently sincerely (and accurately), claim to be good or right. To insist, as the argument would require, that there is in every case, including these, some perhaps hidden or overpowered motivation seems to involve a devotion to theory that goes beyond the evidence.

Or one might argue that the motivational impact of our moral views actually establishes that, at least in some cases, knowing a person's beliefs (specifically her moral beliefs) does tell us something about that person's motivations. So, one might well accept that the special connection between morality and action is to be understood in terms of our moral views having an impact on our motivations and then maintain that this serves to show, contrary to the non-cognitivists' argument, that not all beliefs are motivationally inert. This suggestion is bolstered by the fact that there are other beliefs, say, concerning the prospect of pleasure or pain, that seem to have exactly the sort of connection to motivation that the non-cognitivists attribute to our moral views, though there is little doubt that we can and do have genuine beliefs concerning the prospect of pleasure and pain.⁷

In any case, precisely what the connection is supposed to be between sincere moral views, on the one hand, and motivation, on the other, is controversial. Virtually no one holds that

⁷ Of course not all our beliefs concerning the prospect of pleasure and pain motivate. But it is plausible to think that if the beliefs concern the prospect of our own pleasure or pain, in a context where we can have an influence on the prospect, we will be motivated to secure or increase the pleasure and to avoid or decrease the pain. Similarly, not all our beliefs concerning the prospect of things we think good and bad motivate, since often we take ourselves to have no influence on the prospect. Still, it is plausible to think that if the beliefs concern the prospect of something good or bad, in a context where we can have an influence on the prospect, we will be motivated to secure or increase the good and to avoid or decrease the bad.

thinking something good or right always goes with having an overriding motivation to promote or pursue or protect or do anything in particular.⁸

Many, though, claim that a person who sincerely holds that some act is, say, morally right, will *ipso facto* have some, though perhaps easily overridden, motive to perform it. This view is compatible with the recognition of weakness of will, which seems important to its plausibility, and it does a nice job of highlighting one ground we often appeal to in accusing someone of insincerity – their apparently lack of any motivation to act in accord with the judgments they claim to hold.

Others hold that the connection between sincere judgment and motivation is found only in those who are rational, or moral, or in crucial respects, normal. In holding this they allow that a person might sincerely make a moral judgment and yet fail to have any appropriate motivation, though they argue that whenever the appropriate motivation is missing it's a reflection of some failure or abnormality on the agent's part.⁹

This last view, of course, will not support a direct argument against cognitivism, though, since a cognitivist might consistently hold both that people have moral beliefs and that failure to be appropriately motivated is a reflection (not of the absence of the belief but) of a non-cognitive failure in the character of the person in question.

⁸ Although some people, especially in the context of defending error theories, which are discussed below, have been inclined to hold that in making moral judgments we are committed to thinking that a person cannot genuinely recognize something as morally required without having an overriding motive to act accordingly. This is of course quite a strong commitment and goes with, among others things, ruling out the possibility of being weak-willed (where that is understood as a matter of knowingly doing what is wrong). The use of the word "recognize" in "a person cannot genuinely recognize something as morally required" covers two possible views, one according to which one must have an overriding motive to do what one takes to be right (whether or not one is correct), the other according to which one must have an overriding motive to do what one takes to be right only in cases in which one is correct. These views differ in where they end up putting the burden of explaining the connection between our moral views and motivation. On the first, the connection is to be forged by a proper understanding of what it is to sincerely hold a moral view. This is the approach taken by non-cognitivists who argue that we can understand the motivational impact of our moral views only by seeing the moral views not as beliefs but as, themselves, noncognitive motivations. On the second, the connection needs to be explained by appeal to whatever it is that allows us to distinguish between the views that are correct and those that are not. Most commonly this is done by appeal to the nature of what would have to be the case for our moral beliefs to be true. This is the approach taken by those who think that we do genuinely have moral beliefs. In the hands of the error theorists, they go on to argue that while the truth of our beliefs would require motivationally loaded facts, there are none, so none of our moral beliefs are true.

⁹ See, for instance, Michael Smith's The Moral Problem (Oxford: Blackwell, 1994).

In contrast, those who deny the possibility of having the “belief” absent any appropriate motivation are in a position to argue that the supposed belief may actually not be one at all (if we grant that beliefs are not, by themselves, motivating). Yet they face the significant burden of establishing that the motivation must be present, even if not evident, whenever someone makes a sincere moral judgment.

A different argument against our having genuine moral beliefs focuses not on the intimate connection between moral judgment and motivation but on the difficulty people have in identifying and articulating just what they mean in saying that something is right or good or wrong or bad. The logical positivists pressed this point especially vigorously with the help of their empiricist criterion of meaningfulness. According to them, some term comes to have a meaning only thanks to there being some experience a person might have that would confirm the appropriateness of its application. What cannot, in principle, be tested by experience is, they argued, meaningless. Combine that with the claim that no experience would settle whether something is good, and you have an argument to the effect that ‘good’ is meaningless. So any putative belief that something is good is cognitively empty, a mere pseudo belief.¹⁰

But one need not be a logical positivist to worry that, in the end, there is no good way to identify, articulate, or define the distinctively moral terms of our language. Following G. E. Moore, many philosophers argue that any attempt to define moral terms is bound to fail. By way of argument, they point out that for any proposed definition (say, of ‘good’ as ‘pleasant’ or ‘the object of desire’ or ‘conventionally accepted’ or of ‘right’ as ‘conducive to happiness’ or ‘approved by God’ or ‘consistent with society’s rules’) one can intelligibly ask, of whatever met one’s preferred definition, whether it really was good or right. That these are intelligible questions seems to show that (i) thinking something satisfies the proposed definition and (ii) thinking that thing is good or right involves thinking two different things. Yet if the definitions were successful, the argument supposes, the thoughts would be one and the same (just as thinking something is a triangle is one and the same as thinking that it is a three side plane figure).

The worry is not about any particular proposed definition. Rather, it seems that any attempt to define moral terms by appeal to events or properties that are experientially testable or in other ways cognitively tractable, will leave open an important question: are things that satisfy the definitions actually good or right? Even if we end up answering “yes”, our being able intelligibly

¹⁰ A few logical positivists took this argument in hand, along with the conviction that ‘good’ and ‘right’ etc. were obviously meaningful and argued that they could be defined in experiential terms (say, ‘good’ by ‘pleasant’ and ‘right’ by ‘conducive to overall happiness’). See Moritz Schlick, Problems Ethics (New York: Prentice Hall, 1939).

to raise the question seems to establish that the proposed definitions fail as definitions. They each fail to capture a distinctive and important element of what we are doing in making moral claims.¹¹

What, though, is this missing important element? Non-cognitivists hold that it is not some additional cognitive component, over and above what is captured by some definition. Instead, they maintain, the extra element is best understood in terms of the non-cognitive attitude that goes with making a moral judgment. Yet, we are not, in making our moral judgments, *reporting* the presence of this attitude. We are not, in saying something is good, reporting that we approve of it. Rather, noncognitivists often argue, we are expressing the attitude. A sincere judgment that something is good or right reflects the fact that we do approve of what we take to be good or right (or reflects the fact that we hope others will pursue or do what we call good or right or reflects our preference for them pursuing or doing those things). While the details of various proposals cover a broad range, all of these views account for the distinctive nature of our moral convictions and claims by appeal to the attitudes they express as opposed to the (putative) facts they report.

This line of thought fits naturally with the motivation argument mentioned above, since we come back to the view that sincere moral judgments are tied to the presence of some sort of noncognitive attitude (that many have argued is a motivating attitude). Yet this convergence in the positive view advanced should not lead one to think the arguments for the view are the same. They aren't.¹² Nonetheless, they are nicely complementary, and their fitting together so well is reasonably taken as an additional consideration in favor of both.

Whichever route one takes to the view that in making moral judgments we are not reporting facts or expressing beliefs, but rather giving voice to our non-cognitive attitudes of approval, endorsement, or desire, non-cognitivists can still offer a positive account of moral judgment that plausibly accounts for the systematic nature of our moral thought and talk. To reject the idea that we are reporting facts or expressing beliefs is not to reject the idea that we are

¹¹ Many have suggested that all such definitions fail because they all transgress the fact/value distinction. In attempting to identify via definition rightness, goodness, etc. with some property that is cognitively secure, they each mistake claims concerning the realm of value with claims concerning matters of fact.

¹² In particular it is worth noting that someone might think the relevant attitudes of approval or whatever are not, in any interesting way, especially motivating, even as they are aspects of our affective rather than cognitive nature. Hume is an interesting case in point here, since he pretty clearly ties moral judgment to approbation and disapprobation while arguing that these are not motivating states. And this despite Hume also often being credited as the first to express the motivation argument clearly.

engaged in a practice that is structured systematically by various standards and goals that give point to the practice of moral thought and talk. And this means that there is room for the non-cognitivist to offer a plausible account of why our moral judgments might easily be mistaken for beliefs.

At this point, it is important to register that while non-cognitivists reject the idea that we have moral beliefs, and so (on the assumption that knowledge entails belief) reject the idea that we have moral knowledge, they can and often do offer an account of what we are doing in saying of someone that she has moral knowledge: we are expressing our approval or endorsement of her moral commitments. Just as our moral thought and talk is used to express our attitudes, so too (on this view) is our thought and talk about knowledge and justification. In counting someone as knowing something, or in characterizing her as being justified, we are, on this view, not reporting a fact or expressing a belief, but instead we are expressing some non-cognitive attitude of our own concerning her. This view extends naturally to cases where people do have beliefs, since we can approve of peoples' beliefs just as we can approve of their other attitudes. So it can, in its own way, explain what we are doing in holding, of peoples' true beliefs, that they constitute knowledge: we are approving of their having the beliefs they do. But this view goes with rejecting the idea that we have moral knowledge, if by that we mean we have some justified true beliefs (since it rejects the idea that we have moral beliefs at all).¹³

Do People Have Any *True* Moral Beliefs?

Without denying that it is difficult, or perhaps even impossible, to define moral terms in any way other than by appeal to other moral terms, many people hold that people are nonetheless expressing beliefs, and purporting to report facts, when they make moral judgments. They are

¹³ Quasi-realists, in contrast with traditional non-cognitivists, acknowledge that we do in fact have moral beliefs, although they go on to emphasize that these "beliefs" differ from others in some way that warrants the hedge provided by the "quasi" used to characterize their position. Their central project is to make sense of moral thought and talk, including the extraordinary extent to which it mimics straight descriptive thought and talk, without taking on the ontological commitments that would go with acknowledging that moral just are a species of descriptive judgment. Their challenge is to sustain the contrast. This gets harder and harder the more they succeed in otherwise accounting for the extent to which moral claims seem to express beliefs, report facts, and answer to the world. The quasi-ness of quasi-realism depends on being able to say that while there may be moral beliefs, that purport to report facts, and that are answerable to the world, they nonetheless differ in some important way from non-problematic judgments. See Simon Blackburn Essays in Quasi Realism (Oxford: Oxford University Press, 1993) and Allan Gibbard Wise Choices, Apt Feelings (Cambridge, MA: Harvard University Press, 1990).

not simply expressing their approval or endorsement of certain things rather than others, they are claiming that those things are right, or good, or whatever, and this is a claim that is properly evaluated as true or false in light of the way the world actually is. On this view, the fact that we *seem* to have moral beliefs, and *seem* to be claiming something about the way the world is, morally, when we make moral judgments, is explained by things being as they seem. According to cognitivists, we do in fact have moral beliefs and we are making claims about the way the world is morally. On their side are a number of considerations, most significantly that people regularly take themselves to have moral beliefs and to be making claims about the way they world is, morally, when they express those beliefs. Indeed, the non-cognitivists' suggestion that they are (merely) expressing their tastes or preferences or feelings of approval and disapproval, are often rejected out of hand as misconstruing what is at stake and what is being claimed. It is one thing, people insist, to dislike something, or even to be repelled by it, quite another to think it is bad or wrong. The latter involves forming a belief about the thing, perhaps a belief that justifies the feelings of dislike or revulsion it prompts, but in any case a belief that is distinct from, and something over and above, the feelings in question.

Thinking that people do have moral beliefs, though, is compatible with thinking all the beliefs are false, and the world is not the way people believe it to be. As a result, one might well embrace cognitivism and yet go on to advance an error theory, according to which none of the beliefs we have are true.¹⁴ This view accepts, at face value, the impression that we do have moral beliefs and that, in making moral claims, we are purporting to report facts about how the world is, morally. In that respect, it is less radical than non-cognitivism. In another respect, however, it is much more radical, since error theorists hold that our moral thought is

¹⁴ J.L. Mackie offers an extended defense of an error theory about morality in [Ethics: Inventing Right and Wrong](#) (Penguin Books, 1977). See also Richard Joyce's [The Myth of Morality](#) (Cambridge University Press, 2001). Mackie identified two kinds of arguments for the error theory. One, the "queerness argument" involved arguing that thinking moral beliefs are true commits one to an unacceptably strange ontology which includes, in addition to all the physical and psychological properties we recognize, properties that somehow manage to motivate all who recognize them, regardless of those peoples' desires, and also are discovered by us thanks to a faculty that manages to cognize both the properties and their relations to things in the natural world. It is all too bizarre, he argues, to take seriously, especially since we can explain well why people so consistently develop beliefs with these strange commitments. The other argument was the "argument from relativity" which involves maintaining that the best explanation we have of peoples' moral views, and the extent to which they reflect the cultures in which those people live, is that a person's moral beliefs are simply a reflection of their culture and upbringing (and reactions against those, in some cases) and not at all a response to distinctively moral features of the world.

fundamentally mistaken, whereas the non-cognitivists challenge our understanding of what we are doing in making moral claims, but leave our making those claims utterly unchallenged.¹⁵

A useful parallel here, in coming to understand the error theorist's position, is found in the modern view of the Greek Gods. The ancient Greeks had beliefs about Zeus and the other Olympian Deities, beliefs that were strongly held and tremendously influential – beliefs in light of which they guided their lives. Yet now we see those beliefs as fundamentally mistaken, no matter how sincerely held and no matter how salutary in their effects. Doubting the truth of those ancient religious beliefs – that is, being a skeptic about the existence of Zeus and the others – is of course not a position that everyone accepts, especially the ancient Greeks who were raised in the faith. Nonetheless, we have no problem making sense of the suggestion that the beliefs are false and that elaborate theology that played such a crucial role in Greek life rests on a fundamental mistake about what there is in the world.

Other useful parallels are found in, for instance, current scientists' rejection of beliefs concerning phlogiston, and the now common rejection of the idea that there are witches with occult powers. In these cases too, the idea is that there are beliefs people have sincerely held that we have reason to think are not simply false in a particular case but are such that they systematically involve presuppositions the falsity of which undermine them all. Thus particular claims about Zeus' desires, or phlogiston's effects, or the magical impact of a witch's incantation, are false because they presuppose the existence of things (Zeus, phlogiston, occult powers) that don't actually exist.

Similarly, error theorists about morality argue that our particular moral claims, which they grant often express sincerely held beliefs, are all false because they presuppose the existence of things that don't exist. They argue that our moral beliefs, however strongly held and influential, all rest on a fundamental mistake about what there is in the world.

Different error theorists trace the mistake to different specific presuppositions that, they argue, are part and parcel of holding a moral view. But they all maintain that our moral thought and talk presupposes the presence of something that, simply put, does not exist.

According to some, what we are presupposing in forming a moral belief is that there are some authoritative standards for human behavior. They maintain that if we thought there were no

¹⁵ Of course a non-cognitivist might well disapprove of our various approvals and disapprovals, and so challenge the moral views to which we give voice, but that disapproval is distinct aspect of the person's view, and not an implication of her non-cognitivism. Similarly, a cognitivist who rejects the error theory, and so believes that some moral claims are true, can consistently hold that most peoples' moral beliefs are false.

such standards – if we thought there are no standards at all, or if we thought there are standards, but that they are not authoritative – then it would make no sense to think of people or their actions, in terms of them being morally right or wrong, or morally good or bad. Our moral beliefs, in other words, presuppose the existence of authoritative standards. And, they go on, we have grounds for thinking there are no such standards, and so grounds for rejecting all the beliefs (i.e. all moral beliefs) that commit us otherwise.

According to others, we are presupposing that there are desire-independent reasons for people to act in certain ways, such that whatever desires a person may have, she has reason to do, or refrain from doing, certain things. In some cases, this view is just an elaboration of what would have to be true for there to be the appropriate sort of authoritative standards. In that case, the idea is that we should believe there are the standards presupposed by our moral beliefs only if we also believe there are desire-independent reasons, say because only then could there be relevantly authoritative standards. In other cases, though, the argument by-passes standards all together and depends, instead, on arguing directly that moral claims, say about what a person ought morally to do, could be true only if a person, no matter what her desires, would have reason to act accordingly. Either way, the error theorists go on to argue that we have grounds for thinking that whatever reasons we might have, none of those reasons are desire – independent, and so we have grounds for rejecting all the beliefs (i.e. all moral beliefs) that commit us otherwise.

According to still others, we are presupposing the existence of an all powerful and loving creator who sets the standards for our behavior, or establishes reasons for us to act that are independent of our desires, or in some other ways provides what has to be in place for any moral claims to be true. In some cases, the argument for thinking that moral beliefs presuppose the existence of God go via presuppositions having to do with authoritative standards, or desire-independent reasons, in which case the error theorists argue that those presuppositions themselves presuppose the existence of God. Other times, the error theorists argue directly that if there is no God, there is no morality. Either way, the error theorists go on to argue that we have grounds for thinking that there is no God, and so we have grounds for rejecting all the beliefs (i.e. all moral beliefs) that commit us otherwise.

In whichever way the details are worked out, the common position of error theorists is that we should reject as false or misguided our moral beliefs on the grounds that in holding them we are committed to the existence of something we have grounds for rejecting.

Clearly, the error theorist's position is only as strong as her grounds (i) for thinking the presuppositions she identifies are actually ones to which we would be committed in holding a

moral belief, and (ii) for thinking that the relevant presupposition is false. Predictably, then, cognitivists who reject the error theory regularly argue either (i) that the claims about what moral beliefs presuppose get things wrong or (ii) that, while they are right, the presupposition is satisfied. In taking the first line, the argument is that error theorists saddle those with moral beliefs with commitments that are not theirs to bear. In taking the second line, the argument is that the relevant commitments are, contra the error theorists, commitments we have, on balance, grounds for thinking they are satisfied.¹⁶

Working through these issues is nicely complicated and involves weighing evidence that certain presuppositions are in play, evidence that the presuppositions are not satisfied, and evidence that people might have gone so far wrong. And all of this gets balanced, in the process, against the evidence that some things are right or wrong, good or bad.

But this last bit of evidence needs to be handled carefully. The error theorists are challenging the claim that anything is right or wrong or good or bad. So, to the extent contrary evidence presupposes what is at stake (as presumably most of our moral evidence in favor of particular moral claims, does) there is a pressing worry about begging the question. Of course, error theorists acknowledge, if cruelty is always wrong, then evidence that someone has done something wrong is found in their having been cruel. But the claim that cruelty is always wrong, no less than the claim that a particular person did something wrong, presupposes, they will argue, just what is in question.

At the same time, however, the conviction that some things are genuinely right or wrong, good or bad, provides impetus for showing either that the purported presuppositions are not actually being made, or that they can be vindicated. And that fact understandably structures a great deal of how people set about defending their moral beliefs as not actually undermined by the error theorist's challenge.

In the process, error theorists and their opponents alike work to figure out what people are claiming in saying that something is right or wrong, good or bad, since arguments for or against certain candidate presuppositions turn on what is involved in these claims. As a result, both sides reject the non-cognitivists' claim that our moral claims have no content and do not purport to report the way the world is morally, even as they might disagree about what the content is and, so, about what is being claimed about the world. These debates introduce

¹⁶ In addition, there is room to argue that the very arguments the error theorists offer, that appeal to epistemic reasons that people have, independent of their desires, come with the sort of commitments the error theorists are trying to undermine, so that their position ends up being self defeating.

discussion of whether the putative moral facts would have to be something supernatural, something over and above what is found in the natural world or whether they might properly be identified with, or constituted by, natural entities or properties. And these debates force a discussion as well about what might constitute successful definitions of our moral terms, about whether such definitions are necessary or desirable, and about how it is that our language and our thoughts might come to be about various things in the way they evidently are.

In the context of a discussion of skepticism, however, the important point to recognize is that if the error theorists make their case, they will have established that we lack moral knowledge (assuming that knowledge entails justified true belief) because they will have shown that none of our moral beliefs are true. It is also worth noting, that the error theorists, in making their argument, are not claiming that no one has, or has ever had, justified moral beliefs. Error theorists argue that we have grounds for rejecting all moral beliefs as false, or misguided, on the grounds we have for thinking that they all have false presuppositions. Yet their position is entirely consistent with holding that other people who lack the evidence we have, or have misleading evidence of some kind, might well be perfectly justified in holding the moral beliefs they do.

Of course the error theorist is arguing that we have evidence that makes our beliefs unjustified. But the focus of her arguments is not the justification, or lack of justification, people might have for their beliefs. Her focus is on the (putative) falsity of the beliefs, however well justified they might be.

Are Any Moral Beliefs Justified?

But what would it take for someone's moral beliefs to be justified, assuming people have such beliefs and that some of them are true?

An attractive answer to this question, which usually flies the banner of "foundationalism" in epistemology, starts with the idea that a belief is justified if but only if (i) there is evidence for it or (ii) it is, in some way, self-evident (or not in need of evidence, though able to provide evidence for other beliefs, or, at least, not such that there could be any good evidence against it). And the answer continues with the thought that beliefs that fall in (i) do so thanks to being properly related to those that fall in to (ii).¹⁷

¹⁷ Descartes is especially credited with making the appeal of foundationalism clear and still stands as offering one of the most ambitious attempts to vindicate our ability to know of the external world in light of foundationalism's demands. See Rene Descartes, Meditations on First Philosophy. Descartes' commitment to showing that there is at least one belief – the cogito – that could not possibly be doubted, takes the demands of epistemic privilege as

The common metaphor is found in talk of foundations, with the suggestion being that type (ii) beliefs provide the ground level support for whichever of our other beliefs might count as justified. Thinking of things in this way brings two questions to the fore. Which beliefs fall in to type (ii)? (Here talk of self-evidence, certainty, indubitability, etc. looms large.) And what does it take for a belief to count, for purposes of justification, as being properly related those beliefs? (Here talk of entailment, evidential relations, explanations, inductive support, looms large.)

Put in the most general terms, foundationalists hold that for any belief, moral or not, to be justified, it must either be epistemically privileged in a way that puts it in type (ii), or it must be appropriately related to such beliefs. Among foundationalists, there is a great deal of disagreement both about what it takes to be epistemically privileged and, assuming an answer to that, what it takes to be appropriately related to such beliefs.

According to many, one appropriate relation is explanatory. Indeed, a vast majority of our beliefs are justified, they maintain, thanks to the role those beliefs play in explaining our experiences. The idea is that we have some evidence for our beliefs if their truth contributes to our explanation of other things we believe, which in turn (ultimately) contribute to our explanation of our experiences. So, for instance, we have reason to accept various scientific hypotheses not because they are self-evident, or indubitable, or in any other way, epistemically privileged, considered on their own, but because, on the supposition they are true, we then have an explanation of a great deal of other things we believe, that in turn we have evidence for thanks to their explaining our experiences of the world. Similarly, we are justified in believing that people have certain desires or beliefs because supposing they do contributes substantially to our ability to explain their behavior. In fact, many people think, all of our beliefs that are not either trivial, or directly about our experiences, are justified, if they are justified at all, thanks to the contribution they make to our explanations of what happens in the world and, specifically, our experiences of the world. Whether the beliefs have to do with the abstract and theoretical claims of physics, or the concrete and mundane lessons of cooking and auto mechanics, they all enjoy what justification they have by surviving the test of experience, which they do by contributing to our explanation of what our experiences have been.

being especially rigorous. Other foundationalists have allowed that beliefs might count, thanks to being, say, self-evident, even though they might be doubted and even though, in fact, they might be such that other evidence outweighs what they provide for themselves. One might embrace foundationalism while thinking the relevant foundational beliefs are dubitable, corrigible, and maybe even false. What they do need to be, for the purposes of foundationalism, is such that they are able to justify other beliefs without having to enjoy a similar support from other beliefs.

In thinking specifically about moral beliefs, explanatory relations have become especially salient.¹⁸ The problem, pressed by moral skeptics, is that the truth of our moral beliefs seems not to play any explanatory role at all in understanding of our own experiences nor of what happens in the world or why. Even when evil people do horrible things to virtuous people, we can have (in principle) a full explanation of what the evil people have done and why, and what happens to the virtuous people, and how they suffer, all without having to appeal to our moral beliefs at all. Psychology, sociology, biology, and physics, in some combination, provide the explanation. And while adding talk of evil and virtue may let us categorize those involved, it does not help us explain what has happened. Similarly, we might find ourselves repulsed by the behavior of the evil person and indignant at the treatment the virtuous person has suffered. And while we might be tempted to say that our being repulsed is explained by the person's evil nature and our indignation is explained by undeserved suffering, our reactions can be fully explained by appeal to our having the moral beliefs we do, without having to suppose they are true. We would have been repulsed and indignant, having believed the person evil and the sufferer virtuous, even if, as a matter of fact, our beliefs were in this case false. According to many moral skeptics, this point holds generally and shows that our moral beliefs are explanatorily impotent. True or not, they contribute not at all to our explanations of what happens in the world.

This is going just a bit fast. After all, the truth of our moral beliefs would, it seems, explain the truth of other moral beliefs. The wrongness of cruelty explains why a person who is casually cruel is evil, and the virtue of a victim explains why that person suffering is undeserved. And the moral skeptics in question do not deny this. Yet they point out that all of our moral beliefs are being called in to question. So these explanatory relations among our moral beliefs, supposing they hold, are of no justificatory value unless we have some other grounds for thinking, of some of moral beliefs, that they are justified.

If this is right, moral skeptics argue, then our moral beliefs could be justified only if either they bear some other relation to privileged beliefs, or they themselves fall in to the privileged class. But, according to the skeptics, neither is plausible.

So, for instance, our moral beliefs might count as bearing an appropriate relation to privileged beliefs if they were entailed by those beliefs. Thus, if the fact that something was pleasant entailed that it was morally good, or if the fact that some course of action maximized happiness

¹⁸ Gilbert Harman formulates the argument especially well in *The Nature of Morality* (New York: Oxford University Press, 1977). See also my "Moral Theory and Explanatory Impotence," *Midwest Studies in Philosophy*, Vol. XII (University of Minnesota Press, 1988), pp. 433-457

entailed that it was morally right, then as long as we were justified in believing certain things were pleasant or certain courses of action maximized happiness, our beliefs concerning what was morally good or right could be justified.

Yet Moore's observation that there is always an open question when it comes to whether something that satisfies a non-moral description (as, for instance, being an instance of pleasure) satisfies as well a moral description (as, for instance, being something morally good), renders contentious (to say the least) the claim that there are entailment relations between the content of non-moral beliefs and the content of moral beliefs.¹⁹

Inductive (as opposed to explanatory and entailment) relations seem no better placed to provide justifications for our moral beliefs. After all, inductive relations supporting particular moral beliefs will be present only if we already have on hand, as justified, some moral beliefs that we can use as the basis for the induction. And that means the problem simply shifts back to those beliefs, leaving us unable to make a legitimate induction on their basis unless they are themselves justified (in some way other than by induction, if we are to avoid a regress).

However, if (as it appears) explanatory, entailment, and inductive relations exhaust the relevant relations that moral beliefs might bear to other beliefs, there seems no hope for the justification of our moral beliefs unless some of them count as epistemically privileged – as self-evident, or indubitable, or in some other way such that we are both justified in holding them and able to rely on them in justifying other moral beliefs. And, indeed, virtually all foundationalists, whether they are skeptics or not, end up taking this view. (The skeptics, of course, deny that any moral beliefs qualify, while the non-skeptics maintain that some do.) The challenge facing non-skeptics, at least those who accept foundationalism, is to defend the idea that some moral beliefs qualify as epistemically privileged.

There are some moral beliefs that might immediately come to mind as not possibly false, say that if someone has committed a murder that person has done something morally wrong. If, as one might reasonably hold, to commit a murder is to wrongfully kill someone, then the belief could not possibly be false. But it would be compatible both with no case of killing being a

¹⁹ Although this argument has been tremendously influential, there are reasons to worry that it is less compelling than it has often been taken to be. That there is an open question in these cases does establish, it seems clear, that a person might be competent with the relevant terms and yet not know that the applicability of one (set of) term(s) entails the applicability of the other. But that might be true even if the applicability of one entailed the applicability of the other. And it might be true even if, more dramatically, the content expressed by the terms was the same, if in at least some ways the content might not always be directly available to competent users of the term.

murder (if killing is never wrongful) and with every act of killing being murder (if killing is always wrongful). What guarantees its truth – its being an analytic that murder is wrong – seems to guarantee as well its triviality. In order for a belief of this kind to be of help in justifying other moral beliefs, we would need to have on hand justified substantive beliefs, for instance, to the effect that a certain act was a wrongful killing, or that it was a case of murder. But the justification of these beliefs is just what is at issue, and the (trivial) truth of various analytic claims seems to be of no help with them.

One might argue that the available analytic truths are much more surprising and complicated than this example suggests. And one might hold too that discovering their truth is a kind of revelation even to those who are sophisticated in their moral thinking.²⁰ So it is worth noting that surprising or not, complicated or not, a revelation or not, the availability of analytic moral truths would still leave us with, as of yet, no justification for specific substantive beliefs -- unless the terms of the available analysis are all non-moral ones that we had some justification for applying to specific cases. For this reason, the ambition of those who ground moral epistemology in analytic truths involves finding successful analyses that trade in epistemically unproblematic non-moral terms. Only then, relying on the analyses, and relevant justified non-moral beliefs, might the moral beliefs in question inherit epistemic credentials.²¹

If, for instance, a correct analysis reveals that something is morally wrong if in doing it one is acting contrary to what one's ideal advisor would recommend, and if we can have a justified view of what one's ideal advisor would recommend, then we could have a justified belief about what is wrong. The challenge is to show that the terms of the analysis are epistemically tractable. However, skeptics about morality, if convinced by the analysis, are likely to be as skeptical about beliefs concerning the recommendations of ideal advisors as of the original moral claims. That of course might turn out not to be true, though. There really is room to hope that a convincing analysis will be illuminating in a way that shifts the skeptic, in light of her lack of skepticism concerning the terms of the analysis. Even here, though, we will need an

²⁰ Frank Jackson defends such a view in [From Metaphysics to Ethics](#) (Oxford: Oxford University Press, 1998).

²¹ Along the way, of course, those who take this route will be rejecting the idea that Moore's Open Question Argument shows that no analysis of moral claims, in non-moral terms, could be successful. There are, fortunately for them, a number of reasons for suspecting that the Open Question Argument, as it is usually interpreted, over estimates the ability of those competent with a term to see its implications, and in various ways presupposes an over simple understanding of our concepts, moral and non-moral alike. See my "'Good' on Twin Earth" in [Philosophical Issues](#), Enrique Villaneuva, ed. (Atascadero: Ridgeview, 1997) for a discussion of some of the problems with Moore's argument.

account of what justifies our views concerning, in the case we are imagining, the recommendations of an ideal advisor.

Alternatively, one might, without appeal to any putative analytic truths, defend some substantive moral belief on the grounds that it is *a priori*, even if not analytic. Thus, for instance, one might maintain that an action is wrong if and only if the agent could not consistently will that everyone act on the same grounds she is taking as sufficient for her.

Against such views, a moral skeptic presses two challenges. The first challenge is to show that the claim is, in fact, an *a priori* truth.²² The second is to show that our beliefs concerning what an agent might consistently will are on a better footing, epistemically, than the moral claims concerning an action being wrong seemed to be.²³ Predictably, bearing each of these burdens is a real challenge.

Hoping to meet both burdens, moral intuitionists have argued that there are some substantive moral claims that all who are morally competent, clear-headed, and unbiased, can be brought to recognize as true. Moreover, they maintain, the credentials of these claims are comfortably on a par, epistemically, with the rules of logic and the axioms of mathematics. As the intuitionists emphasize, the rules of logic and the axioms of mathematics admit of no more proof, nor do people think they need more proof, than is found in the fact that those who are logically or mathematically competent, clear-headed, and unbiased, can “see” their truth. The reflective faculties of humans, no less than their sensorial faculties, are, when properly used, sources of information about the world.²⁴

²² Kant, of course, tries to meet this challenge with a transcendental argument. See Immanuel Kant’s Groundwork for the Metaphysics of Morals. Or one might argue that we are, in some way, committed by our nature as practical agents to certain substantive truths. Of course such arguments, even if successful, are not really arguments that the beliefs in question are true, but only arguments that we must (in some sense of “must”) think that they are. Alternatively, one might take a cue from recent work on *a priori* truths concerning, for instance, water being H₂O, and argue that, similarly, we have grounds for thinking certain substantive moral claims are true *a priori*.

²³ Of course, if this putative *a priori* truth is true, and a person knows that it is, then presumably moral claims concerning an action being morally wrong will inherit the justification enjoyed by claims concerning what can and cannot be consistently willed. But the advantage of having this *a priori* truth in hand is that our justification for our beliefs concerning what can be consistently willed is stronger than would be any justification we have for thinking an act right or wrong, absent the availability of the *a priori* truth.

²⁴ See, for instance, W. D. Ross, The Right and the Good (Oxford: Oxford, 1930). See also Michael Huemer’s Ethical Intuitionism (New York: Macmillan, 2005).

Thus, for instance, they suggest that we are justified in believing that, other things equal, if we have made a promise we ought to keep it, and that (again) other things equal, if we are in a position to help someone in need, we ought to help. We are justified in believing these things not because they help us explain our experiences, nor because they are analytic truths, but because, on reflection, if we are morally competent, clear-headed, and unbiased, we will come to see their truth. And once such claims are on hand, the intuitionists note, we are in a position to justify other moral beliefs by appeal to these. So while most of our moral beliefs are not themselves directly intuitive – our moral world is far too complex for that – moral intuition provides the fundamental truths that we can use to justify our other moral beliefs. There is, moral intuitionist acknowledge, plenty of room to get things wrong, in the first instance if we are not morally competent, or not clear-headed, or not unbiased, and in the other instances if we mis-weigh the relevant considerations, misunderstand the facts, or let ourselves think other things are equal when they are not (or vice versa) in applying the intuitively available substantive truths of morality.

Moral skepticism finds a great deal of support in the fact that, seemingly, for every substantive moral claim that might be advanced as an *a priori* (even if not analytic) truth, available to the intuition of some, there are a large number of incompatible claims that strike other people as just as obviously true and just as intuitively available to those who are morally competent, clear-headed, and unbiased. Indeed, what strikes people as intuitive, when it comes to morality, so often reflects how they happen to have been raised, and what their interests are, that there is a persistent, and seemingly well grounded, suspicion that what gets valorized as available to intuition is really simply what one is unwilling to challenge.

More generally, moral skeptics press, there are no substantive moral claims that are unchallengeable and so none that are not in need of some evidence or support. Yet that suggests, they hold, that no moral beliefs are properly seen as epistemically privileged in a way that would allow them to serve as a foundation for our other moral beliefs. And if no moral beliefs are foundational (i.e. appropriately epistemically privileged), and no set of exclusively non-moral beliefs can alone provide a justification for our moral beliefs, then no moral beliefs are justified at all. This argument holds even if some of our moral beliefs are, in fact, true. The problem is not that there are no moral truths, but that if there are we apparently have no reason whatsoever to think that we have got them right.

Putting things this way, brings us back to the explanatory problem. Unless the truth of our moral beliefs figures, at the very least, in an explanation of why we hold the beliefs we do – unless we have grounds for thinking that our moral beliefs have been and are responsive to the (moral) truth – we have no reason to think our moral convictions are not simply free-floating from the moral facts (assuming there are such facts).

Worries that our moral convictions do free-float are amplified by the extent to which psychology, sociology, political science and evolution, with appeals to self-interest, false consciousness, political influence and genetic advantage, seem in some combination or other to offer a compelling explanation of our moral views all, it seems, without ever giving us reason to think any of the views in question are true.²⁵

In order to establish that our moral beliefs are responsive to the moral facts we must justify our view that the facts are one way rather than another. Eschewing a direct appeal to intuitions, and often, along with that, the hope of establishing that some substantive moral claims are epistemically privileged in the way a foundationalist demands, those who reject skepticism offer a variety of arguments in favor of some substantive position or another. Thus, for instance, Kant offers a transcendental argument for the categorical imperative, while Mill defends hedonism on the grounds that our view of the desirability of our own pleasure commits us to acknowledging anyone else's like pleasure as equally desirable (even if it is not equally desired).²⁶ Others appeal to our capacity to engage in practical reason or to some other inescapable feature of our nature arguing that it commits us to accepting some principle or other.²⁷ Still others find arguments for substantive moral claims in our need for shared standards for cooperation, or our commitment to mutual respect, or in the capacity of certain standards to explain and justify our more particular moral judgments.²⁸

But of course in offering arguments for some moral view, as part of justifying our view that the moral facts are one way rather than another (as, in turn, part of showing that our beliefs are responsive to the truth), we are supposing that the evidence we offer in our arguments is sensitive to the truth. A skeptic might be forgiven for suspecting that none of these defenses of one's moral beliefs can work without at least surreptitiously assuming that some views one holds are right.

So far, the discussion has gone along against the background of the foundationalist's view of what it would take for a belief to be justified: it must either be a self-evident (or otherwise epistemically privileged) belief, or be justified, ultimately, by appeal to some such belief. And the worry is that, first of

²⁵ These challenges to moral knowledge go back at least to Thrasymachus, in Book I of Plato's Republic, and they have been given elegant expression through the ages. Nietzsche's The Genealogy of Morals is especially noteworthy here.

²⁶ See my "Mill's 'Proof' – A More than Half-hearted Defense" in *Social Philosophy and Policy* (2001) for a discussion of the extent to which Kant's argument for the second formulation of the categorical imperative is structurally isomorphic to Mill's notorious, and I claim, underappreciated, "proof" hedonism.

²⁷ See, for instance, Christine Korsgaard's Sources of Normativity (New York: Cambridge University Press, 1996)

²⁸ See, for instance, David Gauthier's Morals By Agreement (Oxford: Oxford University Press, 1986), Thomas Scanlon's What We Owe To Each Other (Cambridge, MA: Harvard University Press, 1998), and John Rawls' A Theory of Justice (Cambridge, MA: Harvard University Press, 1971), respectively.

all, no moral beliefs that are suitable substantive seem to qualify as self-evident (or otherwise epistemically privileged) and, second, if there are no such moral beliefs, then the fact that some moral beliefs must be appealed to in order to justify any moral beliefs will mean that no moral beliefs at all are justified.

Interestingly, according to many, moral beliefs are not the only ones that seem to fall short in this way. In fact, they suggest, any attempt to establish that contingently true beliefs are appropriately sensitive to the facts they concern, will implicitly or explicitly appeal to other beliefs about those very facts. After all, it is only if we have those facts right that we can show that our beliefs are appropriately sensitive to the facts. But we are then presupposing that (some of) our beliefs are appropriately sensitive to the facts when we try to establish that they are.

Those who embrace global (unrestricted) skepticism often embrace this point and rely on it as part of their grounds for thinking that virtually all of our beliefs are unjustified and so that we lack knowledge not simply of morality but of the external world, of other minds, even of our own experiences. At the same time, many who reject global (unrestricted) skepticism embrace this point too and then take it as grounds for thinking that in some way or other the foundationalist has misunderstood what is required for justification. In particular, they reject the idea that a belief can be justified only if it is either self-evident (or otherwise epistemically privileged) or justified by appeal, ultimately, to such beliefs. Instead, they maintain, we can succeed in justifying a particular belief, say concerning the taste of a cup of coffee, by appeal to another belief (say, concerning how it tasted before), without supposing we can get a perspective on our beliefs that allows us, from outside (so to speak), to show that our beliefs are responsive to the facts. Such a perspective is, in principle, unavailable, they argue, so if there is justification at all, ever, it must not require our having succeeded in taking this perspective.

With this in mind, coherentists about justification argue that just how justified our beliefs are – moral beliefs, no less than beliefs about the taste of coffee or the effects of penicillin or the fuel economy of a car, or whatever – is dependent on, and a reflection of, how they are related to the other things we believe, with there needing to be in the mix no epistemically privileged beliefs of the sort foundationalists hold are necessary for justification.

This suggestion meets with a familiar objection that is obviously right in the point it makes, but that, in the process, reveals a common misunderstanding of what the coherentist holds. The objection is that the relative coherence of a person's beliefs, the extent to which they happen to hold together, is no evidence that they are true. Indeed, for any set of coherent beliefs, it seems easy to imagine there being an alternative set that is just as coherent, yet incompatible. A piece of fiction need be no less coherent than a historical treatise. That means that one cannot reasonably offer, as evidence of the truth of one's beliefs, that they are coherent. Coherence is not evidence of truth (even if certain kinds of incoherence, especially inconsistency, are evidence of falsity). The misunderstanding here is in thinking that a coherentist is offering coherence as evidence of truth. She isn't, or at least shouldn't be. The coherentist's proposal is, or should be, that (relative) coherence is not evidence of truth, it is a

measure of the strength of the evidence the person in question has. The idea is that each person's evidence for whatever she might be considering is found in the contents of her other beliefs. The more the contents of her beliefs support some claim, the more coherent the claim is with her beliefs and the more justified she would be in believing it.²⁹

In characterizing our epistemic situation, coherentism takes as given, but not as privileged, whatever beliefs we find ourselves with. It then sees the process of testing, refining, and justifying our beliefs as a matter of seeing how the various things we believe hang together, putting ourselves in a situation to gather more information (which becomes available as we form new beliefs), and refining and rejecting various claims in light of our evidence. As the skeptic has pointed out, none of the beliefs count as unchallengeable. Yet, contrary to the skeptic's claim, in the face of challenges a person is justified in appealing to, and indeed has no recourse but to appeal to, whatever she does and might come to believe.

This of course leaves open the possibility that a person's justified beliefs, even those that are extremely well justified, might turn out to be false. Justification is no guarantee of truth on the coherentist's view. But that does not mean, the coherentist insists, that a person had no reason to think of her justified beliefs that they are true – she has all the reason provided by her other beliefs. Each of these may be challenged, but each is also in principle something that might meet the challenge.

While coherentism is a general theory of justification and is meant to account for how our scientific theories and our common sense convictions might be justified even though they are neither self-evident (or otherwise epistemically privileged) nor supported by beliefs that are, the account applies directly to moral beliefs as well. To say that is not, of course, to say that our moral beliefs will, on the coherentist account, turn out to be justified. That cannot be determined a priori and depends on actually seeing how our moral (and other) beliefs hang together, putting ourselves in a situation to gather more information, and refining and rejecting various claims in light of our evidence. Coherentism makes sense of how we might intelligibly address the question of whether our moral beliefs are appropriately responsive to the facts, without thereby ensuring that they do. Far from it.

Part of what makes moral skepticism a challenge especially worth facing is that as we consider the full range of what we take ourselves to know – about how moral convictions are shaped and changed, about how the world operates more generally, and about how we might gather information effectively, we find grounds for suspecting that our moral beliefs do not cohere well with our other commitments. At the same time, our moral convictions are so central to our understanding of ourselves and others that coming to grips with them, either as ultimately justified beliefs or as unjustified beliefs or convictions not properly seen as beliefs at all, is a challenge we cannot reasonably ignore.

²⁹ For an extended defense of this version of coherentism, as a general theory of justification and as applied to moral beliefs in particular, see my "Coherentist Epistemology and Moral Theory," in Moral Knowledge? edited by Walter Sinnott-Armstrong and Mark Timmons (Oxford University Press, 1996), pp. 137-189.