

# From Isolation to Skepticism

Scott Hill

(This is a draft. Please cite the final version which is forthcoming in  
*Erkenntnis*.)

## 1 Introduction

Many people believe that moral properties lack causal powers. Whatever goodness is like, whatever has or does not have goodness, it makes no causal difference to the way in which our moral beliefs are formed. This gives rise to a question: How could it be anything other than an accident that our moral beliefs are true? It is easy enough to see why it isn't an accident that normally formed perceptual beliefs are true. My beliefs about the black mug on my desk, for example, are causally related to the mug in ways that ensure that my mug beliefs are not merely accidentally true. If the mug were a different color, I would not have believed it was black. So far so good. But trouble emerges when we apply this test to our moral beliefs. If the moral realm had been different, our moral beliefs would have been the same. This leads to skepticism.

This familiar sort of worry has been subjected to a number of objections<sup>1</sup>. I show that there is a way of stating the worry that is immune to each of these objections. If moral properties lack causal powers, then skepticism ensues.

## 2 Background

If I decide whether to believe the number of stars in Andromeda is even on the basis of a coin toss, and the outcome of the toss happens to match the truth, then my belief is accidentally true in a way that rules out knowledge. This case is easy. But there are other, more difficult, cases in which it is unclear whether a belief is accidentally true in a problematic way. A test is needed. Some tests require that a subject's belief have the right modal tie to the truth. Let S believe p. Then one such test employs:

**Sensitivity:** If p were false, S would not believe p.

another test employs:

---

<sup>1</sup>See Bealer (2000), Behrends (forthcoming), Clarke-Doane (forthcoming), Enoch (2011), Hanna (2010), Huemer (2005), Lewis (1986), Littlejohn (2011), Skarsaune (2010), Sturgeon (1985), Vavova (forthcoming), and Wielenberg (2011).

**Safety:** If S were to believe p, p would be true.

Sensitivity is understood to say that S does not believe p at the closest worlds at which p is false. Safety is understood to say that p is true at all nearby worlds at which S believes p. Let distance between worlds, for now, be evaluated using the Lewisian<sup>2</sup> closeness relation:

1. It is of the first importance to avoid big miracles.
2. It is of the second importance to maximize the spatio-temporal region of perfect match.
3. It is of the third importance to avoid small miracles.
4. It is of the fourth importance to maximize the spatio-temporal region of approximate match.

An event at one world is a miracle relative to a second world if and only if that event violates the laws of the second world. The greater the violation of the second world's laws, the bigger the first world's miracle is relative to the second world. Call a belief that fails to have the right modal connection to the truth 'modally problematic'.

A different test employs:

**Reliability:** S's belief that p was formed by a reliable process.

A reliable belief forming process is understood to be one that yields a sufficiently high ratio of true to false beliefs. It is unclear exactly how high the ratio must be. But one to one is not high enough. And all to none is.

Finally, consider a metaethical doctrine:

**Isolation:** Moral properties lack causal powers. What moral properties are like, and whether something has a moral property, makes no causal difference to our moral beliefs.

### 3 From Isolation to Skepticism

The ideas in the previous section may be combined to form a skeptical argument:

#### The Argument from Isolation

- (1) Isolation is true.
- (2) If Isolation is true, then our moral beliefs are modally problematic.
- (3) If Isolation is true, then our moral beliefs are unreliably formed.
- (4) Therefore, our moral beliefs are modally problematic and unreliably formed.

---

<sup>2</sup>See Lewis (1979).

- (5) If our moral beliefs are modally problematic and unreliably formed,  
then we lack moral knowledge.
- (6) Therefore, we lack moral knowledge.

(1) is left undefended. I aim to show that Isolation leads to skepticism; not that skepticism is true.

My initial (but not final) defense of (2) relies on the Lewisian closeness relation. Consider two worlds<sup>3</sup>:

**W<sub>0</sub>**: Torturing babies for fun is not wrong and everything else (including our moral beliefs) is the same as it is in the actual world.

and

**W<sub>1</sub>**: Torturing babies for fun is not wrong and we do not believe it is wrong.

Concerning sensitivity: Since everything but whether it is wrong to torture babies is the same,  $W_0$  maximizes the spatio-temporal region of perfect match. And, given Isolation, our belief was formed in a way that is causally independent of wrongness. So no Lewisian miracles are required to keep it the same. Contrast this with  $W_1$ . Since our belief is different at  $W_1$ , the region of perfect match is not maximized. At  $W_0$  there is perfect match and no miracles. At  $W_1$  there is not perfect match. So  $W_0$  is closer to the actual world than  $W_1$ . Therefore our belief is insensitive.

Concerning safety: Since  $W_0$  is a world with perfect match and no miracles, it is as close to the actual world as any world can get. So it is a nearby world. And there is therefore a nearby world at which we believe torturing babies for fun is wrong but that belief is false. So our belief is unsafe.

My considered defense of (2) will rely on dealing with various worries that the Lewisian closeness relation is inapplicable to moral counterfactuals and various attempts to show that our moral beliefs are sensitive or safe stemming from the trivial truth of counterpossibles or the supervenience of the moral on the non-moral.

My initial defense of (3) relies on a comparison due to Field<sup>4</sup>. Suppose we had beliefs about a distant planet. These include elaborate beliefs about the inhabitants of the planet and the goings on there. Since the planet is so distant, it is causally isolated from us in all relevant respects. So however we got our beliefs about the planet, those beliefs were formed in a way that is causally independent of that planet's happenings. Our elaborate beliefs about that planet are unreliably formed. And, given Isolation, there is no relevant difference between our imagined planet beliefs and our actual moral beliefs. So our moral beliefs are unreliably formed.

---

<sup>3</sup>More exactly:  $W_0$  and  $W_1$  are classes of worlds. My claim is that there is a world in  $W_0$  that is closer to the actual world than any world in  $W_1$ . Or put another way: All of the closest worlds at which torturing babies for fun is not wrong are in  $W_0$ . None of the closest worlds are in  $W_1$ .

<sup>4</sup>This is a variant of Field's (1989, pp. 25-30) Nepalese Village case. Similar examples appear in Street (2006), Enoch (2011, p. 158), and Bedke (2009).

My considered defense of (3) will rely on an examination of various attempts to explain how our moral beliefs could be reliably formed and how they are different from our imagined planet beliefs.

My defense of (5) is twofold. First, proponents of Isolation typically adopt either sensitivity or safety or reliability as their test for problematic accidentalness. As we will see below, different authors adopt different tests. But they all seek to show that our moral beliefs satisfy at least one or another of the tests. And, because our moral beliefs succeed in passing whichever of the preferred tests the authors in question adopt, it is claimed that such beliefs are not true in a problematically accidental way. Below, I argue that our moral beliefs fail to satisfy any (suitably refined version) of the tests in question. If I am right, then our beliefs fail to pass the best versions of the various tests that the proponents of Isolation have themselves adopted. Therefore, such proponents should, by their own lights, judge our moral beliefs to be problematically accidentally true and not instances of knowledge. Second, it is not just proponents of Isolation that adopt these tests. They are employed by diverse and established bodies of literature within epistemology<sup>5</sup>.

## 4 Trivially Sensitive, Safe, and Reliable

One might try objecting that (2) and (3) are false. One could do this by appealing to<sup>6</sup>:

**(Trivial):** Beliefs with necessarily true content are trivially sensitive, safe, and reliably formed.

Consider (2). I said that our moral beliefs lack sensitivity and safety and therefore we lack moral knowledge. But someone might doubt this. Someone might suggest that the moral truths we are interested in are necessary truths. And given the standard analysis of counterfactuals<sup>7</sup>, beliefs with necessarily true content are trivially sensitive<sup>8</sup> and safe<sup>9</sup>. Our moral beliefs trivially satisfy the relevant conditions. So (2) is false.

Next, consider (3). I said that our moral beliefs must be unreliably formed and therefore we lack moral knowledge. But one might follow Lewis in holding that beliefs with necessarily true content are not only trivially reliably formed but also trivially infallibly formed. “For if it is a necessary truth that so-and-so, then believing that so-and-so *is* an infallible method of being right. If what I believe is a necessary truth, then there is no possibility of being wrong. That is so whatever the subject matter of the necessary truth, and no matter how it

---

<sup>5</sup>See the footnotes below for relevant citations.

<sup>6</sup>This objection is inspired by Lewis’ (1986, pp. 113-5) answer to parallel concerns about modal realism. Enoch (2011, pp. 172-3), Wielenberg (2010, pp. 460-1), and others make similar points.

<sup>7</sup>See Lewis (1973).

<sup>8</sup>See Sosa (1999).

<sup>9</sup>See Roland and Cogburn (2011).

came to be believed<sup>10</sup>.” Therefore, since our moral beliefs are necessarily true, they are reliably formed. And (3) is false.

This objection is unsound. Start with the question of whether beliefs with necessarily true content can be formed in an unreliable way. Imagine that you are contemplating the existence of God. You decide to let a coin toss determine what you will believe about the matter. If the coin lands heads you will believe that God exists. Otherwise, you will believe God doesn’t exist. Suppose the coin lands heads and you form your belief accordingly. Suppose, in addition, that God exists. Now, note that God, if He exists, exists necessarily. Forming your beliefs on the basis of a coin toss is an unreliable method. So it is possible to form beliefs with necessarily true content using an unreliable method and therefore such beliefs are not trivially reliably formed<sup>11</sup>.

Now consider whether beliefs with necessarily true content are trivially sensitive and safe. There is a genuine problem here. But the problem is not due to sensitivity or safety. The problem is instead due to the Lewisian analysis. According to the analysis, all counterfactuals with necessarily false antecedents (counterpossibles) are trivially true. This is why sensitivity and safety are trivially satisfied by beliefs with necessarily true content. However, the standard analysis goes wrong here<sup>12</sup>. Many counterpossibles are not trivially true. To see this, suppose that God (necessarily) does not exist. Now consider:

**CP1:** If God were to command us to torture babies just for fun, we would be morally obligated to do so.

**CP2:** If God were to try to open a nightclub so exclusive that even He couldn’t get in, He would be unable to do so.

**CP3:** If God were to actualize any world with more than 43 billion significantly free creatures, it would contain at least the amount of evil that is in the actual world.

---

<sup>10</sup>Lewis (1986, pp. 114-5).

<sup>11</sup>Objection: I have not yet shown that Lewis’ claim is mistaken. We know from the generality problem that whether a belief forming method is reliable can vary depending on how that method is described. So although I have described a case in which a necessary belief is unreliably formed under one description, it may remain true that it is reliably formed under a different description.

Reply: It seems to me that Lewis’ claim would be robbed of its interest if he were merely suggesting that for any necessary truth anyone believes, we can cook up some description or other according to which that belief is reliably formed. The same thing could be said of contingent truths. Consider the method of believing that which is true. This is an infallible method. You will never go wrong if you use it whether the belief you end up with is contingent or necessary. But if I decide on the basis of a coin toss to believe that the number of stars in Andromeda is even, then my belief is unreliably formed. The relevant method for assessing reliability is tossing a coin; not believing truly. Similarly, if I decide whether to believe God exists on the basis of a coin toss, the relevant method for assessing reliability isn’t believing a necessary truth. It is instead tossing a coin.

<sup>12</sup>See Brogaard and Salerno (2007) and Nolan (1997). That there are non-trivially true counterpossibles is one reason to posit impossible worlds. As Krakauer (forthcoming) points out, other reasons include the need to make sense of epistemic possibility, the content of the beliefs of non-ideal agents, and (arguably) distinct propositions that pick out the same set of possible worlds.

CP1 through CP3 are not trivially true. But, assuming God does not exist, they are counterpossibles. And the standard analysis judges them to be trivially true. So the standard analysis is false. It needs to be extended to account for counterpossibles<sup>13</sup>.

## 5 The Strangeness of Impossibility

Once it is recognized that the Lewisian analysis must be modified, it seems plausible that the modification should include a closeness relation that discriminates between worlds that are possible and worlds that are impossible. One, therefore, might target (2) by adding:

**(SOI):** It is of the first importance to avoid violations of metaphysical law.

The other respects of similarity would then each be demoted one slot in importance.

With (SOI) in hand, it seems possible to undermine (2). The argument for (2) relied on the idea that the Lewisian closeness relation requires nothing but match with respect to spatio-temporal fact and the avoidance of miracles. And, given Isolation, a world at which our moral beliefs are false but everything else is the same is a world with perfect match of spatio-temporal fact and without any miracles. So such a world is very close to the actual world. It is on this basis that I was able to argue that our moral beliefs are insensitive and unsafe. However, if (SOI) is true, then one might doubt this. Any world at which our moral beliefs are false is an impossible world. There are some possible worlds that are very far away from the actual world. And, given (SOI), any impossible world will be further from the actual world than any possible world. So a world at which our moral beliefs are false is a very distant world. And (2) is false.

This objection is inadequate. Even with the addition of (SOI), my defense of (2) can be sustained. First, consider whether our moral beliefs are sensitive. This is evaluated by going to the closest world at which our moral beliefs are false.  $W_0$ , the impossible world at which the moral realm is different but our moral beliefs remain the same, includes a violation of metaphysical law. But by all other measures it is as close to the actual world as can be.  $W_1$ , an impossible world at which the moral realm is different but our moral beliefs are changed to reflect this, violates metaphysical law as well. But, since our beliefs are different, it also violates exact match of particular matter of fact. So worlds like  $W_1$  are even more distant than  $W_0$ . So, although adding (SOI) yields the result that  $W_0$  is a distant world, it is still the closest of any worlds at which the moral

---

<sup>13</sup>If you don't like the theological examples, consider Maudlin's (2007, p. 22):

**M1:** If one could trisect the angle, then one could square the circle.

**M2:** If one could construct a 23-gon, then one could square the circle.

M1 and M2 are counterpossibles. M1 is true but not trivially true. M2 may be false. Even Lewis in some places (e.g. Lewis (1993, p. 164-5)) is plausibly read as (implicitly) relying on counterpossibles with non-trivial truth-values.

realm is different. And that is all that is required to show that our moral beliefs are insensitive.

Second, the introduction of (SOI) either undermines the motivation for appealing to safety or motivates the introduction of a modification to safety according to which our moral beliefs are unsafe. If (SOI) were true, then all beliefs with necessarily true content would be trivially safe. For a world at which the moral realm is different but everything else is the same is as close to the actual world as any impossible world could be. There is perfect spatio-temporal match. There are no miracles. The only thing that keeps it from being a nearby world is that it is impossible. No impossible world can get any closer than that. So, given (SOI), the closest any impossible world can get to the actual world is still very distant. And if all impossible worlds are very distant, then all beliefs with necessarily true content are trivially safe. Notice how problematic this is for the appeal to safety. Impossible worlds were invoked in this context to rescue sensitivity and safety from triviality. But if (SOI) is true, the problem returns and all beliefs with necessarily true content are trivially safe. So if (SOI) is true, it seems that safety should be abandoned and only sensitivity should be used to diagnose modal problematicness.

Third, the safety theorist might try weakening (SOI) to avoid this problem. But note that any weakening of (SOI) that successfully avoids the problem will yield the result that our moral beliefs are unsafe. For reasons already identified, the world at which our moral beliefs are false but everything else is the same is as close to the actual world as any impossible world can get. So, on the one hand, if being that close *is not* sufficient for being a nearby world, then no impossible world is a nearby world. And the problem would return. On the other hand, if being as close to the actual world as any impossible world can get *is* sufficient for being a nearby world, then there is a nearby world at which our moral beliefs are false but we continue to hold those beliefs. And so our moral beliefs are unsafe.

Fourth, one might modify the safety condition. Some authors have moved from analyzing safety in terms of the modal status of belief to analyzing it in terms of *the modal status of the process*<sup>14</sup> that produces belief. However, even if these revisions avoid the problem<sup>15</sup>, they all articulate safety in terms of the reliability of processes. And whether our moral beliefs are safe ends up hinging on whether we arrive at such beliefs by a process that is reliable in some class of nearby worlds that includes the actual world. And, as I argue below, our moral beliefs are not reliably formed in the actual world. So the relevant modification would yield the result that our moral beliefs are unsafe.

---

<sup>14</sup>Hawthorne (2007) discusses “methods” of belief production. Prichard (2007) discusses “ways” of belief formation. Sosa (2002) and Williamson (2000) discuss “bases” of belief.

<sup>15</sup>Roland and Cogburn (2011) argue that they do not.

## 6 Contingently True Moral Beliefs Can Be Sensitive and Safe

So far we have been focusing on moral beliefs with necessarily true content. Someone might try another strategy. They might focus on moral beliefs with contingently true content together with the metaphysical law that moral properties supervene on non-moral properties. Take, for example, my belief that Hitler is evil. The content of this belief is contingently true. For Hitler might not have perpetrated the Holocaust. He might have been very different than he actually was. And if he had been different enough, he would not have been evil<sup>16</sup>. Now consider two worlds at which Hitler is not evil:

**W<sub>2</sub>:** All non-moral facts are the same as they are at the actual world. But being a Holocaust perpetrator is not evil making.

and

**W<sub>3</sub>:** All necessarily true moral facts are the same as they are at the actual world. But Hitler did not perpetrate the Holocaust.

At W<sub>2</sub> I continue to believe Hitler is evil even though he is not. After all, all the mental facts are the same as they are at the actual world. At W<sub>3</sub> I do not keep my belief that Hitler is evil. For if Hitler did not perpetrate the Holocaust, then I would not believe that he did. And if I didn't believe that, then I wouldn't continue to believe he was evil.

Now, W<sub>2</sub> departs from the actual world by violating metaphysical law. For to change the moral status of Holocaust perpetration without changing any non-moral properties is to violate the supervenience of the moral on the non-moral. And therefore metaphysical law is violated. W<sub>3</sub>, on the other hand, only departs from the actual world with respect to contingent mental and physical fact. For to change my beliefs and to change what Hitler did does not violate metaphysical law. Thus, given (SOI), W<sub>3</sub> is much closer to the actual world than W<sub>2</sub>. Therefore, if Hitler were not evil, I would not have believed that he was evil. And so my belief that Hitler is evil is sensitive. Similarly, at all nearby worlds at which Hitler is evil, I continue to believe that he is evil. For a world like W<sub>2</sub> at which he is not evil but I continue to believe he is evil is a very distant world. Thus, my belief that Hitler is evil is safe.

I do not think this objection is successful. First, the objection fails to address the worry that moral beliefs with necessarily true content are insensitive and unsafe. This failure opens up an alternative route to skepticism about contingently true moral beliefs. The source of skeptical doubt isn't that I might not know Hitler is evil because I might be mistaken about whether he actually perpetrated the Holocaust or had the character he did. The moral skeptic does not call such beliefs into question. Instead, the source of skeptical doubt is that I might not know whether *being a Holocaust perpetrator* or *having the character*

<sup>16</sup>This objection is inspired by Clarke-Doane (forthcoming, p. 13), Sturgeon (1985), and Wielenberg's (2010, p. 455) objections to various skeptical arguments.



*Hitler actually had* are evil making properties. That these are evil making properties is a necessary truth. I may know that Hitler perpetrated the Holocaust and I may know what Hitler's character was like. But I don't know that these things are evil making. Once I see this, I have a defeater for my contingently true belief that Hitler was evil. An adequate solution to the problem at hand will vindicate the belief that if Hitler did what I think he did, and if he had the character that I think he had, then Hitler was evil. The objection in question does nothing to vindicate such beliefs. So it is inadequate.

Second, the objection seems more like a counterexample to the simple account of sensitivity and safety than it does an explanation of how moral knowledge is possible. In particular, I think it makes contingently true beliefs about speculative metaphysical matters sensitive and safe in an objectionably trivial way. Suppose there is a community of people with crazy beliefs about strange ghosts. These ghosts cannot causally effect us in any way. No matter what they are like, it makes no causal difference to the physical world. And no matter what they are like, it makes no causal difference to the way in which we form our beliefs. However, what the physical world is like makes a difference to the ghosts. In particular, whenever someone is over five feet tall, a ghost pops into existence and starts following that person around. Of course, we can't see the ghosts. But they can see us. And they will continue to follow us around so long as we are over five feet tall. In addition, members of this community insist that the connection is not merely contingent. It is metaphysically necessary. The property of being-followed-by-a-ghost supervenes on the property of being-over-five-feet-tall. Necessarily, any and only people over five feet tall are followed by ghosts. So the relevant property supervenes on a causally efficacious property. And that causally efficacious property causally effects our ghosts beliefs.

Now imagine that I am a member of this community. I believe that Hitler was followed by a ghost. It seems to you that my belief is crazy even if it is true. How could I know that Hitler was followed by a ghost? After all, my belief was formed in a way that is entirely causally independent of what the ghost was like and of what its ghostly activities were. So, even if my belief is true, it is not an instance of knowledge. Suppose I then give you the very response I gave regarding my belief that Hitler was evil: I know that Hitler was followed by a ghost because my Hitler belief is sensitive and safe. Consider two worlds at which Hitler was not followed by a ghost:

**W<sub>4</sub>:** All non-ghostly facts are the same as they are at the actual world. But being over five feet tall is not ghost making.

and

**W<sub>5</sub>:** All necessarily true ghostly facts are the same as they are at the actual world. But Hitler was not over five feet tall.

Everything I said on behalf of my non-crazy Hitler moral belief can be said on behalf of my crazy Hitler ghostly belief. At W<sub>4</sub> I continue to believe that Hitler was followed by a ghost even though he was not. For all non-ghostly mental

facts are the same as they are at the actual world. And at  $W_5$  I do not keep my belief that Hitler was followed by a ghost. For if Hitler were not over five feet tall, then I would not have believed that he was. And if I didn't believe he was over five feet tall, then I wouldn't believe he was followed by a ghost.

Now,  $W_4$  departs from the actual world by violating metaphysical law. For, by hypothesis, to change the ghostly status of being over five feet tall is to violate metaphysical law.  $W_5$ , on the other hand, only departs from the actual world with respect to contingent mental and physical fact. For to change my belief and to change Hitler's height does not violate metaphysical law. Thus, given (SOI),  $W_5$  is much closer to the actual world than  $W_4$ . Therefore, if Hitler were not followed by a ghost, I would not have believed that he was. And so my belief that Hitler was followed by a ghost is sensitive. Similarly, at all nearby worlds at which Hitler was followed by a ghost, I continue to believe that he was. For a world like  $W_4$  at which Hitler is not followed by a ghost but I keep my belief is a very distant world. Thus, my belief that Hitler was followed by a ghost is safe. This seems implausible. Even if my crazy metaphysics is true, and even if I am right that ghostly properties supervene on non-ghostly properties, they are not instances of knowledge. That my beliefs satisfy the simple account of sensitivity is a reason to think the simple account needs to be modified or supplemented in some way.

To summarize: The body of literature I am criticizing takes the simple account of sensitivity (or safety) to be its test for problematic accidentalness. If it can be shown that a belief satisfies the simple account, then that is taken to show that the belief in question is not problematically accidentally true. Our moral beliefs satisfy the simple variants of these tests. So, it is argued, they are not accidentally true.

The ghost case, however, is meant to show that this line of argument is mistaken. My ghost beliefs are true in a problematically accidental way. And, since my ghost beliefs satisfy the simple accounts of sensitivity and safety, something more than the mere satisfaction of one or the other of these accounts must be required in order to secure non-accidentalness. Any acceptable refinements of (or additions to) these tests will yield the result that my ghost beliefs are problematic. And, since there is no relevant epistemic difference between the ghost case and the moral case it is designed to parody, any plausible way of refining these accounts that rules out knowledge in the ghost case will also rule out knowledge in the moral case.

For example, one might try adding a causal requirement to these accounts. This would indeed yield tests with the result that my ghost beliefs are problematic. But, by hypothesis, moral properties lack causal powers. So to add to the simple account of sensitivity (or safety) a causal requirement would be to yield an account according to which my moral beliefs are also problematic. For another example, one might try adding a reliability requirement. This might rule out knowledge in the ghost case. But, I argue below, our moral beliefs, given Isolation, are not reliability formed. So to add to the simple account a reliability requirement would be to modify it in such a way that it yields the result that our moral beliefs are problematically accidentally true.

Objection: It is true that my ghost beliefs are problematic. But I have given the wrong diagnosis about why such beliefs are problematic. I say that they are problematic because they are accidentally true. Sensitivity (or safety) is an account of problematic accidentalness. It fails to accommodate the judgment that my ghost beliefs are problematically accidentally true. So it must be modified to accommodate this judgment. However, someone might object, I am mistaken about this suggestion. There are other features a belief might have that can keep it from being knowledge. Accidentalness is just one such feature. Perhaps my ghost beliefs pass the accidentalness test while failing to pass some other test. For example, maybe the problem with my ghost beliefs is that they fail to cohere with other beliefs I have. It is lack of coherence rather than accidentalness that renders my ghost beliefs problematic. Thus, one could keep the simple accounts of sensitivity and safety as our tests for accidentalness. And since our moral beliefs pass those tests, they are not problematically accidentally true. And, the relevant moral beliefs cohere with my other beliefs while the relevant ghost beliefs do not. That is why I have knowledge in the moral case but not the ghost case.

Reply: I am skeptical of this response. It seems to me that there are no plausible diagnoses of what is wrong with my ghost beliefs that deliver the judgment that my ghost beliefs are problematic while my moral beliefs are not. Consider the diagnosis of incoherence. As things actually are, I believe that there are no causally inefficacious ghosts. So adding the belief that Hitler was followed by such a ghost to the beliefs I actually have would indeed yield a contradiction. But we are not imagining that I have the same beliefs I actually have plus some ghost beliefs. Instead, the supposition is that some of my beliefs are different than they actually are. In the imagined case, my beliefs are perfectly consistent with belief in the Hitler following ghost. I believe in the relevant ghosts. And I do not believe anything that contradicts the belief that such ghosts exist. Still, even if such causally inefficacious ghosts happen to exist; and even if they turn out to necessarily supervene on height in the relevant ways, if my beliefs about them are formed without any causal connection to the ghosts or their activities, then it is a mere accident that such beliefs are true. My perfectly consistent ghost beliefs do not count as knowledge. And neither do my perfectly consistent moral beliefs. Is the diagnosis, instead, supposed to be that my ghost beliefs are unreliably formed? If so, then I argue below that our moral beliefs are unreliably formed as well. So again I lack moral knowledge. The challenge I would issue to those who are sympathetic to the strategy in question is this: Find some plausible way to diagnose what is wrong with my ghost beliefs that doesn't also deliver the judgment that something is wrong with my moral beliefs. Coherence cannot play the role of such a diagnosis. Reliability cannot play that role. Sensitivity and safety cannot play that role. Every one of these diagnoses either (i) fails to deliver the judgment that my ghost beliefs are problematic or (ii) fails to exonerate my moral beliefs. Thus, the appeal to moral beliefs with contingently true content cannot support an adequate objection to the Argument from Isolation.

## 7 No Rights Without Beliefs

One might object to (2) and (3) by appealing to<sup>17</sup>:

**(Rights):** Necessarily, a being has rights if and only if that being is capable of having beliefs.

(Rights) seems to spell trouble for (2) and (3). First, if (Rights) is true, then our moral beliefs are sensitive. The idea<sup>18</sup> is this: Go to the nearest world at which I lack rights. Given (Rights), it will be a world at which I lack the capacity to form beliefs. And so, at such a world, I will not believe that I have rights. Thus, if I did not have rights, I would not believe that I did. So my belief is sensitive and (2) is false. Second, if (Rights) is true, then our moral beliefs are reliably formed. The idea<sup>19</sup> is that since (Rights) is true, if I believe I have rights, then I do have rights. And if this is so, then the process by which I come to believe I have rights is reliable. So premise (3) is false.

It seems to me that this objection is unsuccessful. First, given the relevant assumptions, my belief does indeed satisfy the simple account of sensitivity. But this is for the same reason that my belief that Hitler was evil satisfies the simple account. The property of having rights supervenes on the non-moral property of having the cognitive faculties required to form beliefs. But, following the reasoning in Section 6, we can set up a parody case in which the property of being followed by a ghost supervenes on the non-ghostly property of having the cognitive faculties required to form beliefs. So if I were not followed by a ghost, then I would lack the capacity to form beliefs. And therefore I would not believe I was followed by a ghost. Thus, my ghostly belief would satisfy the simple account of sensitivity. But this is not sufficient to secure knowledge in the case of my ghostly belief. So, it is not sufficient to secure knowledge in the case of my moral belief either.

Second, this response still has the result that moral beliefs with necessarily true content are not sensitive. Take, for example, (Rights). It is a necessary truth. So consider two worlds at which it is false. One is the impossible world at which (Rights) is false but everything else is the same. The other is the impossible world at which (Rights) is false but I lack the capacity to form beliefs. Following the reasoning from Section 5, the later world is further from the actual world than the former. So, my belief is insensitive.

Third, some refinements of the simple account of sensitivity deliver the judgment that my belief *is not* sensitive<sup>20</sup>.

---

<sup>17</sup>This is Wielenberg's (2011) objection to skeptical arguments due to Joyce (2001) and (2006), Ruse (1986), and especially Street (2006).

<sup>18</sup>See Wielenberg (2011, p. 455).

<sup>19</sup>Here is how Wielenberg (2011, pp. 449-50) puts it: "[I]f rights exist at all, their presence is guaranteed by the presence of certain cognitive faculties. The cognitive faculties in question are either the very ones required to form beliefs about rights or are closely linked to such faculties. If you think you possess moral barriers, then you do... Therefore... the processes that ultimately generate, say, the belief that one has... right[s]... are significantly reliable."

<sup>20</sup>Consider the refinement Williamson (2000) offers on behalf of the sensitivity theorist: S doesn't know P when there is no Q such that:

Fourth, consider whether my belief is reliably formed: suppose you are agnostic about whether you have rights. You toss a coin. Someone has arranged things so that if the coin lands heads you will instantly believe you have rights and if it lands tails your brain will instantly be damaged in such a way that you lose the capacity to have beliefs. Now consider:

**(Damaging Coin):** The coin lands heads. Your belief that you have rights is due to the outcome of the coin toss.

The earlier reasoning about (Rights) and reliability applies equally to this case. Given (Rights), if you believe you have rights, then you do. If this, by itself, is sufficient for reliability, then your belief in (Damaging Coin) would be reliably formed. But your belief is not reliably formed.

Go back to the ghost case: Imagine that (necessarily) a causally inefficacious ghost follows one around if and only if one has the capacity to form beliefs. You are agnostic about whether such ghosts exist. But you flip a coin that will either damage your brain in such a way that you lose your capacity to have beliefs or you form the belief that such ghosts exist. Even if your belief turns out to be true, it was not reliably formed. So, the appeal to (Rights) is not sufficient to establish reliability.

## 8 Survival is Good

One might object to (2) and (3) by appealing to<sup>21</sup>:

**(SG):** Survival is good.

and

**(Select):** Evolution selects survival. If evolution selects survival, then evolution selects beliefs with content closely related to (SG). Evolution could not have easily failed to select survival.

What is it to be “closely related” to (SG)<sup>22</sup>? For present purposes, it will be sufficient to note two things. First, if (SG) is true, then propositions that are closely related to (SG) are true. Second, many of our most important moral beliefs have content that is closely related to (SG).

- 
- (i) S believes that Q,
  - (ii) S would not have believed Q if Q had been false,
  - (iii) Q entails P, and
  - (iv) Not-P does not explain how S could falsely believe that Q.

Now, let S be me and let P be the belief that I have rights. For any Q you pick, condition (iv) will be unsatisfied. After all, if P were false, then, given (Rights), I wouldn’t believe Q since I wouldn’t have any beliefs at all. So P’s falsity does not explain how I could falsely believe Q. Therefore, my belief that I have rights is not sensitive.

<sup>21</sup>This is Enoch’s (2009) and (2011, pp. 168-76) objection to both Street’s (2006) skeptical argument and a skeptical argument of his own. Skarsaune (2010) and Behrends (forthcoming) defend similar objections. My criticism is meant to extend to Behrends’ and Skarsaune’s variants of Enoch’s objection.

<sup>22</sup>See Enoch (2011, p. 169).

This supports an argument against (2) and (3). Regarding (2)<sup>23</sup>: If evolution gives us any moral beliefs at all, those beliefs will have content closely related to (SG). And since (SG) is true, the moral beliefs evolution gives us will be true. This is a process that results in a high ratio of true to false beliefs. So our moral beliefs are reliably formed. Regarding (3)<sup>24</sup>: (SG) and our closely related moral beliefs are necessarily true. So they are trivially sensitive and safe. There is a way in which evolution could have failed to select true moral beliefs. But this would happen in a bizarre, very distant possible world in which evolution somehow did not select for beliefs closely related to (SG). The existence of a distant possible world like that does not in any way threaten the sensitivity or safety of our moral beliefs. So our moral beliefs are not modally problematic.

It seems to me that this objection is unsuccessful. Regarding sensitivity and safety: In Section 4, we saw that the mere necessary truth of our moral beliefs is insufficient to yield the result that they are not modally problematic. And nothing about the present objection casts doubt on the argument from Section 5 that, given (SOI), our moral beliefs are insensitive and unsafe.

Regarding reliability: To the extent that our beliefs in this case satisfy the simple account of reliability, that is a reason to think the simple account of reliability needs modification. Suppose you decide to toss a coin. Someone has arranged things so that if it lands heads, then you will instantly form a bunch of beliefs with content closely related to (SG). If it lands tails, you will instantly be destroyed. Now consider:

**(Destroying Coin):** The coin lands heads. Your beliefs with content closely related to (SG) are due to the outcome of the coin toss.

The reasoning about (SG) and (Select) applies equally to this case. If the destroying coin gives you any beliefs at all, those beliefs will have content closely related to (SG). And, since (SG) is true, those beliefs will be true. If this, by itself, is sufficient for reliability, then your belief in (Destroying Coin) would be reliably formed. But your belief is not reliably formed.

Again, compare this to an example in which the relevant moral beliefs are replaced with ghost beliefs. Suppose that unbeknownst to anyone previously, it is a necessary truth that:

**(GF):** Causally inefficacious ghosts follow people around.

Imagine you flip a coin that will either destroy you or give you a bunch of beliefs closely related to (GF). If the coin gives you any beliefs at all, they will

---

<sup>23</sup>See Enoch (2011, p. 168).

<sup>24</sup>Enoch (2012, p. 172-4) puts it this way: “One way of putting the [objection to moral knowledge]... is in terms of counterfactual robustness.... It is not completely clear [however] that the thought... here can be made fully coherent. The evolutionary “aim” would have had to not be of any value. And how could that be? Fundamental normative truths are presumably necessary in a fairly strong sense.... It is not immediately obvious, then, how to state the problem.... [A]s a conceptual matter, *evolution* could not have had a very different “aim” in “mind....” The possible worlds in which... there is no correlation between our normative beliefs and the normative truths are quite far.”

be beliefs with content closely related to (GF). But, even if (GF) is a necessary truth, your ghostly beliefs are not reliably formed. In the same way, the appeal to (SG) and (Select) is not sufficient to establish reliability.

## 9 Adequate Grasping

Now we turn to a series of objections that toss out sensitivity and safety and appeal only to reliability<sup>25</sup>. All then hinges on premise (3). Is there an explanation of how our moral beliefs could be reliably formed (even if they are insensitive and unsafe)? One such explanation is that there is a special way of understanding moral properties. When we understand something in this way, our beliefs are guaranteed to be true. And if they are guaranteed to be true, then they are reliably formed.

The plausibility of this explanation will depend on how its details are filled out. For example, the mere guaranteeing of truth is not sufficient for reliability. Consider:

**(Schmunderstanding):** A subject schmunderstands something if and only if that subject's beliefs about that thing are true.

There is a way in which schmunderstanding guarantees truth. By definition, if a subject schmunderstands something, then that subject's beliefs about that thing are true. But this does not preclude unreliably formed belief: Imagine you have a random truth value assigner (RTA). When you type a sentence into the RTA, it randomly produces a judgment about the truth value of the proposition expressed by that sentence. Now suppose you type sentences expressing propositions about some moral property into the RTA. The RTA then randomly produces judgments about the truth values of those propositions. You then conform your beliefs to the deliverances of the RTA. Finally, suppose the deliverances of the RTA just happen to match the truth values of the relevant propositions. So all your beliefs about the moral property in question are true. You schmunderstand the moral property. So there is a way in which your beliefs are guaranteed to be true. However, they are not reliably formed.

Now, consider a way of filling out the account of understanding<sup>26</sup>:

**(Adequate Grasp):** A subject adequately grasps a property if and only if that subject has a consistent, clear, and determinate grasp of that property.

Consistent grasping is left unanalyzed. For a subject to have a clear grasp of a property is for that subject to be able to successfully distinguish between that property and other similar properties. For a subject to have a determinate grasp of a property is for there to be few or no cases about which the subject is unable to successfully determine whether that property is instantiated. Grasping comes in degrees. If a person consistently, clearly, and determinately grasps a property

---

<sup>25</sup>The strategy considered in this section is due to Huemer (2005) and Bealer (2000).

<sup>26</sup>This objection is due to Huemer (2005, p. 125).

to a high degree, then that person adequately grasps that property. Otherwise, the property is grasped but not adequately grasped.

This supports an objection to (3): We adequately grasp moral properties. Adequate grasping guarantees truth. The way in which adequate grasping guarantees truth secures reliably formed belief. So our moral beliefs are reliably formed.

Adequate grasping does guarantee a large body of truth. This follows from the definitions of ‘clarity’ and ‘determinacy’. To have a clear grasp of a property just is to be able to evaluate the truth of property identity statements about it. So you must be able to tell whether goodness is identical to pleasure, for example. Otherwise you wouldn’t have a clear grasp of goodness. And to have a determinate grasp of a property just is to be able to tell what has the property and what doesn’t. So you must be able to tell whether Mother Teresa has goodness, for example. Otherwise your grasp of goodness would not be determinate.

However, adequate grasping does not secure reliability. Go back to the RTA. Suppose you enter sentences expressing propositions about goodness into the RTA. Some of these are sentences of the form ‘Goodness is identical to  $P$ ’ where  $P$  is some property such as pleasure. The others are of the form ‘Goodness is had by  $x$ ’ where  $x$  is some potential bearer of goodness such as Mother Teresa. Suppose that for each proposition expressed by a sentence of one of these forms, you enter the relevant sentence into the RTA. Then you run the RTA and conform your beliefs to its deliverances. Finally, suppose the deliverances of the RTA just happen to match the truth-values of the relevant propositions.

You have a consistent grasp of goodness. Furthermore, you can successfully distinguish between goodness and any other property. So your grasp is clear. Finally, you can successfully determine what has goodness and what doesn’t. So your grasp is determinate. Thus, you adequately grasp goodness. But your beliefs about it are unreliably formed.

Someone might object: Reliability is guaranteed only when beliefs are *caused* by an adequate grasp<sup>27</sup>. In the case under consideration, your beliefs were not caused by an adequate grasp. They were caused by the RTA. So adequate grasping really does ensure reliability. However, consider a parallel defense of the explanation in terms of *schmunderstanding*: Reliability is guaranteed only when beliefs are caused by *schmunderstanding*. In the case under consideration, your beliefs were not caused by *schmunderstanding*. They were caused by the RTA. So *schmunderstanding* really does ensure reliability.

It should be clear what is wrong with this. To *schmunderstand* goodness just is to have true beliefs about goodness. Is the suggestion that you need to have true beliefs and then those beliefs must cause themselves? Assuming such a thing is possible, it is perfectly compatible with the RTA story. You form your beliefs about goodness using the RTA. So you *schmunderstand* goodness. And then those beliefs go on to cause themselves. Is the suggestion that you

---

<sup>27</sup>Huemer (2005, p. 125).



start with a few true beliefs and then go on to infer a bunch of other beliefs from those? That is still compatible with the RTA story. You got some true beliefs from the RTA. Then those beliefs caused you to form more true beliefs. Still, your beliefs are not reliably formed. The same sort of worry applies to the explanation in terms of adequate grasping. To have an adequate grasp of goodness just is to be able to successfully sort out true property identities about goodness and attributions of goodness from false ones. After you run the RTA, you are able to do such sorting. So you adequately grasp goodness. If you go on to form beliefs about goodness based on that adequate grasp, then those beliefs are still unreliably formed. Is the idea that there is something more to adequate grasping than mere successful sorting of true beliefs from false? I have no objections to that. But that something more is not a part of the definition of adequate grasping we have been given. And I do not know what to add to the account that can fix the problem<sup>28</sup>.

## 10 Plenitude

One might deny (3) by appealing to<sup>29</sup>:

**(Plenitude):** For any consistent assignment of truth-values to propositions expressed by sentences of the form ‘ $P$  is identical to  $P^*$ ’ or ‘ $P$  is instantiated in  $C$ ’, there is a property that corresponds to that assignment.

Go back to the case in which you form beliefs about goodness using the RTA. Given (Plenitude), your moral beliefs are reliably formed even in this case. For any consistent assignment of truth-values the RTA might generate, there is a property that corresponds to that assignment. And, according to the response in question, your beliefs are about whatever property most closely corresponds to those beliefs. For this reason, consistent outputs of the RTA produce a high ratio of true to false moral beliefs. All are true and none are false. So your moral beliefs, however you got them, are reliably formed.

This objection entails the truth of an implausible form of moral subjectivism<sup>30</sup>. Consider Kant. He believes that wrongness is identical to the property of treating someone merely as a means, that it is distinct from failing to maximize pleasure, and that it would be wrong to harvest the organs of one healthy person to save others. Now consider Mill. He believes wrongness is identical to the property of failing to maximize pleasure, that it is distinct from treating someone merely as a means, and that an organ harvest would not be wrong. Given (Plenitude), there is a property that corresponds to Kant’s beliefs and a distinct property that corresponds to Mill’s. And if their beliefs are

<sup>28</sup>Note that Huemer’s (2005, p. 125) suggestion is not that moral propositions “are ‘analytic’, or true by virtue of the definitions of the relevant concepts, since in [his] view hardly any concepts are definable.”

<sup>29</sup>See Huemer (2005, p. 126). Huemer’s response, on this interpretation, is similar to Balaguer’s (1998, pp. 48-9) response to a parallel worry in the epistemology of mathematics.

<sup>30</sup>Huemer (2005, pp. 48-65) offers arguments like these in another context.

about different properties, then Kant's beliefs do not conflict with Mill's. But this is an incorrect result. Kant's beliefs do conflict with Mill's. They have competing, incompatible theories of wrongness. The appeal to (Plenitude) does not accommodate disagreement.

Consider Nietzsche. He believes that goodness is identical to the will to power and that it is had by those who dominate and control others. Given (Plenitude), there is a property that corresponds to Nietzsche's beliefs, that property is what his beliefs are about, and so those beliefs are true. This is another incorrect result. Nietzsche has a mistaken theory of goodness. He is wrong about what it is and about what has it. The appeal to (Plenitude) does not accommodate error.

## 11 Moral Skepticism Implies Everything Skepticism

Another objection: Some authors maintain that arguments for moral skepticism prove too much<sup>31</sup>. In some cases the accusation is that worries about moral knowledge extend to all *a priori* knowledge. In other cases the accusation is that the reasoning in support of moral skepticism equally supports perceptual skepticism or that such arguments are mere applications of standard skepticism to morality.

I have different responses to different accusations. First, I think that the skeptical considerations raised here *do* extend to all *a priori* beliefs. But I think it will take hard work that I cannot do here to defend such a claim. However, if the antiskeptical wants to do all the hard work for me, and show that the worries raised here extend to all *a priori* beliefs, then I will gladly accept the result. For it is one that I already believe. Of course, the anti-skeptical will then point out that since my argument is an *a priori* one, it is self-defeating. Again, I agree. The argument is self-defeating. And something, somewhere must be wrong with it. But as I point out in the next section, showing that something, somewhere must be wrong with a skeptical argument is not an adequate solution to the skeptical puzzle motivated by that argument. I don't want to know that some premise or other of the argument is false. I want to know exactly which premises are false and why. Second, these skeptical worries do not extend to perceptual beliefs and are not a mere application of standard skepticism to morality. The argument appeals to standard tools epistemologists use to evaluate beliefs. Simply look at the wide body of literature that employs these tools without skeptical consequences. Here is another way to put the point: if the worry about moral knowledge really is just an instance of more general skepticism, then the standard tools used to resolve perceptual skepticism (e.g. reliability, sensitivity, and safety) should provide equally good resolutions to

---

<sup>31</sup>This objection is inspired by points Cuneo and Shafer-Landau (forthcoming), Hanna (2010), Huemer (2005), Vavova (forthcoming), and Wielenberg (2010) make against skeptical arguments due to Bedke (2009), Joyce (2006), and Street (2006).

the worry about moral skepticism. I have argued that they don't. And if I am right, then either (i) the problem is not just an uninteresting instance of the worry about perceptual skepticism or (ii) it is an instance of the same problem but it has somehow gone unnoticed that the standard responses to perceptual skepticism don't work. And, while (ii) would be a much more interesting result than (i), it is also far less plausible. So it does not seem to me that the worries raised here lead to perceptual skepticism.

## 12 Moorean Modesty

A final objection: We should be more confident about our moral knowledge than we are about the speculative epistemological tools I have employed here. So if we are forced to toss out either the speculative epistemological tools or Moorean modesty, then the speculative epistemological tools should go.

I am not a moral skeptic. And I am in some ways sympathetic to this response. But I am looking for a solution to the puzzle. And it seems to me that Moorean modesty, by itself, cannot support an adequate solution for two reasons. First, a complete solution to a skeptical puzzle must explain how knowledge about the relevant domain of inquiry is possible. Moorean modesty doesn't do this. It just gives us a reason to think that something somewhere must be wrong with the skeptical argument. So, while I am in some ways sympathetic to appeals to Moorean modesty, I don't think such an appeal is enough to solve the problem. Second, the puzzle is not about moral knowledge apart from all else. It is about the conjunction of moral knowledge with Isolation. And Isolation is itself a speculative metaphysical thesis. So the choice is not between speculative epistemological tools and Moorean modesty. The choice is instead between speculative epistemological tools, on the one hand, and Moorean modesty plus a speculative thesis about the metaphysics of moral properties, on the other. Maybe the right choice is to reject the speculative metaphysical thesis while keeping both Moorean modesty and the speculative epistemological tools. Or maybe not. I am perplexed.

## 13 Conclusion

I would like to end with three positive suggestions. First, those of us who are interested in moral epistemology may benefit by thinking about impossibility. In particular, it may be worth trying to come up with a principle that tells us exactly when impossible examples are relevant to the evaluation of a theory. We cannot dismiss an example merely because it is impossible. But it also seems that we cannot appeal to just any impossible case whatsoever when testing a theory. There must be some middle ground. It seems to me that a theory about when impossible examples are relevant could help us think about moral

epistemology<sup>32</sup>. Second, I want to recommend a heuristic: when presented with a moral epistemology, consider what happens when the proposal is applied to non-moral (and non-mathematical) examples (e.g. inefficacious ghosts or distant planets). Applying the theory in this way might yield a counterexample. Third, I want to recommend another heuristic: when presented with a moral epistemology according to which our moral beliefs are not problematically accidentally true because they were formed by a process with features  $F_1$ ,  $F_2$ , ..., and  $F_n$ , look for other processes (e.g. certain kinds of coin flipping or an RTA) that also have  $F_1$ ,  $F_2$ , ..., and  $F_n$  but might yield beliefs that are true in a problematically accidental way. This may show that the process in question does not, in fact, rule out problematically accidental true belief<sup>33</sup>.

## References

- [1] Balaguer, M. (1998). *Platonism and Anti-Platonism in Mathematics*. Oxford University Press.
- [2] Bealer, George (2000). 'A Theory of the *A Priori*'. *Pacific Philosophical Quarterly*, 81, 1-30.
- [3] Bedke, M. (2009). 'Intuitive Non-Naturalism Meets Cosmic Coincidence'. *Pacific Philosophical Quarterly*, 2, 188-209.
- [4] Behrends, J. (forthcoming). 'Meta-Normative Realism, Evolution, and Our Reasons to Survive'. *Pacific Philosophical Quarterly*, 1-17.
- [5] Brogaard, B. & Salerno, J. (2007). 'Why Counterpossibles are Non-Trivial'. *The Reasoner*, 1, 5-6.
- [6] Clarke-Doane, J. (forthcoming). 'Moral Epistemology: The Mathematics Analogy'. *Nous*, 1-26
- [7] Conee, E. & Feldman, R. (1998). 'The Generality Problem for Reliabilism'. *Philosophical Studies*, 89, 1-29.
- [8] Cuneo, T. & Shafer-Landau, R (forthcoming). 'The Moral Fixed Points: New Directions for Moral Nonnaturalism'. *Philosophical Studies*, 1-45.
- [9] DeRose, K. (2010). 'Insensitivity is Back, Baby!'. *Philosophical Perspectives*, 24, 161-187

---

<sup>32</sup>It also seems to me that such a theory could help us think about issues in philosophy of religion. See the impossible examples discussed in Hill (2014).

<sup>33</sup>Thanks to Louise Antony, Lynne Baker, Phil Bricker, Justin Dealy, Maya Eddon, Fred Feldman, Ed Ferrier, Cameron Gibbs, Matt Gifford, Pete Graham, Bob Gruber, Ryan Hebert, Dennis Kavlakoglu, Casey Knight, Hilary Kornblith, Jordan Kroll, Victor Ma, Trish Magalotti, Chris Meacham, Luis Oliveira, Kristian Olsen, Ben Rancourt, Julie Rose, Kim Soland, Miles Tucker, Aaron Washington, Hayley Webster, and three anonymous referees for *Erkenntnis*.

- [10] Enoch, D. (2009). ‘The Epistemic Challenge to Metanormative Realism: How Best to Understand it, and How to Cope With it’. *Philosophical Studies*, 148, 413-438.
- [11] Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- [12] Enoch, D., Spectre L., & Fisher T. (2012). ‘Statistical Evidence, Sensitivity, and the Legal Value of Knowledge’. *Philosophy and Public Affairs*, 40, 197-224.
- [13] Field, Hartry (1989). *Realism, Mathematics, & Modality*. Basil Blackwell.
- [14] Hanna, N. (2010). Cosmic Coincidence and Intuitive Non-Naturalism. *Journal of Ethics and Social Philosophy*, 1-4.
- [15] Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.
- [16] Hawthorne, J. (2007). ‘A Priority and Externalism’. In S. Goldberg (ed.) *Internalism and Externalism in Semantics and Epistemology*. Oxford University Press.
- [17] Hill, S. (2014). ‘Giving up Omnipotence’. *Canadian Journal of Philosophy*, 44, 97-117.
- [18] Huemer, M. (2005). *Ethical Intuitionism*. Palgrave Macmillan.
- [19] Joyce, R. (2001). *The Myth of Morality*. Cambridge University Press.
- [20] Joyce, R. (2006). *The Evolution of Morality*. MIT Press.
- [21] Krakauer, B. (forthcoming). ‘What are Impossible Worlds?’ *Philosophical Studies*, 1-19.
- [22] Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- [23] Lewis, D. (1979). ‘Counterfactual Dependence and Time’s Arrow’. *Nous*, 13, 455-76.
- [24] Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- [25] Littlejohn, C. (2011). ‘Moral Knowledge by Intuition’. In J. Hernandez (ed.), *The New Intuitionism*. Continuum.
- [26] Maudlin, T. (2007). *The Metaphysics Within Physics*. Oxford University Press.
- [27] Nolan, D. (1997). ‘Impossible Worlds: A Modest Approach’. *Notre Dame Journal for Formal Logic*, 38, 325-527.
- [28] Parfit, D. (2011). *On What Matters: Volume Two*. Oxford University Press.

- [29] Prichard, D. (2007). 'Anti-Luck Epistemology', *Synthese*, 158, 277-97.
- [30] Roland, J. & Cogburn, J. (2011). 'Anti-Luck Epistemologies and Necessary Truths'. *Philosophia*, 39, 547-561.
- [31] Ruse, M. (1986). *Taking Darwin Seriously: A Naturalistic Approach to Philosophy*. Blackwell.
- [32] Shafer-Landau, R. (2012). 'Evolutionary Debunking, Moral Realism and Moral Knowledge'. *Journal of Ethics and Social Philosophy*, 7, 1-38.
- [33] Skarsaune, S. (2011). 'Darwin and Moral Realism: Survival of the Iffiest'. *Philosophical Studies*, 153, 229-243.
- [34] Sosa, E. (1999). 'How to defeat opposition to Moore'. *Philosophical Perspectives*, 13, 141-153.
- [35] Sosa, E. (2002). 'Tracking, Competence and Knowledge'. In P. Moser (ed.), *The Oxford Handbook of Epistemology*. Oxford University Press.
- [36] Street, S. (2006). 'A Darwinian Dilemma for Realist Theories of Value'. *Philosophical Studies*, 127, 109-166.
- [37] Vavova, K. (forthcoming). 'Debunking Evolutionary Debunking'. *Oxford Studies in Metaethics*.
- [38] Wielenberg, E. (2010). 'On the Evolutionary Debunking of Morality'. *Ethics*, 120, 441-464.
- [39] Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.