God and Moral Skepticism

Joe Milburn

God and Moral Skepticism

Joe Milburn

In this paper I want to address the question "Must morality be grounded on God?" in a roundabout way. I take it that morality is grounded on God if in order to make our moral discourse and practice intelligible we must refer to God, where the relevant sense of intelligibility is that of *making sense of*, or *understanding as rational*. Morality will be grounded on God if we need to refer to God to make sense of our moral concepts and practices.

Here I wish to look at one aspect of our moral practice that may or may not make sense if we do not make explicit reference to God. This is the practice of making knowledge claims in regards to morality; that is, claiming to have knowledge regarding right and wrong, of what one should and should not do. What I wish to argue for is that it is hard, if not impossible, to make sense of this aspect of our moral practice without supposing theism.

The structure of my paper will be as follows: After making some preliminary remarks, I will look at a recent skeptical argument found in Kieran Setiya's *Knowing Right from Wrong*. We can call this *the argument from epistemic luck*, since it tries to show that if our moral beliefs are true, then it is only a matter of luck that they are so; therefore, we cannot have moral knowledge.[1] I will argue that the theist is in a much better position to respond to the argument from epistemic luck than his atheist counterpart.

0. Preliminaries

To start, it is worth making some comments about the nature of morality and moral belief. I take morality to be an objective system of categorical, universal prescriptions and proscriptions, which can be expressed through "ought" statements. This is to say: morality makes demands of us; it demands us to do certain things, and abstain from doing other things; the demands of morality are categorical, in that they give us reasons for acting that do not essentially depend upon our beliefs, wants, desires, or what we happen to value at a given time. The demands of morality are universal, in that they apply at least to all human beings. One has a *moral belief* just in case he assents to a moral state-

---

[1] Kieran Setiya lays out this argument in full in chapter 3, "Knowledge," of his book, *Knowing Right from Wrong* (Oxford: Oxford University Press, 2012). See especially pp. 89–99 and pp. 111–27.

ment; for example, S believes that one should not lie, just in case he assents to the statement, one should not lie. Moral knowledge is simply knowledge of truths expressed by moral statements.

I take it that the characterization of morality I have provided is more or less the traditional characterization of morality. It is important to note that this characterization of morality makes no explicit mention of God. If, then, we need to refer to God to make sense of morality or moral knowledge so characterized, this would be of interest.

1. The Argument from Epistemic Luck

Setiya, in *Knowing Right from Wrong,* provides a powerful argument which, *mutatis mutandis*, can be used to argue for moral skepticism.[2] We can call this the argument from *epistemic luck* or *non-accidental truth*. Setiya's argument starts from the following truism: knowledge is non-accidentally true belief. Now, Setiya argues that for a belief to be non-accidentally true it must meet what we can call condition K.

> Condition K: for all subjects S, and propositions p, S true belief that p is *non-accidentally true* (i.e., not subject to epistemic luck) only if, (1) S's will has formed and currently sustains the belief that p through some reliable method m, and (2) the reliability of this method m, explains or is explained by the fact that S uses it.[3]

Here we can understand methods of belief formation to be a disposition to form beliefs given certain psychological and environmental conditions. A method m, then, is reliable just in case it regularly produces true beliefs in the actual and all near-by possible worlds. What is it for the reliability of a method m to explain or to be explained by the fact that S uses it? For this to be the case either S must use method m *because* m is reliable, or method m must be reliable *because* S uses it. As Setiya puts it:

> Three kinds of "because" might be relevant here: efficient, final, and formal. Perhaps S uses m because…it reliability follows from the etiology of its use; perhaps these factors have a common cause. Perhaps S uses m for the sake of forming true beliefs: it is the function of m to be reliable and the connection is teleological. Perhaps the conjunction is constitutively explained. A constitutive account of x is an account of its nature, or what it is to be that

---

[2] Setiya, *Knowing Right from Wrong*, 111–27. Setiya's target is ethical knowledge, which would include our knowledge about reasons for acting in general. I am focusing the argument on moral knowledge.

[3] Kieran Setiya, *Knowing Right from Wrong,* 96.

thing…. A constitutive explanation of p and q is one in which p
follows from q together with truths of this kind. Thus the reli-
ability of m might follow from what it is for that method to be
used by S.[4]

Once we have condition K in hand, we can then present a skepti-
cal argument against morality in the following way. Suppose that there is a
non-moral, naturalistic causal explanation for why we have the moral beliefs
that we do. In this case, it seems the reliability of our methods of moral be-
lief formation neither explains nor is explained by the fact that we use these
methods: on the one hand, there appears to be no explanation in terms of
efficient causation for this fact. Given the availability of a non-moral nat-
uralistic explanation of why we have the moral beliefs we do, we have no
reason for believing that the truths of morality account for our using some
method of moral belief formation, and this method being reliable. Neither is
it plausible (without assuming theism) that the function of our moral belief
forming methods is to reliably produce true moral beliefs. Given that there is
a non-moral naturalistic explanation of why we have the moral beliefs we do,
we can assume that our methods of moral belief formation would have been
used whether they were reliable methods or not. But in this case, we cannot
plausibly claim that we use the methods of moral belief formation that we
do because they have the function of being reliable.[5]

Nor does it seem that we can give a constitutive explanation of why
we use the methods we do in terms of the reliability of these methods. On
the one hand, we might give what Setiya calls a constructivist account of
moral truth. In its simplest form constructivism is the doctrine that, for some
"ought" statement to express a moral truth is for us to be disposed to believe
this "ought" statement. Given constructivism in regards to moral truth, we
can see how one can posit an explanatory relationship between our using a
method m, and the reliability of the method. The reliability of the method is
explained by the fact that we use this method, since according to constructiv-
ism, what makes moral statements true is that we are disposed to believe that
they are. But this kind of explanation will not do. The sort of simple moral
constructivism it appeals to is incompatible with the idea that moral demands
are categorical and universal. S might be disposed to believe that murdering
the innocent is permissible if good enough consequences follow. T might be
disposed to disagree. According to constructivism in its simplest form, it will
be true for S that murdering the innocent is sometimes permissible, while it

---

[4] Kieran Setiya, *Knowing Right from Wrong*, 98.
[5] Kieran Setiya, *Knowing Right from Wrong*, 111–12

will be false for T. But in this case morality will be neither universal nor categorical.[6]

On the other hand, one might try to give what Setiya calls an externalist explanation. According to the most simple form of externalism, part of what it is for an agent to have moral concepts and moral beliefs is for the agent to use a sufficiently reliable method m in forming moral beliefs. In other words, it is a condition of having moral beliefs that one's methods for forming moral beliefs are reliable. In this case, there is an explanatory connection between the fact that one uses a method m for forming moral beliefs and its reliability: the method one uses for moral belief formation, is a method for moral belief formation only because it is reliable.

However, appeals to simple externalism are also insufficient. Again, externalism is incompatible with the possibility of widespread, radical disagreement over morality. If S and T seem to disagree over most moral issues, and S seems to be wrong, then according to externalism, S does not have any moral beliefs, since he is not forming his beliefs using a reliable method. But if S has no moral beliefs, then he and T cannot be disagreeing over morality. So they must only seem to disagree. But this is absurd.[7]

In this case it seems that the Explanation Condition cannot be met by our moral beliefs. There seems to be no explanatory connection in terms of either efficient, final, or formal causality, between the reliability of our methods of moral belief formation and the fact that we use these methods. Therefore, moral knowledge seems beyond our grasp.

2. Theistic and Atheistic Responses to the Argument from Epistemic Luck

We can restate the argument from Epistemic Luck as follows:

1. Condition K is a necessary condition for a belief's being non-accidentally true.

2. But our moral beliefs cannot meet condition K.

3. Therefore, we cannot have moral knowledge.

In this section, I want to argue that the theist is in a much better position to deal with the argument from Epistemic luck than is the atheist.

---

[6] Kieran Setiya, *Knowing Right from Wrong*, 115–17. Setiya's criticism of constructivism is somewhat different then how I have put things here. Setiya criticizes constructivism for implying that there cannot be widespread ethical disagreement; the criticism that I have voiced is special to the problem of moral knowledge.

[7] Kieran Setiya, *Knowing Right from Wrong*, 117–22.

Supposing theism, we have good reason for rejecting (2) above. According to theism there is an all good, all powerful, all knowing God who is the creator of man. Given God's goodness we can suppose that he would want man to have knowledge of moral truths; given his wisdom and power, we can suppose that God would create mankind so that it would use methods of moral belief formation that would produce mostly true moral beliefs. So given theism, there is an explanatory connection between the fact that we use a method m for forming moral beliefs and the reliability of m. God sees to it that we use method m in forming our moral beliefs, and he did this because m is reliable.[8]

If we bracket theism, do we have good reason for rejecting either premise (1) or (2) above? Let's start by looking at premise (1). One might try to argue against condition K by pointing to a paradigmatic instance of knowledge that fails to meet condition K. So for instance, one might argue that our perceptual knowledge or our inductive knowledge fails to meet condition K; but since it is obvious that we have perceptual or inductive knowledge, then condition K must not be a real condition on non-accidentally true belief.

In this context, the matter becomes somewhat complicated. Setiya himself attempts to show that our perceptual, inductive, and mathematical knowledge can meet condition K without presupposing theism;[9] but suppose Setiya is wrong about this. Setiya's failure would not imply that there are counterexamples to condition K. It might be that not only must our moral knowledge be grounded in God, but that our knowledge *in general* must be grounded in God. For given theism, it seems condition K will be met in regards to all of our major branches of knowledge. Since God is all good; he is not a deceiver. And since he is all wise and all knowing, he will see to it that the human being will, in general, form beliefs using reliable methods. So given theism, we have reason to believe that in general our beliefs meet condition K or at least can meet condition K.

It is better, then, to judge condition K on its own merits. Setiya argues for condition K by considering existing accounts of epistemic luck or accidental true belief, and showing how such accounts fail in important ways since they lack condition K. Setiya's primary target are *modal accounts* of epistemic luck. The most important of these attempt to understand epistemic luck in terms of a belief's being *safe*.[10] S's true belief that p is safe if and only if at the majority of nearby possible worlds in which S forms his belief p in

---

[8] Setiya makes a similar point. See *Knowing Right from Wrong*, 114–15.

[9] Kieran Setiya, *Knowing Right from Wrong,* 99–110.

[10] See Duncan Pritchard, *Epistemic Luck* (Oxford: Clarendon Press, 2005), for a vigorous defense of such an account.

the same way, p is true. According to the view in question, a true belief will be accidentally true just in case it fails to be safe.

Setiya points out serious problems with safety accounts of epistemic luck. Most importantly, safety accounts of luck cannot deal with necessary truths. Suppose for instance, that one were to believe Goldbach's conjecture using the following method:

> Coin Flip: Flip a coin; if it lands heads believe the Goldbach conjecture. If it lands tails disbelieve it.

Suppose that one flips a coin and it lands heads. On this basis he believes that Goldbach's conjecture is true. Supposing that Goldbach's conjecture is true, one's belief in this case would be accidentally true. Nevertheless, this belief would be safe, since at the majority of nearby possible worlds in which one forms the belief that Goldbach's conjecture is true on the basis of a coin landing heads, this belief is true. After all, Goldbach's conjecture, if true, is true at every possible world.[11]

Nor does it help to make safety relative to some method, ranging over a set of beliefs. For instance, one might try to amend safety in the following way:

> Safety of methods: S's belief that p is safe, just in case S uses some method m to form the belief that p, and in the majority of nearby possible worlds in which S uses method m, S forms a belief in some proposition q which is true.[12]

If we understand safety in this way, then a safety account of epistemic luck would correctly deal with the Goldbach's conjecture case above. This is because in many nearby possible worlds, one uses the same method (e.g., tossing a coin) to form false beliefs. So there will be many possible worlds in which one uses the same method but believes that Goldbach's conjecture is false.

However, understanding safety in this way fails to rule out the following case:

> Cautious coin flip: Sam forms his belief regarding Goldbach's conjecture on using the following method: Flip a coin, if it lands heads, believe that Goldbach's conjecture is true. If not, withhold belief regarding Goldbach's conjecture. Sam flips a coin and it lands heads. On this basis he believes Goldbach's conjecture is true.[13]

---

[11] Masahiro Yamada, "Getting it Right by Accident," *Philosophy and Phenomenological Research,* 83 no. 1 (2011): 102. Yamada gives a similar example.

[12] Kieran Setiya, *Knowing Right from Wrong,* 90.

[13] Kieran Setiya, *Knowing Right from Wrong,* 91.

In this case, Sam's belief that Goldbach's conjecture is true is formed using a safe method. Still, it is obviously an accidentally true belief. If, however, we supplement safety accounts of epistemic luck with condition K, such cases are ruled out. It cannot plausibly be said that Sam is using the cautious coin-flip case because it is reliable.

Now, in defending principle K, Setiya does not consider virtue theoretic accounts of epistemic luck or accidental truth. On these accounts S's belief that p will be accidentally true, just in case, S's belief is true, but not because it was produced by S's intellectual virtue or ability. Greco puts things as follows:

> S's belief that p is [non-accidentally true] just in case S's belief that p is true because it is produced by intellectual virtue or ability.[14]

This formulation, however, is somewhat infelicitous. As Michael Levin has pointed out,[15] this way of putting things makes it seem as though we were asking for a semantic explanation why the belief is true. But this semantic explanation will be something like this: why S's belief that p is true is p. After all, S's belief that p will be true if and only if p. A better way of putting things is as follows:

> Virtue Theory: S's belief that p is non-accidentally true, just in case S's believing truly that p, as opposed to falsely that not p, or withholding belief concerning p, is to be explained by the exercise of S's intellectual virtue .

Even this, formulation, however, is not exactly correct; we can consider cases in which S's believing truly that p is to be explained by the exercise of S's intellectual virtue, but in which S's belief is accidentally true in the relevant sense. Consider for example, the following case due to John Turri:

> (HOBBLED): A competent, though not [masterful] inspection of the crime scene would yield the conclusion that a man with a limp murdered Miss Woodbury. Holmes saw through it and had already deduced that Dr. Hubble poisoned the victim under pretense of treating her.
>
> Holmes also recognized that the scene would fool Watson, whose own inspection of the scene was proceeding admirably competently, though not masterfully... "Look at him," Holmes thought, "measuring the distance between footprints, noting their com-

---

[14] John Greco, "Knowledge and Success From Ability," *Philosophical Studies* 142, no. 1 (2009): 18.

[15] Michael Levin, "Virtue Epistemology: No New Cures," *Philosophy and Phenomenological Research* 69, no. 2 (2004): 405.

parative depth, and a half dozen other things, just as he ought to. There's no doubt where this will lead him—think how discouraged he will be." Holmes then resolved, "Because he's proceeding so competently, I'll see to it he gets it right!"

Holmes sprang into action. Leaving Watson, he hastily disguised himself as a porter, strode across the street to where Hubble was, and kicked him so hard that Hubble was thereafter permanently hobbled with a limp.[16]

In this case, Watson's true belief that the perpetrator is hobbled is to be explained by the exercise of some intellectual virtue. Still Holmes's belief is accidentally true in the relevant sense. Holmes cannot be said to know that the perpetrator is hobbled.

Turri's response to this sort of case, is to suggest that the correct virtue theoretic account of non-accidental truth is the following:

Manifest: S's belief that p is non-accidentally true just in S's believing truly that p is the manifestation of an intellectual virtue.[17]

In this case, if S believes truly that p, but in so doing he does not manifest an intellectual virtue, his belief is accidentally true.

I think Turri's virtue theoretic account of epistemic luck is on the right track. But it raises a number of interesting questions. Something that Turri does not go into is what epistemic virtues must be like on his account of non-accidental truth. According to Turri, S's believing truly that p must be the manifestation of an intellectual virtue in order for S's true belief that p to be non-accidentally true. For simplicity's sake we might suppose that epistemic virtues are canonical dispositions: that is, there will be descriptions of epistemic virtues of the following form:

Canonical: S has epistemic virtue E just in case S is disposed to X in conditions C.

In this case, we must understand X-ing to be the manifestation of epistemic virtue E. But then believing truly that p must be an instance of X-ing; whatever this might be, in order for there to be non-accidental belief. But it is unclear what X-ing must be in order for *truly believing* to be a token of it. One tempting thing to say is that we should substitute *believe the truth regarding some range of truths* for X in the formula above. But something more needs to be said. This is because we can suppose that, for example, the person who believes that the Goldbach conjecture is true when a coin lands tails, is mani-

[16] John Turri, "Manifest Failure: The Gettier Problem Solved," *Philosopher's Imprint* 11 no. 8 (2011): 5.

[17] John Turri, "Manifest Failure," 7.

festing a certain disposition to believe a certain range of truths in conditions in which he flips a coin and sees that it lands heads. What we need to do is separate mere dispositions to believe the truth from virtuous dispositions to do so.

I take it that Setiya's condition K is a minimal condition for a disposition's being an epistemically virtuous disposition. That is, in order to be an epistemically virtuous disposition of S's, S's disposition to believe a range of truths must meet condition K. S's having this disposition must be explained by the fact that it is a disposition to believe the truth; or that it is a disposition to believe a range of truths must be explained by the fact that S uses it. Otherwise, this disposition cannot be thought of as an epistemically virtuous disposition.

If what I have said is correct, then, it seems that condition K is a viable condition on non-accidental belief; neither virtue theoretic nor modal accounts of non-accidentally true belief are complete without it. If we are to reject the argument from epistemic luck without recourse to theism, we must then have reason for believing that (2) is false: that is, reason to believe that our moral beliefs can meet condition K.

This is the tack that Setiya himself takes, but I find little promise in it. On the one hand, as Setiya notes, it seems that there are no plausible etiological explanations relating the reliability of our methods to the fact that we use them. Suppose, for instance, that non-naturalism regarding morality is true. In this case it is hard to see how moral facts or properties could be causally efficacious; but absent this efficacy, how could there be a causal story relating the reliability of our methods of moral belief formation to the fact that we use them? Suppose instead, that naturalism is true. In this case, moral properties and facts are simply complex natural properties and facts. So we can understand how moral properties and facts might be causally efficacious. Nevertheless, we have no reason for believing that the natural property that turns out to be identical to "moral rightness" has played a causal role in our adopting the methods of moral belief formation.

Given that there is no plausible etiological connection between moral facts and properties and our using a reliable method of reliable belief formation, it is hard to countenance the idea that we use some reliable method of moral belief formation m because its function is to be reliable. In this case we want to say that the natural function of our belief forming methods is to be reliable, and this is why we use it. But given that there is a causal explanation of why we use the methods of moral belief formation that we do, and this does not appeal to moral facts and properties, it seems false to say that we are using our method of belief formation because it has the function of being reliable.

This leaves only constitutive explanations connecting the reliability of our methods of moral belief formation to the fact that we use them. But as has been pointed out, simple constitutive explanations like Constructivism or Externalism will not do. Constructivism fails to respect the universality of morality. Externalism fails to respect the possibility of deep and widespread moral disagreement. Could a more sophisticated version of these views avoid these problems?

Setiya's own strategy is to adapt simple externalism so that it allows for disagreement at least amongst individuals and communities. Setiya claims that part of what it is to have the concepts of morality is to belong to a *life form* whose method for identifying moral truths is sufficiently reliable. We can call this thesis, 'Natural Externalism.' Given Natural Externalism, we can give the following explanation for why we use the methods we do, that appeals to their reliability: We use method m for forming moral beliefs, because we are human beings, and human beings use method m for forming moral beliefs. But if m were not reliable, then we would not use it to form moral beliefs, since if it were not reliable, human beings would no longer have the concept of morality. At the same time, this thesis allows for disagreement between members of the human life form regarding moral beliefs. This is because generic statements like "Human beings use method m for forming moral beliefs" can admit of counter-instances but still be true. So unlike simple externalism, Natural Externalism does not imply that there must be widespread moral agreement amongst all users of moral concepts.[18]

While I take it that adopting Natural Externalism allows us to defeat the argument from epistemic luck without supposing theism, I believe that there are serious disadvantages to this strategy. First, the thesis of Natural Externalism is *ad hoc*. Setiya himself does not motivate it, except as a means of avoiding a form of moral skepticism. Second, there are reasons to believe that Natural Externalism is false. If widespread and radical disagreement can occur between individual human beings, or whole communities of human beings, it seems it also could occur between different life forms. But Natural Externalism denies that this is a possibility. Furthermore, consider that there could be an individual that belonged to a life-form that did not have a reliable method for forming moral beliefs; however, while the individual's life form did not use a reliable method in regards to morality, we might have overwhelming evidence that the individual used a reliable method. When asked whether or not, for example, one should ever intentionally kill an innocent human being, this individual would answer correctly and he would always act in accordance with his professed moral beliefs. But according to Natural Externalism, this individual could have no moral concepts and no moral beliefs,

---

[18] Kieran Setiya, *Knowing Right from Wrong*, 128–38.

and so would not have a reliable method for forming moral beliefs. This is counter-intuitive to say the least.

Another possibility is that one adopts what we might call 'Natural Constructivism.' On this view, to be a moral truth just is to be a moral statement that would be believed in ideal conditions, if one were to use a method of moral belief formation *m*, where *m* is whatever method of moral belief formation is used by the human life form. This view seems to avoid, at least on the surface, the problem of relativity. Assuming that there is some method m that the human life form uses, and that in idealized conditions it would yield the same verdicts regarding different moral questions, Natural Constructivism seems to allow an explanation connecting the fact that we use a method of belief formation and its reliability. I use method m because I am a human being and the human being uses method m. But to be a moral truth just is to be a moral statement that method m will judge to be true. So necessarily my method of belief formation is reliable. At the same time, Natural Constructivism seems to respect the universal nature of morality. It is not the case that what is morally true for me or for my community might be false for you or for your community.

Natural Constructivism, however, faces problems similar to those that afflicted Natural Externalism. On the one hand, natural constructivism seems *ad hoc*. Apart from the fact that Natural Constructivism helps us to avoid moral skepticism, why should we believe it? Furthermore, there are reasons for believing that Natural Constructivism is false. On the one hand, Natural Constructivism makes it difficult to understand why morality is categorical. If moral truth just is what Natural Constructivism says it is, why do all human beings have reason to do what morality demands of them? On the other hand, Natural Constructivism seems to deal with problems of relativity at a higher level. After all, it seems to be a contingent fact that the human being uses some particular method of moral belief formation m. So the human being could change so that it is no longer true that he uses method m, but uses method m* instead. But if method m and m* disagree with each other, it seems we have once again lost the universal aspect of morality. Perhaps it could be claimed that it is constitutive of the human life form that the human being uses *m*. But once again, it is hard to see what reason we have for making this claim.

My claim is that what is true for Natural Externalism and Natural Constructivism is true in general. There seems little reason to believe that there is some constitutive relation between the reliability of our methods of moral belief formation and the fact that we use them, apart from the fact that positing this relation would rule out moral skepticism. Furthermore, I take it that all sorts of externalism will implausibly deny the possibility of moral disagreement at some level; whether between individuals, communities, or

life forms. Likewise, all forms of constructivism will either put into jeopardy the universal nature of morality or they will need to claim implausibly that using a certain method of moral belief formation is constitutive of the kinds of beings that we are.

Given this, while I do not claim that it is impossible to resist the argument from epistemic luck without supposing theism, I would contend that the theist is much better placed to deal with the argument of epistemic luck. Once we allow ourselves the central tenets of theism, we can disregard this skeptical argument for free. Must morality be grounded on God then? Perhaps *must* is too strong a word; but I think it is certain to say that morality, our moral practices and beliefs, are best grounded on God.[19]

—*The University of Pittsburgh*