The denial of moral knowledge
Richard Joyce

1. Introduction

Suppose we make the assumption that to know that *p* is to have a true and justified belief that *p*. (We all know, of course, that there are problems with this JTB assumption, but for the purposes of this chapter they will be put aside.) On this assumption, there are three ways of denying that a person knows that *p*: one can deny that *p* is true, or deny that the person believes *p*, or deny that the person's belief that *p* is justified. The *moral skeptic* denies that anyone has moral knowledge, and thus there are three forms of moral skepticism: (i) the *error theorist* denies that any moral judgments are true; (ii) the *noncognitivist* denies that moral judgments are beliefs; and (iii) the *justification skeptic* denies that any moral beliefs are justified. This chapter will discuss these three skeptical options in turn, but will not argue for or against any of them. First, some preliminary points will be made.

The skeptical positions outlined concern non-negative first-order moral judgments.

1. Stealing is morally wrong.

is an example of such a judgment, but the following are not:

2. It is not the case that stealing is morally wrong.
3. John believes that stealing is wrong.
4. You are not justified in judging that stealing is wrong.

The error theorist thinks that (1) is untrue, but may hold that (2), (3), and (4) are true. Likewise, *mutatis mutandis*, the noncognitivist and the justification skeptic. What it takes for something to count as a "moral judgment" in the first place is something that different theorists will answer in different ways. Perhaps moral judgments are those that involve the employment of any of a range of certain concepts (*wrongness*, *obligation*, *evil*, etc.), or ontological commitment to any of a range of certain properties (wrongness, obligation, evil, etc.), or the expression of any of a range certain attitudes (disapproval, approbation, the desire to punish, etc.). And various answers are available for how these different kinds of lists might be drawn up. In the case of properties, for example, one might say that moral properties are those that (putatively) place certain authoritative practical demands upon agents.

Moral skepticism is not a form of moral anti-realism. If we accept the traditional view that moral realism is the position that moral judgments are beliefs that are sometimes true, and whose

truth-value is an objective matter (under some to-be-specified understanding of *objectivity*), then skepticism and moral anti-realism may come apart. In short: realism says nothing about justification, and skepticism says nothing about objectivity. One may be a moral realist while maintaining that moral judgments lack justification (thus also being a moral skeptic). Alternatively, one may deny moral skepticism while maintaining that moral truths are entirely subjective (thus also being a moral anti-realist).

Moral skepticism is not a form of moral eliminativism. If we take eliminativism to be the position that we should stop engaging in moral discourse, then skepticism and eliminativism may come apart. Even if moral judgments are all false or all unjustified, the question of whether we should entirely abolish moral discourse remains a live one. A certain kind of moral fictionalist, for example, maintains that we should retain moral discourse even while believing that the moral error theoretic position is true (see Joyce 2001, 2017). Alternatively, one may deny moral skepticism but still hold that moral discourse is, pragmatically speaking, a damaging practice that is best eliminated. (Such a view forces us to treat the notion of *justification* with care. In denying moral skepticism one must hold that moral judgments may be *epistemically* justified; but this is consistent with the eliminativist contention that employing moral discourse is *instrumentally* unjustified.)

2. Error theory

The error theorist denies the existence of moral knowledge by denying the truth criterion of the JTB analysis. The error theorist affirms the belief criterion—holding that moral judgments are beliefs (when we consider them as mental entities) and are assertions (when we consider them as speech-acts)—and may remain silent regarding the justification criterion (i.e., the error theorist may or may not hold that our moral judgments are unjustified). The error theorist usually argues for this denial via a two-step argument: (i) that making a moral judgment ontologically commits the speaker to the instantiation of certain properties, but (ii) that these properties are not actually instantiated.

Perhaps the simplest way of understanding this view is via analogy with a more familiar form of error theory: atheism. The atheist maintains that when people engage sincerely in theistic discourse they commit themselves to the existence of certain entities (Gods, divine providence, post-mortem paradise, etc.) but the world simply doesn't contain these entities; hence, the atheist thinks, theistic discourse suffers from a systematic failure to state truths. (And it is worth noting in passing that the atheist may or may not hold that these theistic beliefs are epistemically justified.) Even more simply: If I assert "The book is blue" then I am claiming that the book instantiates the property of blueness; thus if the book is not blue I have said something false. The difference in the case of moral error theory is that the error theorist argues that the moral properties that are ascribed to things in the world—wrongness, obligatoriness, evil, etc.—are *never* instantiated. A weak form of error theory will hold that this lack of instantiation is merely

a contingent affair; a stronger form will hold that these properties cannot possibly be instantiated. (And it is worth noting in passing that there are analogous weak and strong versions of atheism.)

Let's think a little more, in very general terms, about why someone might endorse atheism. First, such a person does need to have some kind of conceptual grasp of the entities whose existence is being denied. She thinks of God, say, as an entity that is (inter alia) omnipotent, omniscient, and the creator of the universe. Obviously, if a theist were to respond "No, God doesn't have those characteristics at all," then the debate would break down; in order for the theist and atheist to be in disagreement, they must first reach a threshold of conceptual agreement. Having sufficiently grasped the relevant concepts, the atheist then denies that the world contains anything answering to them. Perhaps she looks around and concludes that there is no phenomenon whose explanation requires the existence of God—there are better explanatory hypotheses across the board—and she couples this with some principle of parsimony that permits (or demands) disbelief in anything so explanatorily unnecessary. Perhaps she reflects on the many different world religions and their varying conceptions of God (even while agreeing on the core divine attributes), and the natural question of who's got it right and who's got it wrong (and how this epistemic discrepancy would be explained) arouses in her the suspicion that *none of them* have got it right, because there's nothing to *get* right. Or alternatively perhaps she just finds the characteristics ascribed to God outlandish and utterly far-fetched, and she is committed to a worldview that banishes such weirdness. (Perhaps she even finds one or more of the characteristics contradictory, in which case only the stronger version of atheism is available). Needless to say, the atheist is well aware that billions of people across the world and throughout history have believed in God, but this doesn't particularly move her—she probably has a picture of humans as epistemically vulnerable in this respect.

The picture just sketched of the atheist's likely reasoning broadly matches that of the moral error theorist. First, the moral error theorist needs to have some conceptual grasp of the entities whose existence is being denied: properties like moral wrongness, obligatoriness, and evil. She may then maintain that there is no phenomenon whose explanation requires the existence of these properties—there are better explanatory hypotheses across the board—and she couples this with some principle of parsimony that permits (or demands) disbelief in anything so explanatorily unnecessary. Perhaps she reflects on the many different moral systems in the world, and the natural question of who's got it right and who's got it wrong (and how this epistemic discrepancy would be explained) arouses in her the suspicion that *none of them* have got it right, because there's nothing to *get* right. Or alternatively perhaps she just finds the characteristics ascribed to moral properties outlandish and utterly far-fetched, and she is committed to a worldview that banishes such weirdness. (Perhaps she even finds one or more of the characteristics contradictory, in which case only the stronger version of error theory is available). The error theorist will be well aware that billions of people across the world and throughout history have believed in morality, but this doesn't particularly move her—she probably has a picture of humans as epistemically vulnerable in this respect.

Most of the aforementioned argumentative moves can be found in the work of John Mackie, who coined the term "error theory" (1977). First, Mackie has to establish the conceptual step of the argument. He maintains that our moral conceptual framework commits us to the existence of "objective values" and "objective prescriptions": actions that *must* be done, whether we like it or not—where this categoricity is not a human construction but rather something that supposedly transcends human institutions. According to Mackie, we think that murdering innocent people is wrong (for example) not because it undermines the perpetrator's interests (i.e., we do not think of the prohibition as a hypothetical imperative), and not because some human institution has explicitly or implicitly decreed that it is wrong (i.e., we do not think of the prohibition as an institutional norm)—rather, we think of murdering innocent people as "wrong in itself" (1977: 34), that the action has "not-to-be-doneness somehow built into it" (1977: 40). Mackie maintains that any system of norms that didn't support these features would lack the peculiar but crucial kind of authority with which we imbue our moral prescriptions. It is not going too far, he asserts, "to say that this assumption [that there are objective values] has been incorporated in the basic, conventional, meanings of moral terms" (1977: 35).

Having established (to his own satisfaction) this conceptual step, Mackie is well-placed to execute some of the arguments that constitute the substantive step of the error theorist's case: that these properties are not actually instantiated. His *argument from relativity* starts with the observation of variation and disagreement among moral systems, and asks which explanatory hypothesis is more plausible: (i) that cultures disagree because they vary in their epistemic ability to discover objective moral facts, or (ii) that cultures disagree because they are, essentially, inventing the moral facts. Mackie thinks the latter hypothesis is preferable. His *argument from queerness* states that properties with such unusual practical authority would be simply too bizarre for us to countenance their existence—a skepticism that is supplemented by the question of by what means we would have epistemic access to such properties. ("A special sort of intuition" is quickly dismissed by Mackie as "a lame answer" (1977: 39).) Mackie is aware of but unmoved by the fact that billions of people have believed in morality. He explains the widespread error by speculating that humans have a tendency to "objectify" their subjective concerns and values (invoking Hume's metaphor of the human mind's "great propensity to spread itself on external objects"). If it is true that humans have a tendency to project their feelings onto their experience of the world, resulting in their erroneously judging certain actions to be required or prohibited with an objective authority, then billions of people believing in morality is exactly what the error theory would predict.

Opposition to Mackie's arguments (or error theoretic arguments with a similar structure) can take several forms. Probably the most vulnerable premise of the skeptical argument is the conceptual step, which maintains that moral properties must have a special kind of objective authority. One might respond that even if it is true that many people have thought moral properties to have this kind of authority, it is nevertheless not *necessary* that they do. Analogously, at one time everyone thought that the Earth is the center of the universe, but it

didn't follow that it was part of the very concept *Earth* that it has this feature, such that the discovery that the Earth is not at the center could be discounted on *a priori* grounds.

It is worth noticing that it is not just the precise content of Mackie's conceptual step that is potentially problematic here. The error theorist might alight on a different puzzling feature of moral properties: not that they involve objective authority (as Mackie thought), but, say, that they appear to imply the existence of a kind of pure free will. (And then the substantive step of the error theorist's argument would aim to establish that the world just doesn't contain the requisite kind of free will.) But the opponent can reply as before: granted that many people have thought that moral properties involve this odd kind of pure free will, it is nevertheless not *necessary* that they do. This problem generalizes and reiterates simply because there is no consensus on how we should distinguish between "X is widely (/universally) believed about Y" and "X is an essential part of the concept *Y*." Whatever problematic characteristic of moral properties the error theorist offers, it would appear that it is always possible to doubt that it is an *essential* characteristic, and at this point the error theorist and her opponent will not know where to turn to settle their dispute.

A typical example of this kind of opponent is the moral naturalist who identifies moral rightness (say) with something along the lines of *whatever maximizes happiness* (just to choose an extremely simple version). Naturalistic properties of this kind do not appear to have the kind of objective practical authority the existence of which Mackie doubts: there seems nothing inherently irrational about someone who doesn't care whether happiness is maximized, and so it would appear that a person might have no reason to care about moral rightness (ditto, *mutatis mutandis*, for any naturalistic property). The naturalist accepts this, but holds that such authority is an unnecessary extravagance that Mackie has tried to foist on moral ontology—in fact (the naturalist asserts) moral discourse was never committed to anything so grand. (See, e.g., Railton 1986; Brink 1989; Copp 2010.)

A different kind of opponent will concede that Mackie has gotten the conceptual step right— yes, it is an essential characteristic of moral properties that they have objective authority—but resist the substantive step. Such an opponent would argue that there *are* properties with this objective authority instantiated in the actual world. Certain forms of moral rationalism— according to which moral facts are identified with what is practically rational or irrational to do—are instances of this kind of opposition. If the rationalist holds a non-instrumentalist account of rationality (that is, denies that rational imperatives are hypothetical) and, moreover, maintains that rational norms exist and are not the product of human construction, then he or she can reject Mackie's substantive step. (See, e.g., Kant [1783] 1985; Nagel 1970; Gewirth 1978.) Certain forms of moral non-naturalism might also oppose Mackie by accepting the conceptual step but rejecting the substantive step. It is, after all, Mackie's commitment to a certain naturalistic worldview that powers his argument from queerness; if, therefore, one lacks that standing commitment to place a constraint on ontology—as the non-naturalist does—then one might countenance the existence of properties with objective authority. (See, e.g., Moore 1903; Enoch 2011.) (As mentioned before, even if the error theorist puts forward a different puzzling feature

of moral properties—e.g., that they imply the existence of pure free will—opposing arguments with a matching format will be possible.)

3. Noncognitivism

The noncognitivist denies the existence of moral knowledge by denying the belief criterion of the JTB analysis. Although we speak naturally of "moral beliefs" and say things like "Mary believes that telling lies is wrong," such talk is, according to the noncognitivist, at best a kind of shorthand that doesn't withstand scrutiny as a literal description. Rather, the kind of mental state associated with making a moral judgment is something other than belief, and, correspondingly, the kind of speech-act associated with publicly declaring a moral judgment is something other than assertion (the expression of belief). One can interpret noncognitivism as an attempt to avoid the error theory. Sympathizing, perhaps, with worries that moral properties would be ontologically very odd indeed, so much so that asserting their instantiation would involve speakers in falsehood, the noncognitivist instead proclaims that when we make a moral judgment we simply aren't making reference to moral properties: moral judgments cannot be false because they were never the kind of thing designed to have truth-value in the first place.

Classic noncognitivism appeared in the 1930s, espoused by such thinkers as Rudolf Carnap (1935), A. J. Ayer (1936), and Charles Stevenson (1937). More contemporary noncognitivism emerged in the 1980s, especially through the work of Simon Blackburn (1984, 1993) and Allan Gibbard (1992, 2003). One of the main differences between the two is in their differing attitudes toward commonsense folk intuitions. The early noncognitivists were more willing to accuse ordinary speakers of confused thinking about their own moral discourse; later noncognitivists offer a more conciliatory tone which attempts to accommodate as much of folk discourse as possible. So, for example, the statement with which I began this section—that noncognitivists maintain that when people say things like "Mary believes that telling lies is wrong" they are saying something false—is something which early thinkers like Ayer were willing to accept, but about which later noncognitivists will probably feel uncomfortable and attempt to deny. (And the same thing goes for statements like "Mary knows that it would be wrong for her to break that promise," "Mary asserted that breaking promises is wrong," and "The sentence 'Promise-breaking is wrong' is true." Ayer would likely accuse anyone who asserted any of these claims of simply making a blunder, whereas Blackburn would try to rescue appearances.)

At this point it would be natural to wonder whether later "noncognitivists" should even count as noncognitivists, given their desire to deny what appears to be the central tenet of the theory. And this is indeed a significant issue, with several commentators claiming that the noncognitivist can take attempts at accommodation too far (see Dreier 2004; Cuneo 2008). (It is worth noting that the contemporary metaethicist whose name is most closely associated with noncognitivism—Blackburn—has steadfastly eschewed the label.) This is not a debate to be adjudicated here; suffice it to say that if a so-called noncognitivist aims to accommodate moral beliefs—actual literal *beliefs*—then this position will no longer count as a denial of the belief

criterion of the JTB analysis, and therefore will not be of concern to this chapter. We are concerned here only with the belief-denying forms of noncognitivism.

Strictly speaking, noncognitivism is an entirely negative doctrine: it denies that moral judgments are beliefs, and thus denies the existence of moral knowledge. But there are a variety of positive theses offered by noncognitivists, too. Carnap states that moral judgments function as commands. Ayer claims that moral judgments function to express our feelings of (dis)approval while also attempting to arouse similar feelings in others. Stevenson has a mixed view, according to which moral judgments both state our feelings on the matter (a cognitive element) and enjoin others to share those feelings (a noncognitive element). Blackburn argues that moral judgments express our conative attitudes. And Gibbard maintains that moral judgments function to evince our allegiance to a given normative framework.

The attractions of noncognitivism are several. One has already been mentioned: that it sidesteps thorny puzzles about the ontology of moral properties and the nature of our epistemic access to them. According to the noncognitivist, there simply *aren't* any such properties with which we could be in epistemic contact. Another apparent advantage of noncognitivism is that it can accommodate the close connection between moral judgment and motivation. If the mental state a person is in when he makes a moral judgment is a conative state, then it is no mystery that he should have some (though defeasible) motivation to act in accordance with that judgment. Noncognitivism accounts easily for the phenomenon of moral variation: different individuals and different cultures approve and disapprove of different things, so moral variation is to be expected. Noncognitivism also accounts for the apparent heatedness and intractability of moral disagreement: being conative, moral judgments are both emotional and potentially impervious to rational debate and empirical evidence.

There are also several problems for noncognitivism. One has already been mentioned: that it would appear to deny the existence of certain phenomena that are commonly spoken of: moral beliefs, moral truths, moral knowledge, moral assertions. In this the noncognitivist doesn't appear to be worse off than any other kind of moral skeptic. All*, ex hypothesi*, deny moral knowledge—and while it may be true that the error theorist doesn't deny the existence of moral beliefs, whatever positive mark for "intuitiveness" the error theorist thereby gains, it is more than compensated for, most would think, by the accusation of massive falsehood.

The most famous problem for the noncognitivist is the *embedding problem* (a.k.a. the Frege-Geach problem). Suppose that when a speaker declares "Stealing is morally wrong" she is not ascribing a property to stealing—indeed, she is not asserting anything truth-evaluable at all—but rather saying something akin to "Boo to stealing!" or "Don't steal." That may be all very well, but what are we going to say about her embedding this judgment in logically complex contexts, such as "If stealing is wrong, then it is wrong to encourage your brother to steal"? And what are we going to say about how these two utterances—the freestanding and the embedded forms—combine in order to entail the conclusion "Therefore it is wrong to encourage your brother to steal"? The obvious first pass will look something like this:

1. Boo to stealing.
2. If boo to stealing, then boo to encouraging your brother to steal.
3. Therefore, boo to encouraging your brother to steal.

But if (1) is not even truth-evaluable, then what sense can we make of the conditional (2) being truth-evaluable (which is to say, how can we make sense of its really *being* a conditional?), and how can we understand the argument as valid, since *validity* is defined as a truth-preserving relation?

One strategy for the noncognitivist's reply is to earn the right to speak of the components of such arguments as truth-evaluable, after all. (This is what Blackburn refers to as "fast track quasi-realism" (1993: 184-6).) If this program also involves earning the right to speak of the components of such arguments as items to be *believed*, as it presumably will, then (as stated earlier) it will no longer count as the form of moral skepticism presently under discussion, and is thus beyond the purview of this chapter. Another strategy for the noncognitivist would be simply to bite the bullet and declare that we're doing something incoherent when we embed moral judgments in logically complex contexts. However, ascribing such widespread incoherence to a linguistic population will usually be thought of as a very unattractive option. (Note that it is qualitatively different to the widespread mistake that the error theorist attributes to speakers. The error theorist thinks we're all mistaken about the nature of the world—a possibility that we know, when surveying the history of human thinking, has frequently turned out to be the case. The noncognitivist, by contrast, thinks that we're all mistaken about the nature of our own linguistic practices—a possibility that is considerably less familiar.)

The more interesting strategy for the noncognitivist's reply is to deny that the components of the argument are truth-evaluable (deny that they are items of belief) while maintaining that in the moral context relations akin to, but distinct from, logical connectives and validity are in play. So, for example, while an entirely non-moral sentence like "If the book is blue, then the book is colored" uses a conditional connective (understood however we would ordinarily understand it), one that includes a moral claim, such as "If it is wrong to steal, then it is wrong to encourage your brother to steal," can be interpreted as using something that is *like* a conditional, but different in that it need not connect truth-evaluable relata but can connect, say, expressions of conative attitudes. The enterprise then becomes to define this quasi-conditional by reference to its relation to other quasi-connectives (such as quasi-negation, etc.) and its contribution to quasi-validity. What the noncognitivist seeks to develop, in other words, is a logic of attitudes. (See Blackburn 1984: 193-196; Schroeder 2010.)

There are further challenges for the noncognitivist that might be mentioned, but I shall instead end this section by discussing what bearing empirical evidence may have on the debate between the cognitivist and the noncognitivist. Suppose for the sake of argument that we could build a machine that reliably detects the occurrent conative state of *approval* (say): when the machine is aimed at someone in that state, a little green light turns on. Suppose we observe a great many conversations involving moral topics (and a bunch of control conversations involving

nothing moral), and we witness the green light go on whenever someone makes a moral judgment (and in conversations involving nothing moral, the light never goes on). It's important to recognize that this would not count as empirical evidence in favor of noncognitivism. The data would not support the conclusion that the state of approval *is* the moral judgment—that state might be something that reliably accompanies the moral judgment (e.g., that causes it or is caused by it). And the data would not support the conclusion that the utterance "X is morally good" *functions to express* that state of approval. Indeed, even if that utterance did function to express that mental state, we should hardly predict a constant correlation between the two. By comparison, we know that the act of apologizing functions to express regret (for that is part of how we define the speech-act of apology), but we don't therefore expect that any apology will be accompanied by occurrent regret (some apologies are insincere, for instance, and some are sincere but presumably too hurried or habitual for emotional arousal), nor that any regret will be accompanied by an apology.

If we want to know what kind of mental state a kind of utterance functions to express, the type of empirical evidence to be consulted is not neuroscientific but socio-linguistic. What determines whether a kind of speech-act S expresses a kind of mental state M (in the sense of "expresses" in which we're interested) are the conventions surrounding S, accepted by speaker and audience. The kind of "acceptance" here is more like knowledge-how than knowledge-that—the conventions governing our speech-acts are not always transparent to competent speakers—therefore the best guide to understanding the conventions regulating speech-acts is to scrutinize their use in an array of real-life settings. In other words, the noncognitivist who maintains that moral judgments express some kind of conative state rather than belief needs to locate socio-linguistic *evidence* that this is how moral judgments function.

Mark Kalderon (2005) presents such an argument in the course of advocating a view he calls "moral fictionalism." When one believes something, Kalderon claims, then upon encountering an epistemic peer who firmly disagrees, one has a "lax obligation" to examine one's reasons for believing as one does. Kalderon calls this "noncomplacency." However, the norms surrounding morality, he argues, permit complacency: we feel no embarrassment in steadfastly maintaining our moral views in the face of disagreement from epistemic peers. Were this argument to succeed, then there would be grounds for doubting that moral discourse is belief-expressing.

But evidence can be mustered in the other direction, too. David Enoch (2011) argues that cases of ordinary conflict of interpersonal preferences (e.g., you want us to play tennis and I want us to go to a movie) are governed by an impartiality norm: it is wrong to stand one's ground—rather, some kind of egalitarian compromise ought to be sought (e.g., tossing a coin). Cases of straightforward factual disagreement, by contrast, lack this characteristic. And the norms surrounding moral disagreement, Enoch goes on to argue, appear much more like those surrounding factual disagreement. If you and I disagree on the moral status of some kind of action—you think it's permissible and I think it's wicked—then impartial solutions like tossing a coin, or alternating doing it your way and then doing it my way, etc., seem inappropriate.

Enoch's conclusion is that moral judgment appears to be governed by norms indicating that it is not a matter of personal preference, which counts against many forms of noncognitivism.

I have sketched Kalderon's and Enoch's views in the briefest terms possible (for more detailed and critical commentary, see Joyce 2011 and Joyce 2013). And it will be noticed that the two views are not really opposed to each other, for it is entirely possible that moral judgments are neither expressions of belief nor mere reflections of personal preference. The intention here is not to support the evidence in either direction, but rather provide illustration of the kind of empirical investigation of socio-linguistic norms that might lead to progress in the debate between the cognitivist and the noncognitivist.

4. Justification skepticism

The justification skeptic denies the existence of moral knowledge by denying the justification criterion of the JTB analysis. As mentioned earlier, this form of skepticism is compatible with moral realism. According to this kind of skeptic, moral judgments are/express beliefs, and the beliefs may or may not be true, but we lack justification for holding these beliefs. Everyone accepts that people *sometimes* lack justification for their moral judgments, so this form of skepticism can be seen as something familiar writ large. One should distinguish a weak version of the view—according to which everyone as a matter of fact lacks justification for their moral beliefs (but where justification could in principle be instated)—from a strong version—according to which we are permanently stuck lacking justification. The fact that there are many competing theories of what it takes for a belief (or believer) to have justification makes it challenging to sum up or assess this view succinctly. I shall proceed by reflecting on some familiar folk opinions about what might render a moral judgment unjustified, and then sketch how some of these might be "writ large" to undermine moral knowledge in general.

Suppose someone, Sally, judges that an available action, *A*, is morally good. We will assume that cognitivism is true, so Sally's judgment is a belief. *What kind of factors might ordinarily lead us to doubt that Sally's belief is justified?*

First, we might think that Sally is insufficiently impartial in making this judgment—perhaps we see that a self-serving bias has distorted her thinking on the matter. We might think that other factors have distorted her thinking, too, such as her being overly emotional, or upset, or being drunk, or being hypnotized, and so forth. We might doubt the epistemic status of Sally's belief if we note that other people in the same situation (armed with the same evidence and powers of reflection as Sally) have confidently judged *A*-type actions to be morally abhorrent. If we notice that Sally's judgment about *A* appears to be inconsistent with some of her other beliefs, then we might doubt that her moral judgment is justified. Perhaps we know that Sally thinks that *A* is good simply because she's parroting what she's been told on the matter without sufficiently reflecting on it herself. Perhaps the process that has led Sally to this belief—including the process by which the belief gained purchase in the community from which Sally picked up the belief—is an unreliable one, in the sense that the process doesn't appear to be sensitive to

whether *A* actually is morally good. (Perhaps the process would have led to Sally's believing that *A* is good even if it is false that *A* is good, or perhaps we cannot see how the process even could track facts about moral goodness.) This is not an exhaustive list, and I'm not claiming that any of these factors alone would be sufficient for us immediately to declare Sally's moral judgment to lack justification, but it seems fair to say that everyday opinion allows that these factors all may bear on whether, and to what extent, Sally's belief that *A* is morally good is justified.

If someone thought that moral judgments *in general* manifest one or more of these justification-undermining features, then this may be ground for endorsing justification skepticism. Walter Sinnott-Armstrong (2006, chapter 9) argues that moral judgment exhibit several of these features.

First (Sinnott-Armstrong argues), which moral beliefs one holds is likely to affect one in a variety of material ways—sometimes very substantively—so we are all, to a greater or lesser extent, potentially biased in our moral judgments.

Second, while influences like drunkenness and hypnosis don't play a widespread role in human moral judgments, emotions certainly do. There is a body of evidence from experimental and social psychology suggesting that manipulating a subject's emotions will often result in altering his/her moral attitudes (see Moretti & di Pellegrino 2010; Horberg et al. 2011; Olatunji et al. 2016). There is evidence from neuroscience revealing the critical role that emotional arousal plays in moral judgment (see Young & Koenigs 2007; Decety & Wheatley 2015; Liao 2016). The idea that emotions influence moral judgment is of course entirely mundane, but empirical investigations may reveal the influence to be far more ubiquitous than commonly realized. The view from social intuitionism of moral judgment as an emotional knee-jerk reaction, with moral reasoning as a post-hoc rationalization (see Haidt 2001), is one that would very much cast the epistemic credentials of moral beliefs into doubt.

Third, moral disagreement among equally well-informed and reflective individuals/cultures appears to be a common phenomenon. In an earlier section I mentioned Mackie's argument that the ubiquity of moral disagreement might support the view that moral facts do not exist at all, but here we are countenancing something different: we are not being offered skepticism as an explanatory hypothesis of moral disagreement (as Mackie does), but rather we are wondering whether the existence of epistemic peers who disagree with you over X should undermine the justification of your beliefs about X. There is a very long tradition, stretching back to the ancient skeptics, of thinking that it should: references to widespread disagreement on all matters were central weapons in the classical skeptics' arsenal. But we are not here interested in those radical global forms of skepticism that deny *all* knowledge; we are interested in *moral* skepticism in particular—a skepticism that might be maintained within a broader non-skeptical worldview—raising the question of whether there is something special about *moral* disagreement in particular. It seems that there might be. After all, finding people who disagree with us about moral matters is much easier than finding people who genuinely disagree about whether this table is bigger than this chair. I interact on a daily basis with people who I know disagree with me about important moral matters. Whole cultures have thought that slavery is morally

acceptable, to say nothing of human sacrifice, violence toward women, killing foreigners, and so on. It might be responded that individuals and cultures have held some really wacky non-moral beliefs, too (e.g., that leeches are a cure-all for disease). However, there remains room for arguing that moral disagreements still seem different insofar as they seem particularly resistant to resolution through rational debate or consultation of evidence (see Doris & Plakias 2008; Adams 2013). This may be significant, since my justification for believing that *p* isn't undermined if I encounter someone who disagrees with me *if the disagreement could be resolved through reasoned discussion and examination of evidence*. It is only if the disagreement is intractable in the face of such discussion that the presence of an epistemic peer who believes that not-*p* might show my belief to lack justification.

Fourth and finally, significant doubts can be raised about the origin human moral beliefs. The crucial question is: Are the processes that lead to the formation of moral beliefs sensitive to the moral facts? The focus of this question is adjustable: it might pertain to how individuals come to internalize moral norms (a question of developmental psychology), or to how cultures come to adopt their moral norms (a question of cultural anthropology), or to how humans came to have the capacity to wield moral norms in the first place (a question of evolutionary psychology). Suppose that the best empirically-supported answers to all these questions turn out to be hypotheses that the moral error theorist can happily endorse. For example, perhaps evolutionary psychology will reveal that the reason that our ancestors developed a brain with the capacity to categorize their social world in terms like *right*, *good*, *evil*, (etc.) was not that it allowed them to track a special class of useful facts, but rather because it afforded them the advantage of strengthening social cohesion (perhaps by bolstering the motivation in favor of certain cooperative activities). This, clearly, is something to which an error theorist need have no objection. (This is the basis of the "evolutionary debunking argument"; see Joyce 2006, 2016.) But if all moral beliefs are best explained by a hypothesis which a moral error theorist may endorse, then the process that has given rise to moral beliefs can hardly be claimed to be a reliable one. And if moral beliefs are the product of an unreliable process—a process by which beliefs are not connected to the relevant facts—then moral beliefs lack justification.

It is possible that none of the aforementioned grounds for questioning the justification of moral beliefs is alone sufficient, but that together (or in combination with further considerations not mentioned here) they may entail that skeptical conclusion. It is worth noting, though, that the kind of justification skepticism that would be established would be the weak kind. The considerations raised at best show that our moral judgments lack justification; they do not appear to demonstrate that our moral judgments are unjustifiable. In other words, these arguments, if successful, would show that we lack moral knowledge, not that moral knowledge couldn't possibly be forthcoming.

An argument for the stronger skeptical conclusion is a version of the regress argument—well-known to ancient skeptics like Sextus Empiricus and Agrippa, but here aimed at moral knowledge in particular. According to the traditional regress argument, any belief must be justified by other beliefs, but in order for a belief to provide justification it must itself be

justified, which leads to an infinite regress. The ancients used this argument to generate a global skepticism, but we are interested in the possibility that the regress argument might succeed in showing that moral knowledge is impossible but fail to show that ordinary knowledge about the world is impossible. Here the argument can be sketched only very briefly.

The foundationalist attempts to defeat the regress argument by identifying a privileged class of beliefs whose justification does not depend on other beliefs. Candidates for these basic beliefs include beliefs about our own experiences and perceptual beliefs about the world. Let us suppose for the sake of argument that some such version of foundationalism is plausible, and so the regress argument for global skepticism is defeated. Now consider moral beliefs. Are they basic foundational beliefs or are they justified by other beliefs? While there are certainly attempts in the philosophical literature to argue that moral beliefs may be foundational (e.g., Ross 1930, 1939; Tolhurst 1990), these views have not won widespread acceptance (recall Mackie's snub of moral intuitionism as "a lame answer"); let us suppose then, again for the sake of argument, that these arguments fail. It is important at this point to recognize that the success of foundationalism generally (regarding, say, perceptual beliefs) is entirely compatible with the abject failure of *moral* foundationalism. If this were the case, then the only way left by which moral beliefs might be justified is by inference from non-moral beliefs. But this avenue also immediately runs into severe generic problems, for it is difficult to see how a body of entirely non-moral beliefs (or propositions, if you prefer) could even possibly non-trivially entail a moral belief (/proposition). Again: while there are certainly attempts in the philosophical literature to argue that moral conclusions may be validly derived from non-moral premises (e.g., Prior 1960; Searle 1964), these views have not won widespread acceptance.

Obviously, no attempt has been made here to make the several steps of this case at all plausible; my intention is just to outline the structure of an argument that leads to the strong skeptical conclusion that justification for our moral beliefs is absent and will remain so.

5. Conclusion

This chapter has revealed the many ways that a person may deny the existence of moral knowledge. The skeptical view has been organized into three major metaethical positions, and we have seen that there are often different versions of these positions and numerous argumentative paths leading to each. For the sake of simplicity, *knowledge* has been assumed to be amenable to the JTB analysis, but, clearly, if this presupposition were relaxed then the field of moral skeptical possibilities would be proportionally more complicated.

References

Adams, Z. 2013. "The fragility of moral disagreement." In D. Machuca (ed.), *Disagreement and Skepticism*. NY: Routledge. 109-130.

Ayer, A. J. 1936. *Language, Truth and Logic*. London: Victor Gollancz.

Blackburn, S. 1984. *Spreading the Word*. Oxford: Oxford University Press.

Blackburn, S. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.

Brink, D. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.

Carnap, R. 1935. *Philosophy and Logical Syntax*. London: Kegan Paul, Trench, Trubner & Co.

Copp, D. 2010. "Normativity, deliberation, and queerness." In R. Joyce & S. Kirchin (eds.), *A World Without Values*. Dordrecht: Springer. 141-165.

Cuneo, T. 2008. "Moral realism, quasi realism, and skepticism." In J. Greco (ed.), *The Oxford Handbook of Skepticism*. Oxford: Oxford University Press. 176-199.

Decety, J. & Wheatley, T. (eds.) 2015. *The Moral Brain: A Multidisciplinary Perspective*. Cambridge, MA.: MIT Press.

Doris, J. & Plakias, A. 2008. "How to argue about disagreement: Evaluative diversity and moral realism." In W. Sinnott-Armstrong (ed.), *Moral Psychology Vol. 2: The Cognitive Science of Morality: Intuition and Diversity*. Cambridge, MA: MIT Press. 303-331.

Dreier, J. 2004. "Meta-ethics and the problem of creeping minimalism." *Philosophical Perspectives* 18: 23-44.

Enoch, D. 2011. *Taking Morality Seriously*. Oxford: Oxford University Press.

Gewirth, A. 1978. *Reason and Morality*. Chicago: University of Chicago Press.

Gibbard, A. 1992. *Wise Choices, Apt Feelings*. Cambridge, MA.: Harvard University Press.

Gibbard, A. 2003. *Thinking How to Live*. Cambridge, MA.: Harvard University Press.

Haidt, J. 2001. "The emotional dog and its rational tail: A social intuitionist approach to moral judgment." *Psychological Review* 108: 814-834.

Horberg, E., Oveis, C., & Keltner, D. 2011. "Emotions as moral amplifiers." *Emotion Review* 3: 237-244.

Joyce, R. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.

Joyce, R. 2006. *The Evolution of Morality*. Cambridge, MA.: MIT Press.

Joyce, R. 2011. "Review essay: Mark Kalderon's *Moral Fictionalism*." *Philosophy and Phenomenological Research* 85: 161-173.

Joyce, R. 2013. "Review of David Enoch's *Taking Morality Seriously*." *Ethics* 123: 365-369.

Joyce, R. 2016. "Reply: Confessions of a modest debunker." In U. Leibowitz & N. Sinclair (eds.), *Explanation in Mathematics and Ethics*. Oxford: Oxford University Press. 124-145.

Joyce, R. 2017. "Fictionalism in metaethics." In D. Plunkett & T. McPherson (eds.), *Routledge Handbook of Metaethics*. NY: Routledge.

Kalderon. M. 2005. *Moral Fictionalism*. Oxford: Oxford University Press.

Kant, I. [1783] 1985. *Groundwork to the Metaphysics of Morals*. H. J. Paton (trans.). London: Hutchinson.

Liao, M. (ed.) 2016. *Moral Brains*. Oxford: Oxford University Press.

Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. NY: Penguin.

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

Moretti, L. & di Pellegrino, G. 2010. "Disgust selectively modulates reciprocal fairness in economic interactions." *Emotion* 10:169-180.

Nagel, T. 1970. *The Possibility of Altruism*. Princeton: Princeton University Press.

Olatunji, B., Puncochar B. D., & Cox, R. 2016. "Effects of experienced disgust on morally-relevant judgments." *PLoS ONE* 11(8): e0160357. doi:10.1371/journal.pone.0160357.

Prior, A. N. 1960. "The autonomy of ethics." *Australasian Journal of Philosophy* 38: 199-206.

Railton, P. 1986. "Moral realism." *Philosophical Review* 95: 163-207.

Ross, W. D. 1930. *The Right and the Good*. Oxford: Oxford University Press.

Ross, W. D. 1939. *The Foundations of Ethics*. Oxford: Oxford University Press.

Schroeder, M. 2010. *Noncognitivism in Ethics*. NY: Routledge.

Searle, J. R. 1964. "How to derive 'ought' from 'is.'" *Philosophical Review* 73: 43-58.

Sinnott-Armstrong, W. 2006. *Moral Skepticisms*. Oxford: Oxford University Press.

Stevenson, C. L. 1937. "The emotive meaning of ethical terms." *Mind* 46: 14-31.

Tolhurst, W. 1990. "On the epistemic value of moral experience." *Southern Journal of Philosophy* suppl. vol. 29: 67-87.

Young, L. & Koenigs, M. 2007. "Investigating emotion in moral cognition: A review of evidence from functional neuroimaging and neuropsychology." *British Medical Bulletin* 84: 69-79.

Suggested Reading

Audi, R. 1999. "Moral knowledge and ethical pluralism." In J. Greco & E. Sosa (eds.), *The Blackwell Guide to Epistemology*. Oxford: Blackwell. 271-302.

Jones, K. 2005. "Moral epistemology." In F. Jackson & M. Smith (eds.), *The Oxford Handbook of Contemporary Philosophy*. Oxford: Oxford University Press. 63-85.

Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. NY: Penguin.

Miller, A. 2013. *An Introduction to Contemporary Metaethics*. 2nd edition. Malden, MA.: Polity.

Shafer-Landau, R. 2003. *Moral Realism: A Defense*. Oxford: Oxford University Press.

Sinnott-Armstrong, W. 2006. *Moral Skepticisms*. Oxford: Oxford University Press.

Sinnott-Armstrong, W. & Timmons, M. (eds.) 1996. *Moral Knowledge? New Readings in Moral Epistemology*. Oxford: Oxford University Press.

Timmons, M. 1998. *Morality without Foundations*. NY: Oxford University Press.

Zimmerman, A. 2010. *Moral Epistemology*. NY: Routledge.