

Epistemology, Moral

Geoffrey Sayre-McCord

Moral epistemology, as a field, is concerned with (i) whether, (ii) how, and (iii) what (if anything) we know about morality, about right and wrong, justice and injustice, virtue and vice.

Generally, people seem to have high confidence in at least some of their moral opinions, a confidence that goes with thinking one knows right from wrong, good from bad, etc. Understanding what grounds people might have for their confidence is a central challenge in moral epistemology.

Despite people's confidence in their moral views, developing a plausible account of how they might know what they think they do is a real challenge, one made all the more pressing by the familiar, and disturbing, fact that there is such dramatic, deep, and apparently irreconcilable disagreement about morality. In fact, there are substantial arguments for thinking that, when it comes to morality, there is really nothing to know, because there are no moral facts, and that even if there were such facts, we would have no way to discover them, so that, when it comes to moral knowledge, we should recognize that people do not have it (*see* REALISM, MORAL).

Of course, the confidence we might have in our own moral views will weigh against this conclusion. But our confidence, standing alone, and especially when others have equal confidence in opposing views, is not really evidence that we know what we think we do.

Models for Moral Knowledge

Those who think we do have moral knowledge appeal to a variety of models to understand its nature.

Thus, for instance, some see moral knowledge as rightly understood on the model of perceptual knowledge (*see* A POSTERIORI ETHICAL KNOWLEDGE). They maintain that our capacity to know what is moral, and to distinguish virtue from vice, is much like our capacity to know colors, and to distinguish blue from red. Just as we know about color thanks to having vision, we know about morality (people argue) thanks to our having a moral sense. On this model, if we are in the proper circumstances, we can "see" that causing pain is bad and that helping others is a virtue. It is because we have a moral sense – a conscience – that we are able to distinguish right from wrong, and are in a position to learn about morality (Hutcheson 1725; McNaughton 1988; McGrath 2004).

Alternatively, some suggest that moral knowledge is better understood as analogous to mathematical and logical knowledge, accessible not because of some special sense, but because of the nature or structure of our reason or the distinctive

powers of intuition (*see A PRIORI ETHICAL KNOWLEDGE*). They maintain that our capacity to know what is moral, and to distinguish virtue from vice, is much like our capacity to recognize that $2 + 2 = 4$, and that certain claims express logical truths. This knowledge, it seems, is not culled *from* experience (as moral sense theorists would have it), but brought *to* experience. It is available to all whose rational or intuitive powers are sufficiently developed, whatever their experience might be (Plato 2004; Kant 1785; Ross 1930).

Others hold that moral knowledge should be understood as similar to our scientific knowledge of, say, unobservable entities and natural laws, or of cultural artifacts and sociopolitical features of the world. On this view, our capacity to know what is moral, and to distinguish virtue from vice, is, as with scientific knowledge of the physical and social world, not primarily knowledge of what is directly perceptible, nor of what can be recognized by intuition or through reason alone. Yet it depends, nonetheless, on both experience and reason since what we end up knowing we infer from experience, using reason (Rawls 1971; Boyd 1985).

Still others recommend seeing moral knowledge as not primarily being a matter of knowing *that* something is the case, but of knowing how to do something. Moral knowledge, on this model, is more like our knowing a language than knowing the color of things around us. It is expressed in how we act and feel in response to the circumstances in which we find ourselves. Thus, our knowledge of morality, like our knowledge of a language, is largely a matter of having acquired a broad, yet distinctive, range of abilities. In both cases, the ability in question may well come with various kinds of knowledge *that* (say, that something is wrong, or that a sentence is ungrammatical), but the core capacities that constitute the knowledge in question do not depend on this additional knowledge. One can respond with courage without knowing that acting otherwise would be cowardice; one can speak grammatically without knowing that speaking otherwise is ungrammatical (Aristotle 1999; Mikhail 2011).

As it happens, most people agree that if there is moral knowledge at all, to a large extent it involves knowing *how* (to respond appropriately to one's circumstances). At the same time, though, a full account of moral knowledge seems to require as well a fair amount of knowing *that*, and for at least two reasons. First, in many people's view, what morality requires is not merely that we do what is right, but that we do it because we recognize that it is right. Morality requires not merely that we do our duty, but that we do it *because* it is our duty. And this, many people think, involves knowing *that* acting in a certain way is our duty. If so, we need an account of that knowledge. Second, in identifying some people as knowing how to act morally, and others not, we presuppose moral distinctions among the various abilities, as better or worse, good or bad, virtuous or vicious, and knowing that abilities fall on one side or the other of these distinctions itself calls for an explanation not provided by the model of knowing how.

So, inevitably, attempts to explain moral knowledge need to explain how it is we might know *that* certain things are right, others wrong, *that* certain traits or habits or reactions are virtuous and others vicious, and *that* some things are more

valuable morally than others. The other models mentioned above (of a moral sense, rational intuition, or theoretical inference) find their role when offering such explanations.

Justification

A primary focus of such explanations is on whether, when, and how we might be justified in holding the moral views we do. What might count as evidence for one position and against another? In the context of moral thought, the idea of justification has two distinct roles. It is important to be clear on the difference.

On the one hand, questions of justification often ask of things (say actions, attitudes, or institutions) whether they are *morally* justified, that is, whether the appropriate moral considerations weigh in their favor or not. Indeed, much of substantive moral theory is concerned with articulating and defending more or less general accounts of when and why, and in what way, various things might meet morality's requirements and so count as morally justified.

On the other hand, questions of justification also often ask of our moral beliefs (including, specifically, beliefs that something is morally justified) whether they are *epistemically* justified, that is, whether the available evidence weighs in their favor. Much of substantive epistemic (as opposed to moral) theory is concerned with articulating and defending more or less general accounts of when and why, and in what way, various beliefs (moral and otherwise) might meet epistemic requirements and so count as epistemically justified.

These different roles for questions of justification reflect quite different dimensions of evaluation. Often, it seems, we might be epistemically justified in thinking something is, say, morally justified, even though (as a matter of fact) it is not; and one can be morally justified in doing something even though one might not be epistemically justified in believing it is morally justified. Still, moral justification and epistemic justification interact in interesting ways. In particular, sometimes it seems as if the fact that one does not have good evidence for one's moral beliefs itself counts against the moral permissibility of acting on them.

In any case, in thinking about moral *epistemology*, epistemic justification is of primary importance. The theories of epistemic justification developed in other contexts – say in accounting for when and why we are justified in trusting our perceptions, or, more generally, our experience, or justified in relying on our intuitions or on our reasoning – carry over to thinking about the epistemic justification of our moral beliefs. Predictably, the general positions and challenges that emerge concerning epistemic justification in these other contexts all have their place in thinking specifically about the standing of our moral views.

Thus, for instance, one might hold that the structure of justification requires an unshakable foundation. Those who hold this view, epistemic foundationalists, hold that we cannot be justified in believing anything unless at least some of our beliefs are certain, or indubitable, or infallible, or self-evident. If this is right, then defending the possibility of our moral beliefs being justified requires showing either that some

such beliefs enjoy the required privileged status or that they are properly related to beliefs that do (Huemer 2005).

Alternatively, one might reject the need for such foundations and defend the idea that justification is a matter of how one's beliefs fit together. Those who hold this view, coherence theorists, hold that a person's belief stands as justified (or not) in light of the relations it bears to the person's other beliefs. If this is right, then defending the possibility of our moral beliefs being justified requires showing that they might actually stand in the relations that are required either to each other or to our other beliefs (Sayre-McCord 1996; *see* COHERENTISM, MORAL).

Or, to consider a third approach, one might hold that what justification requires is not centrally a matter of relations to privileged beliefs (as foundationalists hold), nor even a matter of what relations hold among beliefs (as coherentists hold), but is instead a matter of being appropriately responsive to the facts (as reliabilists hold). If this is right, then defending the possibility of our moral beliefs being justified requires showing that they might be responsive to the moral facts in the required way.

Each of these views of justification carries specific burdens for those hoping to show that some moral beliefs are epistemically justified. And they each play out against a broad set of considerations that shape how we might best understand what is involved in our having good evidence for our beliefs, whether those beliefs are perceptual, mathematical, logical, scientific, or otherwise.

Different theories of justification, and different models of knowledge, go with different accounts of what counts as evidence for, or justification of, our moral views, and with different accounts of when and how we might acquire moral knowledge. Moreover, they go with different understandings of what is known, when someone has moral knowledge.

Challenges

Models and analogies are all well and good, and when it comes to moral knowledge, each of these models seems to get at something important about how we come to our moral views. Yet, their value, when it comes to moral knowledge, depends on our being able to work out, in some detail, how, in the proffered model, we come actually to know something about morality.

Just how hard this is depends, in part, on what it is that we are supposed to know when we have moral knowledge. If moral facts are of a piece with, say, perceptual, or mathematical, or sociological, or psychological, or physical, facts, then (on the perceptual, mathematical, or scientific models) the details relating to our knowledge of morality are the very same details involved in accounting for our knowledge of these facts (*see* NATURALISM, ETHICAL). Moral knowledge would pose no special problem. However, many think that moral facts are (or would have to be, if there were any) distinctive in various ways that mean they are not *simply* "of a piece" with these other, more familiar and epistemically tractable, facts.

Thus, for instance, virtually no one holds that value (as opposed to the things that are valuable) is literally visible, or tangible, or audible, or that value itself has a taste or an odor. Even those who think that perception offers a good model for moral knowledge are overwhelmingly likely to hold as well that our “moral sense” is distinct from our other senses and – crucially – that what is sensed by it is different from what is available to the other senses.

Similarly, virtually no one holds that when we have moral knowledge what we have is merely a case of mathematical or logical knowledge, even if they think we come to know moral truths much as we know truths of mathematics and logic. Morality is not mathematics, nor merely a matter of logic, even though numbers might matter and logic constrains what morality might demand. Consequently, those who embrace our knowledge of mathematics and logic as their model for moral knowledge treat what we know as distinctive in important ways.

In contrast, and moving to the third model, some people do think it is plausible to hold that what we know, when we have moral knowledge, is some empirically accessible fact. So, for instance, cultural relativists have maintained that an action is morally right if, but only if, it is in accord with the norms of one’s society. On this view, discovering that an action is right is a straightforward sociological discovery, to be understood on a par with other such discoveries, and to be explained in exactly the same way. Or, to take other examples, hedonists and preference theorists hold that discovering something has value is discovering, respectively, that it produces pleasure or that it satisfies someone’s preferences. On either suggestion, discovering the value of something is a matter of making an empirical discovery that is utterly explicable as a piece of scientific knowledge. Indeed, one of the main arguments offered in favor of these accounts is that they remove the mystery from moral knowledge (Perry 1926; Brandt 1979; Railton 1986).

Of course, while these views make moral knowledge tractable, they invite the question “how do we know that rightness is simply a matter of conforming to the norms of one’s society, or that value is a matter of causing pleasure or, alternatively (and incompatibly), of satisfying preference?” Is this (putative) knowledge itself a bit of empirical knowledge?

Many think not, and for two related reasons. The first traces back to David Hume and his insistence that one can never derive an “ought” from an “is” (*see IS–OUGHT GAP*). As he saw it, whatever we might discover about the way the world is nonmorally, judgments then made concerning what ought or ought not to happen do not follow directly, but presuppose some moral principle to justify the inference.

In every system of morality, which I have hitherto met with, I have always remarked, that the author proceeds for some time in the ordinary ways of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when all of a sudden I am surprised to find, that instead of the usual copulations of propositions, *is*, and *is not*, I meet with no proposition that is not connected with an *ought*, or an *ought not*. This change is imperceptible; but is however, of the last consequence. For as this *ought*, or *ought not*, expresses some new relation or

affirmation, 'tis necessary that it should be observed and explained; and at the same time that a reason should be given; for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it. (Hume 1978: 3.3.1; *see* FACT-VALUE DISTINCTION)

Making a closely related point, G. E. Moore (1903) offered what has come to be called the “Open Question Argument,” which was meant to show that it is always a mistake to equate moral terms (say “good”) with empirical terms (say “pleasant” or “preference satisfying” or “the object of a desire one desires to have”). Any such equation, Moore argued, must be abandoned in light of the observation that, for each such proposal, it is always intelligible to ask of what is, in fact, pleasant, or such as to satisfy someone’s preference, or the object of a desire one desires to have, whether it is actually good. As Moore saw things, if “morally good” actually referred to what was pleasant (as it would if “morally good” meant “pleasant”), then there would be no room to wonder, as there seems to be, whether the mere fact that something is pleasant establishes that it is morally good. So there would be no room to wonder, as there seems to be, whether sadistic pleasure is valuable. Wondering such things would be like wondering whether a triangle has three sides; it would show that one doesn’t understand either “triangle” or “has three sides.” However, wondering whether something pleasant is morally good does not show that one doesn’t understand either “good” or “pleasure” (Moore 1903; *see* OPEN QUESTION ARGUMENT).

The idea is not that being morally good and being pleasant are unrelated. For all the argument might show, we could discover that all and only pleasures are good. However, if we did, Moore held, it would constitute a substantive discovery concerning either pleasures – that they are all, actually, good – or good things – that they are all, actually, pleasant. Such discoveries would not be a reflection of the words “pleasure” and “good” having the same meaning nor of pleasure and goodness being the same properties (*see* AUTONOMY OF ETHICS).

But then, in whatever way we learn of the nonmoral facts concerning pleasure, social sanctions, the causal effects of various actions or policies, or the nature of God’s will, we need an additional account of our knowing the connection between such facts and the status of something as (for instance) right or good or virtuous. This suggests that the scientific model of knowledge, unless it is in some way extended, cannot account for moral knowledge.

How the scientific model, or any of the other models, might be extended to account plausibly for our knowledge of morality is one of the main challenges facing moral epistemology.

On virtually all views that see us as actually having moral knowledge, a crucial first step involves taking the content of (some of) our moral convictions or “intuitions” at face value. Thus, we might suppose that killing innocent people for entertainment is wrong, or that we have an obligation to keep our promises, or that facing danger for a good cause is a virtue, or that respecting others is morally required, or that we ought to treat others as we would have them treat us. Standardly, moral claims that strike

people as obviously true play two roles: they are taken to be something we do know, and they are used as evidence to support other moral claims. Without some such claims, we would not be in a position even to start wondering about whether we have moral knowledge. And once we do start wondering, it seems that without them we would have nothing that would count as evidence for any substantive moral view at all (*see* INTUITIONS, MORAL). Needless to say, different models of moral knowledge and of justification end up giving such claims different roles. Sometimes they are treated as having the kind of certainty that allows them to serve as epistemic foundations (*see* INTUITIONISM, MORAL). Other times they are offered simply as the most likely to be true, if any are. However such beliefs are treated, they end up figuring crucially in any account of the epistemic justification of our moral beliefs.

Put this way, of course, things might sound problematically circular. If we only get a justification of our moral beliefs by relying on some of them, are we not illegitimately simply dealing ourselves a victory?

But this is too quick. Every candidate domain of knowledge – of the physical world, or other minds, or mathematics, of biology, of society – relies unavoidably on claims within the domain as both what is to be justified as known and as what might serve as evidence for other things we might know about the domain. The mere fact that we have to rely on our observations of the physical world in order to develop and justify our theories of the physical world does not by itself show that those theories do not constitute knowledge. Similar points hold for the other areas. In each case, it seems, the process of moving from the content provided by our initial observations, convictions, or intuitions to justified beliefs is a matter of seeking a “reflective equilibrium” between what we have at the start and what we are able to relate systematically to other things we believe or (as it seems to us) discover (*see* REFLECTIVE EQUILIBRIUM).

Still, there is no escaping worries about these “initial inputs” (as well as the further “discoveries”), and the process of reconciling them, that raise questions about whether they might just be the start of elaborate fictions with no basis in fact.

When it comes to morality, these worries are amplified by the nature and depth of moral disagreement. Virtually any moral conviction or intuition that one person finds obvious, another is prepared to deny as misguided.

Of course, mere disagreement does not by itself constitute a serious obstacle either to justification or to knowledge. In many cases, disagreements (about, say, evolution, the effects of vaccines, or the impact of different economic institutions) look to be resolvable, *at least in principle*, by appeal to evidence, argument, or reason. Such disagreements might in fact resist resolution, but only because someone party to the disagreement (and perhaps all parties to it) lack relevant evidence, or refuse to acknowledge the force of the arguments, or fall short in their reasoning.

Moral disagreements, in contrast, strike many as being *in principle irresolvable*. No amount of evidence, argument, or reasoning, many think, work to settle moral disagreements about war, abortion, human rights, etc., even in principle. Yet, if this is right, many hold, the different sides to the various debates are in no position, however strongly they might hold their opinions, to claim knowledge. If the evidence the arguments, and the relevant reasoning, all fail to settle the matter, *even*

in principle, then (they hold) no one has evidence, argument, and reason on their side in the way knowledge requires.

Again, the point is not that knowledge is incompatible with actual disagreement, nor that knowledge requires that one have in hand the decisive evidence, argument, or reasoning that is in principle available. Rather, the idea is that in matters where, even in principle, there is no decisive evidence, argument, or reasoning, there is no knowledge. Needless to say, there is room to wonder just how extensive or deep moral disagreement really is, and room too to hold that whatever disagreements there are might actually be, at least in principle, resolvable. But if (as it seems) there are extensive and deep moral disagreements, and they are not even in principle resolvable, many conclude that moral knowledge is simply not available (*see* DISAGREEMENT, MORAL).

Moreover, even if one leaves to one side the question of whether moral disagreements are in principle resolvable, the fact that people who are apparently epistemic equals disagree seems to provide grounds for them all to reduce the credence they give to their own views. After all, if others equally able and informed come to different conclusions, that seems a reason to wonder whether one's own conclusion is as well supported as one thought (Kelly 2005; Christensen 2009).

Moral disagreements challenge moral knowledge in another way as well. If we look honestly at when and why people disagree morally, it seems as if the best explanation of people's moral views, and of why they disagree when they do, is found in a complex combination the norms of their society, their personal histories, their emotional responses, and perhaps more deeply in the evolutionary advantages that these provide (*see* EVOLUTION, ETHICS AND). Assuming, as seems plausible, that these norms, histories, responses, and evolutionary advantages are not, in turn, explained by morality, there seems no reason to think people are actually responding (albeit with differing success) to some distinctively moral dimension of the world, even indirectly. It looks as if our moral views would be what they are, regardless of whether or not the world has a moral dimension. Our moral beliefs seem to free-float from what they are putatively about.

Put another way, people's moral views seem to be reflections of the nonmoral features of the world they find themselves in, not responses to the moral facts those views purport to be about. As a result, it seems we have no reason to think our moral views are sensitive to the moral facts, even assuming there are such facts (Harman 1977; Joyce 2001; Mackie 1977; Street 2006). Importantly, this concern would remain in place even if there were widespread and deep moral agreement. Although the problem might arise naturally when we ask why people disagree in the ways they do, it persists regardless of whether people disagree, as long as the best explanation of their agreeing need make no appeal to their responding to some distinctively moral dimension of the world. The underlying concern is that if moral facts play no role in explaining our moral convictions and intuitions, the considerations we appeal to in support of our moral views will, it seems, have no claim to being evidence (Sturgeon 1985; Sayre-McCord 1988).

The assumption that there are moral facts, an assumption that comes naturally with our having confidence in our moral views, is itself called into question by these

considerations. For all we know, it appears, morality might be a figment of our collective imagination. Indeed, it looks as if our moral beliefs might simply be about a culturally cultivated fiction.

Such a fiction would no doubt serve many people's interests, working to help keep order and to shape behavior that would be different in the absence of moral convictions. Yet, those effects, for better or worse, do not require that the beliefs be true. What seems to matter is just that people have the beliefs.

Given the many advantages of having effective methods of social organization and control in place, and the obvious ways in which having people believe in morality might serve various people's interests, it is no surprise that parents, preachers, and teachers all (usually with complete sincerity) cultivate moral convictions and encourage a belief in morality's authority. But, as predictable and explicable as a belief in morality might be, we have in these considerations and explanations no evidence that the convictions are true or the authority real (Sinnott-Armstrong 2006).

This line of argument raises deep challenges to the idea that we can reasonably rely on our moral convictions and intuitions in justifying our other moral beliefs. After all, if there are good grounds for suspecting that we would have the convictions or intuitions we do, regardless of whether they are accurate, it seems that relying on them in justifying our other beliefs (or our actions) is relying on something we have no reason to think true. However, if we are barred from relying on these, it seems we have no resources whatsoever for coming to an informed view about what morality requires.

Worries are compounded when we reflect on what morality is supposed to be like, and in particular its supposed authority and a kind of objectivity that allows it to apply to all. The law of gravity applies equally to all and, importantly, forces compliance, but the moral laws, which are also supposed to apply to all, are thought (by many) to have an authority that demands, or calls for, compliance, even in the absence of enforcement. Just what this authority might be, and how the laws of morality might come to have it in a way that would allow them to apply to all, is a mystery.

There might be a way to make sense of the authority of morality, many suggest, if that authority were tied to people's desires or aims, but then it is unlikely to be an authority that backs the same requirements for all. Alternatively, there might be a way to make sense of the morality applying to all, many suggest, if it were abstracted from all that makes individuals different from one another, but then it is unlikely to consist in requirements that had authority over all. Now, of course, there might actually be a way to combine authority and objectivity (and many have tried). And there might be reason to mitigate one or the other that would render them more obviously compatible. But, according to many, morality's pretensions are to a combination of authority and objectivity that nothing can have.

Many find, in these considerations, decisive grounds for thinking that moral knowledge is an illusion. Advancing an error theory, they maintain that in making moral judgments and embracing moral principles, we are making a mistake. We are believing in things – moral facts, or an authoritative system of moral laws, or a set of supernatural commands – that do not exist (*see* ERROR THEORY; PROPERTIES, MORAL). When it

comes to moral knowledge, there is nothing to know, they argue, except that, in holding moral views, we are making a mistake (albeit, perhaps, a salutary mistake that we should be happy to encourage in others and, maybe, in ourselves as well).

Others, though, see the challenges to moral knowledge as suggesting that we rethink what is going on when people have moral convictions. Rather than see us all as forming moral beliefs that involve (thanks to what they presuppose) some deep mistake, they argue, we should see moral convictions not as beliefs (that might be false) but as some other sort of attitude, more akin to emotions or desires that might powerfully shape one's behavior and concern the behavior of others, but which are neither true nor false. Moreover, our moral convictions might well reflect attitudes that we have toward all (in a way that explains apparent objectivity) and feel compelled to comply with (in a way that explains apparent authority) (Stevenson 1944; Ayer 1952; Gibbard 1990; Blackburn 1993; *see* NON-COGNITIVISM; PROJECTIVISM; QUASI-REALISM). On this view, often called "non-cognitivism," none of the models of knowledge mentioned above really fit our moral views, because those models all apply specifically to beliefs, not to emotions, affective, or desiderative attitudes that are not evaluable as true or false. We might, of course, distinguish among the various moral convictions that people have, counting some as knowledge and others not, but in doing so we will (on this view) be expressing our own convictions (in this case about the moral convictions of others). Thought of in this way, valorizing a position as moral knowledge is like thinking an action is right, or a character trait virtuous – it is an expression of our approval or our endorsement or our convictions.

None of this is to say that our moral convictions are arbitrary or misguided. Far from it. However, the considerations that we might appeal to in defending our own convictions will presuppose others (thus harmonizing with Hume's observations about "is" and "ought"), and should not be seen as constituting evidence for the truth of certain beliefs, nor as attempts to get at some independent moral reality that will independently serve as a standard for our convictions. On this sort of view, our thinking about moral knowledge is itself an evaluative matter to be understood on the non-cognitivist model as being, fundamentally, a matter of a certain affective attitude toward the views one valorizes as knowledge.

Prospects

Not surprisingly, many resist the idea that morality is a fiction and reject the idea that our moral convictions are fundamentally just expressions of our non-cognitive attitudes or commitments. The main grounds for this resistance is found in the apparent fact that we do have some moral knowledge concerning, for instance, the wrongness of slavery, the obligation to keep one's promises, and the importance of justice. It is, of course, not sufficient simply to pound the table and insist that this is knowledge. Yet, rejecting the idea that it is genuine knowledge comes at a very high cost to a central aspect of our understanding of the world. As a result, many see reason to challenge, as inadequate, the arguments offered in defense of thinking we lack moral knowledge.

If indeed slavery is wrong, and (as it seems) we are in a position to know it, then a full understanding of what we know will have to include that knowledge. Developing such an understanding involves tackling head-on the challenges described above in a way that (i) provides an intelligible and plausible account of what it is for slavery (or whatever) to be wrong, and (ii) explains how, given that account, we are in a position to know that slavery (or whatever) is wrong.

Doing this successfully requires capturing the distinctive features of morality in a way that shows it to be accessible to creatures like us, with the limited capacities we have to learn about our world. Not surprisingly, there is deep and persistent disagreement as to how such an account should go, even among those who are more convinced of our having moral knowledge than they are of the skeptical arguments. But the reality of this disagreement is not itself an argument against moral knowledge, at least if the disagreement is in principle resolvable by appeal to evidence, argument, and reasoning. So the burden facing those who are not skeptics is to provide the evidence, arguments, and reasoning that should be persuasive (even as we recognize in advance that it will not always actually persuade).

The main strategy, which goes naturally with giving substantial roles to the models with which we began, is to show the extent to which our moral convictions are on the same footing as our views concerning the existence of an external world, of other minds, of the laws of nature, of mathematical principles, or of standards of reasoning. Almost surely, “being on the same footing as” will not be “being the very same as” any of these other kinds of knowledge. Morality is different from mathematics, just as they are both different from other minds and the material world. However, to a surprising extent, as standard skeptical arguments concerning these other areas show, the influential challenges to moral knowledge have their twins concerning these other areas of (putative) nonmoral knowledge. The best prospects for moral epistemology are likely found in showing that our moral views may be no less well supported by evidence, argument, and reason than (some of) our other views that nonetheless have standing as knowledge (Boyd 1985; Sayre-McCord 1996; Huemer 2005).

In pursuing the strategy, those interested in moral epistemology will quickly get involved in, and take on commitments concerning, the nature of knowledge more generally. Following this strategy, as a way of coming to understand how our moral beliefs might fare in the face of significant skeptical challenges, involves developing views of how our perceptual, mathematical, and scientific beliefs sometimes earn their status as knowledge.

In the process, one must sort out the connection, if any, between knowledge, on the one hand, and (among other things) justification, evidence, reliability, truth, and certainty, on the other. Generally speaking, the skeptical challenges, whether against moral knowledge or other kinds of knowledge, work by first identifying what knowledge (were there to be any) requires, and then arguing that, in the relevant area, the requirements are not, in fact, met. Replies then fall into two groups, those that accept the account of what knowledge requires and then argue that those requirements are actually met, and those that argue for different requirements, in the conviction that they are met.

Either way, the hope of those attempting to meet the skeptical challenges to moral knowledge is that successful replies to skepticism in other areas will carry over naturally to morality, despite the admitted dramatic differences in substance between morality, mathematics, mind, and matter.

So far, the strategy has focused on potential parallels between the status of our moral beliefs and that of our perceptual, mathematical, and scientific beliefs, both in trying to account for how these beliefs might earn their epistemic credentials and in understanding how they might resist various skeptical challenges. It is worth noting that a promising alternative – or addition – is to explore the parallels between our moral beliefs (about the value of various actions, characters, and institutions) and our evaluative epistemic beliefs (about the epistemic value of various beliefs, arguments, and theories, and their status as justified or not, as well as about the nature of evidence). These epistemic beliefs are in important ways very much like our moral beliefs, and they invite many of the same skeptical challenges. In particular, our epistemic commitments, no less than our moral commitments, seem to presuppose the presence of evaluative facts (about evidence and justification). As a result, even those who argue that we *should not* believe that we have moral knowledge appear to suppose that there are facts about what one should (and should not) believe. Those who pursue this strategy stress the ways in which defenses of the possibility of moral knowledge can build on however we end up understanding epistemic facts and our knowledge of them (including our knowledge of when and why certain beliefs, say, do not qualify as knowledge, while others do; Putnam 1981; Sayre-McCord 1988; Cuneo 2007; Enoch 2011).

As mentioned above, non-cognitivists might consistently apply their view to claims concerning justification, evidence, and knowledge, treating them all as, fundamentally, expressions of our non-cognitive attitudes, rather than of beliefs that might themselves be straightforwardly true or false in virtue of getting the facts right or not. But the resulting position strikes many as misconstruing the nature of epistemic claims concerning evidence, justification, and knowledge, by failing to do justice to the ways in which such claims purport to report facts and are, as a result, true or false depending on whether they get those facts right.

Alternatively, a thoroughgoing skeptic who rejects the idea that we know anything, including anything about knowledge, will simply sidestep this strategy by avoiding any epistemic commitments whatsoever. Yet, anyone who thinks we do have some knowledge, say of mathematics, or of physics, or of the external world, or (as error theorists hold) of what would have to be true for there to be moral facts, will face many of the challenges identified above. And if there are ways to meet those challenges when it comes to our having knowledge (or even justification) concerning anything, there is reason to think that those ways might provide good models for similar responses to the challenges facing those who think we (sometimes) have moral knowledge.

Admittedly, this is all by way of hopeful prospects, not proven possibilities, let alone accomplished facts. Yet, many think the strategies and prospects suggested do offer the promise of finding satisfying, or at least illuminating, accounts of what

would be required in order for our (often firmly held) moral commitments to count as moral knowledge.

see also: A POSTERIORI ETHICAL KNOWLEDGE; A PRIORI ETHICAL KNOWLEDGE; AUTHORITY; AUTONOMY OF ETHICS; COHERENTISM, MORAL; DISAGREEMENT, MORAL; ERROR THEORY; EVOLUTION, ETHICS AND; FACT–VALUE DISTINCTION; INTUITIONISM, MORAL; INTUITIONS, MORAL; IS–OUGHT GAP; NATURALISM, ETHICAL; NON-COGNITIVISM; OPEN QUESTION ARGUMENT; PERCEPTION, MORAL; PROJECTIVISM; PROPERTIES, MORAL; QUASI-REALISM; REALISM, MORAL; REFLECTIVE EQUILIBRIUM; SKEPTICISM, MORAL

REFERENCES

- Aristotle 1999. *Nicomachean Ethics*, 2nd ed., trans. and ed. Terence Irwin. Indianapolis: Hackett.
- Ayer, A. J. 1952 [1935]. *Language, Truth, and Logic*. New York: Dover.
- Blackburn, Simon 1993. *Essays in Quasi-Realism*. New York: Oxford University Press.
- Boyd, Richard 1985. “How to Be a Moral Realist,” in Geoffrey Sayre-McCord (ed.), *Essays in Moral Realism*. Ithaca: Cornell University Press, pp. 181–228.
- Brandt, Richard 1979. *A Theory of the Good and the Right*. Oxford: Oxford University Press.
- Christensen, David 2009. “Disagreement as Evidence: The Epistemology of Controversy,” *Philosophy Compass*, vol. 4, pp. 756–67.
- Cuneo, Terence 2007. *The Normative Web*. Oxford: Oxford University Press.
- Enoch, David 2011. *Taking Morality Seriously*. Oxford: Oxford University Press.
- Gibbard, Allan 1990. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Harman, Gilbert 1977. *The Nature of Morality*. New York: Oxford University Press.
- Huemer, Michael 2005. *Intuitionism*. Houndmills: Palgrave Macmillan.
- Hume, David 1778 [1739–40]. *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, 2nd ed. Oxford: Oxford University Press.
- Hutcheson, Francis 2004 [1725]. *An Inquiry into the Original of our Ideas of Beauty and Virtue*, ed. Wolfgang Leidhold. Indianapolis: Liberty Fund.
- Joyce, James 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
- Kant, Immanuel 2003 [1785]. *Groundwork of the Metaphysics of Morals*, ed. Thomas E. Hill, trans. Arnulf Zweig. Oxford: Oxford University Press.
- Kelly, Thomas 2005. “The Epistemic Significance of Disagreement,” *Oxford Studies in Epistemology*, ed. J. Hawthorne and T. Gendler. Oxford: Oxford University Press, pp. 167–96.
- Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. New York: Penguin.
- McGrath, Sarah 2004. “Moral Knowledge by Perception,” *Philosophical Perspectives*, vol. 18, pp. 209–28.
- McNaughton, David 1988. *Moral Vision*. Oxford: Blackwell.
- Mikhail, John 2011. *Elements of Moral Cognition: Rawls’ Linguistic Analogy and the Cognitive Science of Moral and Legal Judgment*. Cambridge: Cambridge University Press.
- Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Perry, Ralph Barton 1926. *A General Theory of Value*. Cambridge, MA: Harvard University Press.
- Plato 2004. *Republic*, trans. and ed. C. D. C. Reeve. Indianapolis: Hackett.

- Putnam, Hilary 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- Railton, Peter 1986. "Moral Realism," *Philosophical Review*, vol. 95, pp. 163–207.
- Rawls, John 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Ross, W. D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Sayre-McCord, Geoffrey 1988. "Moral Theory and Explanatory Impotence," *Midwest Studies*, vol. 12, pp. 433–57.
- Sayre-McCord, Geoffrey 1996. "Coherentist Epistemology and Moral Theory," in Walter Sinnott-Armstrong and Mark Timmons (eds.), *Moral Knowledge?* Oxford: Oxford University Press.
- Sinnott-Armstrong, Walter 2006. *Moral Scepticisms*. New York: Oxford University Press.
- Stevenson, Charles 1944. *Ethics and Language*. New Haven: Yale University Press.
- Street, Sharon 2006. "A Darwinian Dilemma for Realist Theories of Value," *Philosophical Studies*, vol. 127, pp. 109–66.
- Sturgeon, Nicholas 1985. "Moral Explanations," in David Copp and David Zimmerman (eds.), *Morality, Reason, and Truth*. Totowa, NJ: Rowman & Allanheld, pp. 49–78.

FURTHER READINGS

- Anderson, Elizabeth 1995. "Knowledge, Human Interests, and Objectivity in Feminist Epistemology," *Philosophical Topics*, vol. 23, pp. 27–58.
- Audi, Robert 1997. *Moral Knowledge and Ethical Character*. Oxford: Oxford University Press.
- Brink, David 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Dancy, Jonathan 1993. *Moral Reasons*. Oxford: Blackwell.
- Daniels, Norman 1979. "Wide Reflective Equilibrium and Theory Acceptance in Ethics," in D. Copp and D. Zimmerman (eds.), *Morality, Reason, and Truth*. Totowa, NJ: Rowman & Allanheld.
- Nietzsche, Friederich 1998 [1887]. *On the Genealogy of Morals*, trans. Maudmarie Clark and Alan Swensen. Indianapolis: Hackett.
- Shafer-Landau, Russ 2003. *Moral Realism: A Defense*. Oxford: Oxford University Press.
- Sinnott-Armstrong, Walter 2006. *Moral Scepticisms*. New York: Oxford University Press.
- Tersman, Folke 2006. *Moral Disagreement*. Cambridge: Cambridge University Press.
- Timmons, Mark, and Terence Horgan 1992. "Troubles for New Wave Moral Semantics: The 'Open Question Argument' Revived," *Philosophical Papers*, vol. 21, pp. 153–75.
- White, Morton 1981. *What Is and What Ought to Be Done: An Essay on Ethics and Epistemology*. New York: Oxford University Press.