MORAL KNOWLEDGE BY DEDUCTION

Declan Smithies

The Ohio State University

How is moral knowledge possible? This paper defends the *anti-Humean thesis* that we can acquire moral knowledge by deduction from wholly non-moral premises.[1]

It is not controversial, of course, that we can acquire moral knowledge by deduction from moral premises. What is more controversial is the suggestion that we can acquire moral knowledge by deduction from non-moral premises alone. Indeed, this conflicts with *Hume's Law*, which states that we cannot infer an 'ought' from an 'is', since it is "altogether inconceivable how this new relation can be a deduction from others, which are entirely different from it" (Hume, 1739, 3.1.1). There is some debate about whether Hume really endorsed Hume's Law, or merely raised it as a challenge to be surmounted, but we can set this exegetical debate aside.[2] Either way, Hume's Law is widely endorsed in contemporary moral epistemology.

This paper explores the prospects for an inferentialist theory of deductive moral knowledge that rejects Hume's Law. My defense of the anti-Humean thesis relies on the following assumptions:

(1) *Minimal moral realism*: There are some moral facts.

(2) *Moral supervenience*: For any moral fact M, there is some non-moral fact, N, such that N entails M: that is, it's necessary that if N, then M.

---

[1] I am indebted to Aaron Zimmerman's (2010, Ch. 5) pioneering defense of this view, and also to ideas in the work of Kieran Setiya (2012) and Ralph Wedgwood (2019), as I explain in §11.

[2] See Qu (2019) for an overview of the debate and a defense of the thesis that Hume rejects both deductive and probabilistic inferences from 'is' to 'ought'.

(3) *General anti-skepticism*: We can know some non-moral facts, including some non-moral fact N that entails some moral fact M.

We can make these background assumptions in moral metaphysics and general epistemology without settling disputed questions in moral epistemology. For instance, we can assume there are some moral facts without prejudging whether and, if so, how we know them. Similarly, we can assume we have some non-moral knowledge without prejudging the existence and nature of moral knowledge.

The main thesis of the paper is that we can know moral facts by deduction from our knowledge of non-moral facts on which they supervene. In other words, we can know moral facts by making deductions of the form: N; therefore, M. Consider Gilbert Harman's (1977, 4) example: when you see some teenagers torturing a cat just for fun, how do you know that what they are doing is wrong? My answer is that you can acquire this moral knowledge by performing the following deduction:

Premise: Those teenagers are torturing a cat just for fun.

Conclusion: Therefore, they are doing something wrong.

According to the anti-Humean thesis, you can perform this deduction without antecedent knowledge of the general moral principle that it's always wrong for anyone to torture cats just for fun. Indeed, the direction of explanation is the other way around. Our knowledge of general moral principles is explained by our capacity to deduce moral conclusions from non-moral premises, not vice versa.

Let's apply our three assumptions to Harman's example. First, the conclusion is true, since we're assuming *minimal moral realism*: if there are any moral facts, this is surely one of them. Second, the premise entails the conclusion, since we're assuming *moral supervenience*: if the teenagers are acting wrongly, then it's a necessary truth that anyone who does the same thing in the same circumstances is

also acting wrongly. Third, the premise can be known, since we're assuming *general anti-skepticism*: if we can know anything at all, then we can know non-moral facts about what people are doing and why they're doing it. Anyone who suspects that 'torturing' is a thick moral term may substitute their chosen non-moral description of the action (e.g. 'intentionally inflicting suffering').

What the anti-Humean thesis adds is that you can know that the teenagers are doing something wrong by deduction from the premise that they are torturing a cat just for fun. If you know the non-moral premise, then you're in an epistemic position to know the moral conclusion by deduction from this known premise. And if you exploit your epistemic position by competently performing the deduction, then you thereby know the conclusion. This follows from the more general principle that knowledge is closed under competent deduction:

> *The Closure Principle for Knowledge*: Necessarily, if p entails that q, and you know that p, and you competently deduce that q from the known premise that p, then you thereby know that q.[3]

In general, competent deduction is a way of extending your knowledge. Moral deduction is no exception to this general rule.

The first half of the paper builds a positive case for the anti-Humean thesis by appealing to its explanatory power. First, it blocks the regress argument for moral skepticism without collapsing into moral foundationalism (§1). Second, it is economical, rather than extravagant, in the sense that it avoids positing distinctive sources of moral evidence (§2). Third, it explains how our moral knowledge about cases depends causally and epistemically on our non-moral knowledge about those cases (§3).

---

[3] See Williamson (2000, 117-18) and Hawthorne (2004, 31-6) in defense of the closure principle. We can focus on single-premise closure for now, since we're considering single-premise deductions, but I'll extend the anti-Humean thesis to multi-premise deductions in §8.

Fourth, it explains how we can acquire moral knowledge about particular cases without relying on antecedent knowledge of general moral principles (§4). Fifth, it explains how we can acquire a priori knowledge of general moral principles by deducing moral conclusions from non-moral premises that are assumed, rather than known through experience (§5). And sixth, it explains why we cannot be justified in having or acting upon morally repugnant beliefs (§6).

The second half of the paper defends the anti-Humean thesis against a series of objections. First, the *over-intellectualization problem*: is our moral knowledge about cases based on any psychological process of inference at all (§7)? Second, *the no-entailment problem*: is our moral knowledge entailed or merely made probable by our non-moral knowledge (§8)? Third, *the wrong-entailment problem*: if our moral knowledge is entailed by our non-moral knowledge, is it the right kind of entailment to ground knowledge by deduction (§9)? And fourth, *the problem of hard cases*: can we know moral facts by deduction in hard cases where we encounter moral disagreement or uncertainty (§10)? Finally, the paper concludes by situating the anti-Humean thesis defended in this paper within the recent literature in moral epistemology (§11).

1. The Regress Argument

The first motivation for the anti-Humean thesis is that it provides an attractive way to block the regress argument for moral skepticism.[4]

The regress argument purports to show that *inference* cannot be our most basic source of moral knowledge. The key premise is that we cannot know moral conclusions by inference from non-moral premises alone: as Hume's Law states, we cannot infer an 'ought' from an 'is'. To know a moral conclusion by inference, we must infer it from at least one moral premise. But we cannot know a

---

[4] See Sinnott-Armstrong (1996) for a detailed version of the regress argument for moral skepticism and see Zimmerman (2010, Ch. 5) for the response that rejects Hume's Law.

moral conclusion by inference from a moral premise unless we know the moral premise to be true. But then how can we know the moral premises from which our moral conclusions are inferred? It cannot be that all our moral knowledge is based on inference, the argument continues, since the chain of inference would either continue forever or loop around in a circle. Plausibly, however, we cannot acquire knowledge by chains of inference that are infinitely regressive or viciously circular. Hence, the chain of inference must come to an end somewhere.

The interim conclusion is that inference cannot be our most basic source of moral knowledge. To know moral conclusions by inference, we must have some other way of knowing the moral premises from which we infer them. So, either moral knowledge is impossible, or it must be possible to acquire moral knowledge without inference. Moral skepticism follows given the further premise that it's impossible to acquire moral knowledge without inference.

Moral foundationalism blocks the regress argument for moral skepticism by claiming that we can acquire moral knowledge without inference. Unfortunately, this claim is hard to sustain. No doubt, we can acquire some non-inferential knowledge on the basis of sensory perception. Plausibly, however, we cannot acquire non-inferential knowledge on the basis of sensory perception unless it represents the known fact in question. And it is doubtful that sensory perception represents moral facts at all. For example, if it seems to you that the teenagers are doing something wrong by torturing a cat, when they are really just acting in a play, you are not thereby suffering a visual illusion. Arguably, therefore, sensory perception cannot give us moral knowledge without inference.[5]

Moreover, sensory perception is the only plausible candidate to play a foundational epistemic role in justifying beliefs without standing in need of justification. Some argue that we can acquire non-inferential moral knowledge on the basis of quasi-perceptual experiences of moral intuition or moral

---

[5] See McGrath (2004), McBrayer (2010), and Audi (2013) in defense of moral perception.

emotion.[6] But even given that such experiences can represent moral facts, it is doubtful that they can play the foundational epistemic role of sensory perception. Imagine someone whose moral intuitions and emotions are wildly deviant: for example, when they see the teenagers torturing the cat, their intuitive and emotional reactions represent that the teenagers are morally required to act as they do. Is this enough, in the absence of defeaters, to justify the morally repugnant belief that the teenagers are acting as they should? As I'll explain in §6, this result is hard to stomach. A more plausible response is that these deviant intuitive or emotional reactions cannot justify belief because they are unjustified themselves. Intuitions and emotions are unlike sensory perceptions because they are subject to normative assessment.[7] As such, they cannot play a foundational epistemic role in justifying beliefs without standing in need of justification.

This is not the place to prosecute the case against moral foundationalism in any more detail. Nevertheless, there are serious challenges to be overcome, and I am not optimistic about the prospects for success. Moral foundationalism is sometimes regarded as the only viable option when it comes to explaining our most basic moral knowledge. After all, the regress argument presents us with an unpalatable choice between moral foundationalism and moral skepticism: if moral knowledge is possible at all, then it must be possible without inference. By rejecting Hume's Law, however, we can block the regress argument for moral skepticism without collapsing into moral foundationalism. Instead, we can endorse a version of *moral inferentialism*, which says that inference is our most basic source of moral knowledge. On this anti-Humean view, the regress of moral knowledge comes to an end in moral conclusions known by inference from wholly non-moral premises.

---

[6] See Huemer (2005, Ch. 5) in defense of moral intuition and Döring (2007) or Tappolet (2016) in defense of moral emotion.

[7] Sosa (2007, 44-56) makes this point about intuition, while Brady (2013, 112-16) makes the corresponding point about emotion.

2. Theoretical Economy

A second motivation for the anti-Humean thesis is its theoretical economy. Moral inferentialism is *economical*, rather than *extravagant*, in the sense that it subsumes our most basic moral knowledge within the familiar category of inferential knowledge. Alex Byrne (2005) motivates an inferentialist theory of introspective knowledge in much the same way. Personally, I am more sanguine about the prospects for inferentialism in the case of moral knowledge than introspective knowledge. Nevertheless, theoretical economy is a strong though defeasible consideration in each case.

The anti-Humean thesis is economical because it promises a version of what Kieran Setiya calls *reductionism* about moral evidence, which "reduces evidence in ethics to evidence elsewhere" (2012, 49). On this view, moral knowledge doesn't require any special kind of moral evidence, since it merely requires deducing moral conclusions from non-moral premises. The background assumption here is that deductively valid arguments transmit evidential support from premises to conclusion (§9.4). If non-moral premises entail moral conclusions, then the evidential support for these non-moral premises transits across the entailment to the moral conclusion. Hence, the evidence for any moral conclusion can be reduced to the evidence for some non-moral premise that entails it.

In contrast, moral foundationalism is extravagant insofar as it appeals to distinctive sources of moral evidence, such as moral perception, moral intuition, or moral emotion. Quite apart from the problems that plague specific versions of moral foundationalism, they all impose costs in theoretical economy by inflating our sources of evidence. If we can, we should avoid positing distinctive sources of moral evidence. This counts to some degree against moral foundationalism and establishes moral inferentialism as the view to beat.

3. Causal and Epistemic Dependence

A third motivation for the anti-Humean thesis is that it explains how our moral knowledge about cases depends causally and epistemically on our non-moral knowledge about those cases.

In Harman's example, you know two things: a moral fact and a non-moral fact. Moreover, these two items of knowledge are not completely independent of each other. On the contrary, your moral knowledge depends counterfactually on your non-moral knowledge: if you didn't know the non-moral fact that the teenagers were torturing the cat, then you wouldn't know the moral fact that they were doing anything wrong. After all, we can assume that you have no other way of knowing the moral fact: for instance, no one tells you that the teenagers are doing anything wrong. In this example, you know the moral fact because, and only because, you know the non-moral fact.[8]

To illustrate the point, suppose a trusted expert gives you compelling but misleading testimony that the teenagers are not really torturing the cat, but merely acting in a play. Once you abandon your belief that they're torturing the cat, you lose your moral knowledge that they're doing anything wrong. Moreover, you cannot retain your moral knowledge by dogmatically maintaining your non-moral belief in the face of misleading counterevidence. To know the moral fact that the teenagers are doing something wrong, you must know the non-moral fact that they are torturing a cat. Again, you know the moral fact because, and only because, you know the non-moral fact.

Moral inferentialism explains why our moral knowledge about cases depends causally and epistemically on our non-moral knowledge about cases. It is because our moral knowledge depends inferentially on our non-moral knowledge. This is much harder for moral foundationalism to explain. If you can know without inference that the teenagers are doing something wrong, then why do you

---

[8] McGrath (2017) argues on similar grounds that our recognitional knowledge that something is an avocado depends inferentially on knowledge of how avocados look. As McGrath allows, however, we can know non-inferentially how avocado looks based on sensory perception alone.

need to know that they are torturing a cat? Why should the misleading testimony that the teenagers are acting in a play destroy your knowledge that they are doing something wrong? After all, there is no reason in principle why this misleading testimony should preclude you from having some moral perception, intuition, or emotion, which represents that the teenagers are doing something wrong. Given that the teenagers are in fact doing something wrong, why you can't you know this moral fact without inference on the basis of moral perception, intuition, or emotion?

I anticipate the reply that your moral knowledge is *defeated* by the misleading evidence about what the teenagers are doing. But this is simply to restate the intuitive datum to be explained without explaining it. Why should misleading evidence about the non-moral facts have any tendency to defeat your knowledge of the moral facts? Moral inferentialism explains why your moral knowledge is defeated when your non-moral knowledge is defeated: it is because your moral knowledge depends inferentially on your non-moral knowledge. Moral foundationalism cannot offer the same explanation of why your moral knowledge depends on your non-moral knowledge. Indeed, it's not clear that moral foundationalism can explain it at all. And yet there is surely a datum here that needs to be explained.

4. Moral Knowledge about Particular Cases

A fourth motivation for the anti-Humean thesis is that it explains how we can acquire moral knowledge about cases without relying on any antecedent knowledge of general moral principles. Sarah McGrath (2004) explains this by appealing to moral perception. She argues that moral inferentialism reverses the correct order of explanation, since it implies that we cannot acquire moral knowledge about particular cases unless we already know general moral principles. I'll argue in reply that moral inferentialism can avoid this objection by rejecting Hume's Law.

The anti-Humean thesis explains how you can know moral conclusions by deduction from non-moral premises without any antecedent knowledge of general moral principles. Moral perception

has no advantage over moral deduction in this respect. When you see some teenagers torturing a cat for fun, for example, you can know the moral conclusion that they are doing something wrong by deduction from the non-moral premise that they are torturing a cat for fun. You don't need to rely on knowledge of the general moral principle that it's always wrong to torture cats for fun.

The anti-Humean thesis diverges from *the modus ponens model* of moral deduction. According to the modus ponens model, you cannot know a conclusion by deduction from some premises unless the argument has a valid logical form. This means you cannot acquire moral knowledge by deduction from wholly non-moral premises, since there is no logically valid argument from non-moral premises to moral conclusions. To secure logical validity, we need to add a conditional moral premise linking the non-moral premises with the moral conclusion:

Moral Premise: It's always wrong for anyone to torture a cat just for fun.

Non-Moral Premise: Those teenagers are torturing a cat just for fun.

Moral Conclusion: Therefore, those teenagers are doing something wrong.

I'll argue in §9.2 that the modus ponens model is unmotivated, since not all deductive arguments are valid in virtue of logical form. I'll now explain why the modus ponens model of moral deduction should be abandoned and replaced with the anti-Humean model.

First, the modus ponens model makes it too hard to acquire moral knowledge about cases. It's much easier to know that some particular action is wrong than to know some general moral principle that explains why it is wrong. Imagine you witness the murder of an innocent person. You know it's not always wrong to kill, since it is permissible in self-defense and just war, but you that know this act of killing is wrong. Even so, you may be unable to articulate any general principle that explains why this act of killing is wrong, while others are permitted. Indeed, some philosophers argue that there are

no general moral principles that apply across all cases without exceptions. If there are any such exceptionless principles, they are very hard to articulate.[9]

Second, the modus ponens model cannot block the regress argument for moral skepticism without collapsing into moral foundationalism (cf. Zimmerman 2010, 126-7). On this view, you know a moral conclusion by deduction from a non-moral premise only if you know a conditional premise that links the non-moral premises with the moral conclusion. Since this conditional premise has moral content, we need to ask where this moral knowledge comes from. How do you know, for example, that it's always wrong to torture cats for fun? Perhaps you know this by deduction from a more general moral theory, such as act utilitarianism. At some point, however, the regress must come to an end in moral premises that are known without inference – say, on the basis of moral intuition. The modus ponens model implies that you cannot acquire moral knowledge by deduction unless you already have moral knowledge acquired non-inferentially in some other way. This is to abandon moral inferentialism in favor of moral foundationalism as a solution to the regress argument for moral skepticism. As such, it confronts the usual problems associated with explaining how we can acquire moral knowledge without inference.

Third, the modus ponens model limits the explanatory power of moral deduction. It cannot explain how we know general moral principles by means of deduction, since it presupposes that we already have knowledge of general moral principles, which figure as premises in our moral deductions. Once we know some general moral principles, we can acquire knowledge of new moral principles by deduction from the ones we already know. But deduction cannot explain how it is possible to acquire knowledge of general moral principles in the first place.

---

[9] See Dancy (2004) in defense of moral particularism and Jackson, Pettit, and Smith (2000) in reply. I revisit this issue in connection with the no-entailment problem (§8).

In contrast, the anti-Humean thesis avoids all of these problems. First, it explains how you can acquire moral knowledge about particular cases without any antecedent knowledge of general moral principles. Second, it blocks the regress argument for moral skepticism without collapsing into moral foundationalism. And finally, as we'll see in §5, the anti-Humean thesis explains how we can acquire a priori knowledge of general moral principles by deduction from non-moral premises that are assumed, rather than known on the basis of empirical evidence.

5. A Priori Knowledge of Moral Principles

A fifth motivation for the anti-Humean thesis is that it explains how we can acquire a priori knowledge of general moral principles by deducing moral conclusions from non-moral premises that are assumed, rather than known on empirical grounds. In §4, I argued that the anti-Humean thesis explains how we can acquire moral knowledge about particular cases without antecedent knowledge of general moral principles. In this section, I'll argue that the anti-Humean thesis explains how we can acquire a priori knowledge of general moral principles by exercising the same capacity for deductive reasoning that we use in acquiring moral knowledge about particular cases.

The key point is that we can acquire knowledge by *suppositional reasoning* – that is, by reasoning from assumptions as well as from known premises. Suppose you acquire deductive knowledge by following the introduction rule for disjunction: you know the disjunctive conclusion that *p* or *q* by deduction from the known premise that *p*. In much the same way, you can deduce the disjunctive conclusion that *p* or *q* from the assumption that *p*. Moreover, by conditional proof, you can discharge the assumption that *p* and deduce the conditional that if *p*, then *p* or *q*. This is a way to acquire deductive knowledge of the conditional by reasoning from assumptions, rather than known premises. Your knowledge is *certain*, rather than merely *probable*, since it is based on a deductive proof from no

undischarged assumptions. Moreover, your certain knowledge is justified *a priori*, since it doesn't rely on premises known through experience.

Similarly, the anti-Humean thesis explains how we can use suppositional reasoning to acquire a priori knowledge of general moral principles. Suppose you can know the moral conclusion that someone is acting wrongly by deduction from the known non-moral premise that they are torturing a cat just for fun. In much the same way, you can deduce the conclusion that someone is acting wrongly from the assumption that they are torturing a cat just for fun. By conditional proof, you can discharge the assumption and deduce the conditional that if someone tortures a cat just for fun, then they are doing something wrong. Once again, your knowledge of the conditional is certain, rather than merely probable, since it is based on a deductive proof from no undischarged assumptions. Moreover, your certain knowledge is justified a priori, since it doesn't rely on premises known through experience.

The anti-Humean thesis explains how we can acquire a priori knowledge of general moral principles with certainty and not just a high degree of probability. Some philosophers argue that we can use suppositional reasoning to acquire a priori knowledge of contingent facts about the reliability of perception and induction.[10] The suppositional reasoning in this case is inductively strong, but not deductively valid, so contingent facts about the reliability of perception or induction can be known with high probability, but not with certainty. In the moral case, by contrast, general moral principles can be known with certainty by deductive proof. Given the anti-Humean thesis, moral reasoning is deductive rather than inductive when moral conclusions are not just made probable by non-moral premises but entailed by them. As such, the anti-Humean thesis provides the basis for an attractive form of *moral rationalism*, according to which we can acquire certain knowledge of general moral principles on purely a priori grounds.

---

[10] See, for example, Hawthorne (2002), Cohen (2010), and Wedgwood (2013).

Any version of moral rationalism faces the challenge of explaining how it is possible to acquire a priori knowledge about morality. Can we explain the possibility of a priori knowledge without commitment to an obscure faculty of intuition or rational insight? Inferentialist theories of a priori knowledge promise to demystify the phenomenon by explaining it in terms of inferential capacities that we are committed to on independent grounds. The idea is that we can acquire a priori knowledge by making inferences from premises that are assumed and discharged, rather than known by empirical means. Inferentialist theories of a priori knowledge are *economical*, rather than *extravagant*, in the sense explained in §2: they make no appeal to distinctive sources of evidence, such as intuition or rational insight into the nature of necessity.

The anti-Humean thesis enables us to extend the inferentialist theory of a priori knowledge from logic to morality.[11] Just as we can acquire a priori knowledge of logical principles by using our capacity for logical deduction in suppositional reasoning, so we can acquire a priori knowledge of moral principles by using our capacity for moral deduction in suppositional reasoning. When we engage in suppositional reasoning about imaginary cases, we exercise the same inferential capacities that we use in reasoning about actual cases.[12] This demystifies the capacity for a priori knowledge by explaining it in terms of inferential capacities that we use in reasoning about real life scenarios. In this way, we can defend moral rationalism against the charge of obscurantism. The anti-Humean thesis answers the explanatory challenge by showing how we can acquire a priori knowledge of general moral principles by suppositional reasoning about imaginary cases.

---

[11] Wedgwood (2015) defends inferentialism about the a priori, and Wedgwood (2019) argues that fundamental moral principles can be known with a priori certainty, but he doesn't seek to explain our a priori knowledge about morality by extending his inferentialist theory to the moral domain.
[12] Williamson (2007) highlights this continuity between reasoning about actual and hypothetical cases, but he is more skeptical than I am about the epistemic significance of the distinction between a priori and a posteriori knowledge.

6. The Problem of Morally Repugnant Beliefs

A final motivation for the anti-Humean thesis is that it provides an attractive solution to *the problem of morally repugnant beliefs*: that is, it explains why we cannot have epistemic justification to hold morally repugnant beliefs, such as the belief that you ought to torture cats.

Some may be inclined to question the datum. After all, it's a familiar point that an evil demon can deceive you about the non-moral facts without making your beliefs unjustified. Indeed, this is the basis of Stewart Cohen's (1984) *new evil demon problem*. So perhaps an evil demon can deceive you about the moral facts too without making your beliefs unjustified. As I argue in more detail elsewhere, however, this position is unsustainable (see Smithies, forthcoming a).

The first premise is that if you have epistemic justification to believe that you should torture cats, then you should believe that you should torture cats. More generally, you should always believe whatever you have epistemic justification to believe. There can be no conflict between epistemic and moral demands on what you should believe. You can be morally faulted for your epistemic failings, but epistemic failings are never morally required. Morality cannot require you to refrain from believing what you know or have epistemic justification to believe.

The second premise is that if you should believe that you should torture cats, then you should torture cats. More generally, you should always do whatever you should believe you should do. Otherwise, you are sometimes permitted to be *akratic* – that is, to refrain from doing something that you should believe you should do. But akrasia is never permissible.

The third premise is that if you should torture cats, then you need no excuse to avoid culpability for torturing cats. Excuses are only needed to avoid culpability for wrongdoing, but you do nothing wrong when you act exactly as you should.

The fourth premise is that you are always culpable for torturing cats unless you have an excuse. It is hard to imagine a legitimate excuse for this kind of morally repugnant action, but perhaps we can

excuse those who cannot reasonably be expected to know any better – say, because they have been corrupted by their education within a sadistic society that has no respect for animal rights. In any case, an excuse is surely needed to avoid culpability, whether or not such an excuse is ever available.

These four premises imply that you can never have epistemic justification to believe that you should torture cats. Here is the argument in outline:

(1) If you have epistemic justification to believe that you should torture cats, then you should believe you should torture cats.

(2) If you should believe that you should torture cats, then you should torture cats.

(3) If you should torture cats, then you need no excuse to avoid culpability for torturing cats.

(4) But you always need an excuse to avoid culpability for torturing cats.

(5) Therefore, you never have epistemic justification to believe that you should torture cats.

Of course, the argument generalizes to other morally repugnant beliefs too. But it is one thing to argue for this conclusion and it is another thing to explain it. What explains why morally repugnant beliefs cannot be epistemically justified?

Standard versions of moral foundationalism cannot explain this. Suppose, for example, that we can acquire non-inferential moral knowledge on the basis of experiences that accurately represent the moral facts (such as moral perceptions, intuitions, or emotions). If so, then presumably we can experience moral illusions that misrepresent the moral facts too. For example, a moral deviant who sees the teenagers torturing a cat might have an illusory experience that represents that they are doing something morally required. In the absence of defeaters, such illusory experiences might give them

epistemic justification to believe that they should torture cats. And so moral foundationalism seems committed to the problematic result that morally repugnant beliefs can be epistemically justified.[13]

The anti-Humean thesis explains why you cannot have epistemic justification for morally repugnant beliefs. If you know that someone is torturing cats just for fun, then you're in an epistemic position to know by deduction that they're acting immorally. Moreover, you're always in an epistemic position to know by suppositional reasoning that anyone who tortures cats just for fun is thereby acting immorally. But if you're always in an epistemic position to know that p, then you can never have epistemic justification to believe that not-p.[14] Hence, you can never have epistemic justification to believe that it's morally permissible to torture cats just for fun.

The problem of morally repugnant beliefs raises a more general question about the connection between epistemic justification and truth. Some philosophers solve the problem by invoking a *global truth-connection* in epistemology, according to which a belief is epistemically justified only if it is true.[15] And yet this seems like overkill, since it has the implausible consequence that an evil demon cannot deceive you about the external world without thereby making your beliefs unjustified. It solves the problem of morally repugnant beliefs at the cost of reinstating the new evil demon problem.

In contrast, the anti-Humean thesis solves the problem of morally repugnant beliefs by invoking a *local truth-connection* in moral epistemology. The key point is that we're always in a position to know the necessary truths of morality, but we're not always in a position to know contingent moral or non-moral truths. This is why we can have justified false beliefs about contingent moral or non-

---

[13] Littlejohn (2011) and Williamson (2019) give similar arguments against Huemer's (2005) moral intuitionism, while Kelly and McGrath (2010) make the same complaint against moral coherentism.
[14] This principle follows from two further assumptions: (i) you're in an epistemic position to know that p only if you have epistemic justification to believe that p; and (ii) you can never have epistemic justification to believe contradictory propositions.
[15] See, for example, Littlejohn (2012) and Williamson (2013).

moral truths, but not about necessary truths of morality. Hence, the anti-Humean thesis solves the problem of morally repugnant beliefs without reinstating the new evil demon problem.

7. The Over-Intellectualization Problem

Having made a case for the anti-Humean thesis, let me now defend it against some objections. First, why suppose that our moral knowledge is based on inference at all? When we see the teenagers torturing the cat, for example, why do we need to *infer* that they are doing something wrong? Gilbert Harman articulates the challenge as follows:

> If you round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it, you do not need to *conclude* that what they are doing is wrong; you do not need to figure anything out; you can *see* that it is wrong. (Harman, 1977, p. 4)

The suggestion here is that you don't need to make an inference because it's obvious that what the teenagers are doing is wrong. As Harman puts the point, you can just *see* that it is wrong. The objection is that moral inferentialism over-intellectualizes the psychological process involved in acquiring moral knowledge about cases. My response is that the objection over-intellectualizes the psychological process of inference. I'll make three points in turn.

First, semantic considerations have no force against moral inferentialism. When you see the teenagers torturing the cat, it would be perfectly acceptable to say: "I see that those teenagers are doing something wrong." And yet this doesn't mean that your knowledge is not based on inference. There is no semantic entailment from the premise (i) that you see that *p* to the conclusion (ii) that you know that *p* non-inferentially on the basis of visual perception. To illustrate the point, consider a paradigm example of inferential knowledge in which you see your neighbor's car parked in her driveway and

infer that she is back home. In this case, it would be perfectly natural to say, "I see that my neighbor is home." And yet this doesn't imply that your knowledge is not based on inference. On the contrary, your knowledge is based on abductive inference: the hypothesis that your neighbor is home provides the best explanation of what you know non-inferentially on the basis of perception: namely, that her car is parked in the driveway.[16]

Second, introspective considerations have no force against moral inferentialism. You may have no awareness of making an inference to the conclusion that the teenagers are doing something wrong. But the same is true of much of our inferential knowledge. When you see the car parked in the driveway, for example, you needn't be aware of making an inference to the conclusion that your neighbor is home. And yet it surely doesn't follow that you're not making an inference at all. Rather, you're not aware of making the inference because it is unconscious and automatic. You make the inference, we might say, without thinking. Similarly, in Harman's example, you can know without thinking that the teenagers are doing something wrong. You don't need to "figure out" this conclusion, since it is obviously true, but that doesn't mean your moral knowledge isn't based on inference at all.

When you're not aware of making an inference, what is the difference between unconscious inference and no inference at all? My third point is that your knowledge of a conclusion is inferential when it depends causally and epistemically on your knowledge of the premise in the right kind of way. For example, your knowledge that your neighbor is home is inferential because it depends causally and epistemically on your knowledge that her car is in the driveway. Consider the following counterfactual test: if you didn't know that her car was parked in the driveway, then you would have no way of knowing that she was home. As we're imagining the example, this is your only way of knowing that she is home: for example, your neighbor didn't share her travel plans with you. You

---

[16] Compare McGrath (2004, 221-2) and Zimmerman (2010, 85) for similar points.

know that your neighbor is home because, and only because, you know that her car is in the driveway. That is why your knowledge is based on inference.

Similar considerations apply in Harman's example. Why suppose that your moral knowledge is inferentially based on your non-moral knowledge? My answer, as before, is that your moral knowledge depends causally and epistemically on your non-moral knowledge in the right kind of way. Once again, you satisfy the counterfactual test: if you didn't know that the teenagers were torturing a cat, then you wouldn't know that they were doing anything wrong. As we're imagining the example, this is your only way of knowing they are doing something wrong: for example, no one tells you this. You know that the teenagers are acting wrongly because, and only because, you know that they are torturing a cat for fun. That is why we should conclude that your moral knowledge is inferentially based on your non-moral knowledge.

In summary, moral inferentialism is undermined by neither semantic nor introspective considerations. On the contrary, it is supported by the consideration that our moral knowledge about cases depends causally and epistemically on our non-moral knowledge about those cases. As we saw in §3, this is much harder for opposing views to explain.

8. The No-Entailment Problem

Assuming your moral knowledge is based on inference, why suppose it is based on deductive inference, rather than probabilistic inference? This section addresses the objection that the inference cannot be deductive, since there is *no entailment* from our non-moral knowledge about particular cases to our moral knowledge about those cases.[17]

---

[17] McGrath (2004, 218-20) raises this objection to the modus ponens model of moral deduction, while Setiya (2012, 44-6 & 52 n. 1) extends it to Zimmerman's (2010, Ch. 5) anti-Humean thesis.

According to the anti-Humean thesis, moral knowledge can be acquired by means of deductions that satisfy the following three conditions:

(1) The premise is wholly non-moral.

(2) The premise is known to be true.

(3) The premise entails the moral conclusion.

The objection is that these conditions cannot be jointly satisfied: if the premise is wholly non-moral, and known to be true, then it doesn't entail the moral conclusion. This is because there are no general moral principles – or very few – that hold across cases without exceptions. To secure the entailment, we need to build these exceptions into our specification of the premises. If we specify the exceptions in moral terms, then the premise entails the conclusion, and we can know the premise to be true, but the premise is not wholly non-moral. But if we specify the exceptions in non-moral terms, then the premise entails the conclusion, and the premise is wholly non-moral, but we cannot know the premise to be true. Either way, the deduction fails to satisfy all three conditions.

Arguably, there are some general moral principles that hold across cases without exceptions. Consider the principle that it's always wrong to torture cats just for fun. Perhaps it's morally permissible to torture a cat to assuage an evil demon who threatens otherwise to torture many cats. In that case, however, I'm not torturing a cat just for fun, since I'm doing it to save other cats.

Is there any possible scenario in which it's morally permissible to torture a cat just for fun? Suppose an evil demon threatens to torture many cats unless I torture a single cat and, moreover, I do it just for fun. Some might say this gives me a moral reason to torture a cat just for fun – namely, because it will save other cats. But this is the wrong kind of reason for me to act on. If I act for this reason, then I'm not torturing the cat just for fun, since I'm doing it to save other cats. Therefore, I

cannot act for this reason at all. Arguably, it is not a reason to torture cats just for fun, but a reason to get myself into a state in which I torture cats just for fun. Perhaps this reason makes it morally permissible to get myself into a state in which I do something morally impermissible.[18] And yet this is no counterexample to the principle that it's never morally permissible to torture cats just for fun.

I don't rest much weight on this response, however, since it doesn't generalize far enough. Suppose that almost all general moral principles are subject to exceptions aside from a handful of rare exceptions. This would significantly limit the scope of our capacity to acquire moral knowledge through deduction. Any robust defense of the anti-Humean thesis needs to explain how we can acquire moral knowledge by deduction in a much wider range of cases. Suppose, for example, that I know that my friend has wronged me by breaking a promise. Can I acquire deductive moral knowledge by means of the argument below?

My friend has broken a promise.

Therefore, he has done something wrong.

The problem is that the premise doesn't entail the conclusion, since it's not always wrong to break a promise. Everyone (except perhaps Kant) can agree that it's morally permissible to break a promise for good enough reasons, such as rescuing a drowning toddler. The challenge for the anti-Humean thesis is to explain how we can acquire moral knowledge by deduction in cases like this one.

One response is to modify the principle by building in the exceptions in explicitly moral terms: for example, perhaps it's always wrong to break a promise unless you have a good enough reason to do so. So we might reconstruct the argument as follows:

_____

[18] Compare Parfit (1984, 12-13) on rational irrationality: just as it's sometimes rational to induce yourself to act irrationality, so it's sometimes morally permissible to induce yourself to act immorally.

My friend has broken a promise, but he has no good reason to do so.

Therefore, my friend has done something wrong.

The problem is that even if the premise entails the conclusion, and we know the premise to be true, the premise now contains moral content. After all, a good enough reason for breaking a promise is just a condition that makes it morally permissible to do so.

Another response is to modify the principle by listing all the exceptions in non-moral terms: for example, it's always wrong to break a promise unless you do so in order to rescue a drowning toddler, or to assuage an evil demon, and so on. Now we can reconstruct the argument as follows:

My friend has broken a promise, but not in order to rescue a drowning toddler, or to assuage an evil demon, or . . .

Therefore, my friend has done something wrong.

Although the premise entails the conclusion, and the premise is wholly non-moral, the problem is to explain how the premise can be known, since there is no finite way to state all the possible exceptions in non-moral terms. After all, there are infinitely many ways in which you might permissibly break a promise and there is no finite way to specify in non-moral terms what they all have in common. If the non-moral premise cannot be finitely stated, then how can it be known by any finite human mind?

In reply, this is a general feature of our epistemic situation, rather than anything unique to the moral case. As I look around my office, for example, there are infinitely many things I know about

the contents of the room: there are no elephants, no gorillas, no zebras, and so on.[19] And it is not just that there are no zoo animals: there are no tightrope walkers, no ghosts, no waterfalls, and no black holes. I know all these facts about my office, although I cannot list them in any finite way. There is no finite way of stating everything I know about my office right now.

Much the same point is familiar from the literature on abduction. Suppose I know that it rained last night by abduction from the observation that the streets are wet. My abductive knowledge relies on vast quantities of background knowledge: for example, I know it's vastly improbable that the streets are wet because they were sprayed with a hose or a water cannon. Again, there may be no finite way of stating everything I know that is relevant to the justification of my abductive inference. As Harman writes, "It is doubtful that anyone has ever fully specified an actual piece of inductive reasoning, since it is unlikely that anyone could specify the relevant total evidence in any actual case" (1973, 28-9).

The same point applies in moral cases too. Suppose I know, as I surely can, that my friend has wronged me by breaking a promise. To know this, I must know on independent grounds that he isn't rescuing drowning toddlers, he isn't assuaging an evil demon, and so on. To know that my friend has wronged me, I must know that none of these exempting conditions obtain, although there is no finite way to list them all. How do I know these facts? In many different ways: perception, testimony, abduction, and so on. My total evidence makes it vastly improbable that any of these exempting conditions obtains. In any case, this is a question in general epistemology, rather than moral epistemology, so we can set it aside here.[20]

---

[19] Wittgenstein famously denied these claims in conversation with Russell, but he was rightly mocked: should we go ahead and check under the table?

[20] Likewise, we can set aside concerns about multi-premise closure arising from the preface paradox: how can I know an infinite conjunction of non-moral facts when there is some non-zero evidential probability that each conjunct is false? One tempting solution is that knowledge requires evidential probability 1, although perhaps in a contextually restricted probability space (Lewis, 1996).

Sometimes, of course, my beliefs are mistaken. I might reasonably believe that my friend has no good reason for breaking his promise when some unforeseen emergency has arisen. Even when my belief is true, it is *insensitive* because I would still believe that my friend has no good reason in the unlikely event that he actually has one. But this is not enough to impugn my knowledge. I can know that my friend has no good reason so long as my belief is not only justified, and true, but also *safe from error* in the sense that it couldn't easily have been false. This safety condition is satisfied when these unforeseen circumstances are not only statistically unlikely, but also modally remote.[21]

On a probabilistic version of moral inferentialism, my moral knowledge need not be entailed by my non-moral knowledge so long as it is made sufficiently probable.[22] On this view, I can know that my friend has wronged me by breaking a promise without knowing that he has no good reason to do so. But this is extremely puzzling. How can I know that my friend has wronged me if I don't know, for instance, that he isn't rescuing toddlers? Surely, I am not oblivious to the fact that it is wrong for him to break the promise only if he isn't rescuing toddlers. But if I know this conditional, and I know the antecedent, then I must know (or be in an epistemic position to know) the consequent. Otherwise, we violate a plausible closure principle for knowledge. Moreover, I cannot know the non-moral fact that my friend isn't rescuing toddlers by deduction from the moral premise that he has wronged me, since this reverses the direction of epistemic dependence. Rather, I must know the non-moral fact antecedently in order to know the moral fact in the first place.

The residual problem is to explain how finite creatures like us can know infinitely many facts. The key point is that much of our knowledge is dispositional. We know and believe many things we are not currently thinking about and have never previously considered. It's enough that we're disposed

---

[21] Sosa (1999) and Williamson (2000, Ch. 7) argue that knowledge doesn't require sensitivity, but only safety from error.

[22] Zimmerman (2010, 123) considers this view as a fallback option, while Setiya (2012, 52 n. 1) endorses it as an alternative to the deductive view.

to judge them true when we consider them. Not all knowledge is acquired through the conscious mental act of judgment. Consider, for example, the proposition that you lust after Donald Trump. Presumably, you learn nothing new when you consider this for the first time. You already knew it was false, although it was too obvious to be worth considering. This is dispositional knowledge, rather than knowledge you are merely disposed to acquire. Finite creatures can have infinitely many dispositions even if we cannot exercise them all at once. This is how finite creatures can have dispositional knowledge of infinitely many facts.[23]

In summary, the no-entailment objection fails because we know more than we can state. In particular, we can know moral facts by deduction from the non-moral facts on which they supervene even if there is no finite way of stating these non-moral facts in wholly non-moral terms. This is a general feature of our epistemic situation, rather than anything unique to the moral case.

9. The Wrong-Entailment Problem

Assuming our moral knowledge is entailed by our non-moral knowledge, is it the right kind of entailment to explain how we can acquire moral knowledge by deduction? The problem is that there is no logical or analytic entailment from non-moral facts to moral facts, but metaphysical entailment is the wrong kind to explain deductive knowledge. I argue that this is a problem for everyone and that we need an epistemic concept of entailment to solve the problem. Hence, the anti-Humean thesis should be understood as the thesis that all moral facts are *evidentially entailed* by non-moral facts on which they supervene.

---

[23] On a dispositional theory, believing that *p* doesn't require storing the information that *p* in some explicit format. Indeed, this true even on Fodor's (1987) language of thought hypothesis, since you can implicitly believe that *p* without storing a mentalese sentence that means that *p*.

9.1. Metaphysical Entailment

Moral supervenience implies that moral facts are metaphysically entailed by non-moral facts on which they supervene: for any moral fact M, there is some non-moral fact N such that it's metaphysically necessary that if N, then M.[24] In other words, there is a metaphysical entailment from the non-moral premise that N to the moral conclusion that M, since it is metaphysically impossible for the premise to be true when the conclusion is false. Nevertheless, this is the wrong kind of entailment to explain how we can know the moral conclusion by deduction.

Sarah McGrath (2004, 217) makes the point with the following example:

x contains water. Therefore, *x* contains oxygen.

There is a metaphysical entailment from premise to conclusion, since it is metaphysically necessary that water molecules are composed of oxygen atoms: any superficially similar liquid with a different chemical composition is not water, but fool's water.[25] And yet this is the wrong kind of entailment to explain deductive knowledge. An ancient scientist who knows that his glass contains water cannot know by deduction that the glass contains oxygen, since he doesn't know the metaphysically necessary truth that water always contains oxygen.

We shouldn't conclude from this example that you can know that Q by deduction from the premise that P only if you know the further premise that P entails Q. This principle is viciously regressive, since it implies that you know Q by deduction only if you know not only (i) P, and (ii) P entails Q, but also (iii) P, and P entails Q, entails Q; and so on ad infinitum. Lewis Carroll (1985) illustrates this regress problem in his classic tale, "The Tortoise and Achilles".

---

[24] See Jackson (1998, 24-6) for the argument.
[25] See Putnam (1975) and Kripke (1980) on natural kind essentialism.

Paul Boghossian (2003) argues on similar grounds that some deductions are *blind* in the sense that they yield knowledge without antecedent knowledge of their own validity. But which deductions are blind in this sense? And is there any plausible answer that includes anti-Humean deductions from non-moral premises to moral conclusions? Metaphysical entailment is the wrong kind of entailment to explain how deductive knowledge is possible. Therefore, we must look elsewhere.

9.2. Logical Entailment

The obvious proposal is that you can know a conclusion by deduction from a known premise only if there is a *logical entailment* from premise to conclusion. Of course, this proposal spells doom for the anti-Humean thesis. After all, the grain of truth in Hume's Law is that there is no logical entailment from non-moral premises to moral conclusions.[26] For example, the argument from the premise that the teenagers are torturing a cat just for fun to the conclusion that they are doing something wrong has the logically invalid form: P, therefore Q. To make the argument logically valid, we need to add a conditional premise linking the non-moral premise with the moral conclusion. And this is to abandon the anti-Humean thesis in favor of the modus ponens model, which I criticized in §4.

Nevertheless, this proposal is much too restrictive. Not all deductively valid arguments are valid in virtue of their logical form. Consider the following examples of materially valid deduction:

x is red. Therefore, *x* is colored.

x is a bachelor. Therefore, *x* is unmarried.

x knows that *p*. Therefore, *x* believes that *p*.

---

[26] Exceptional cases aside; for example: "Tea-drinking is common in England; therefore either tea-drinking is common in England or all New Zealanders ought to be shot" (Prior 1960, 205).

We can deduce the conclusion of each argument from its corresponding premise without relying on antecedent knowledge of a conditional premise linking the premise with the conclusion. To suppose otherwise is to privilege logical expressions over non-logical expressions in a theory of deductive knowledge without any principled epistemological motivation. Moreover, it makes it hard to explain how we can know the conditional premise in the first place. How can we know, for example, that anything red is colored? Arguably, we know this by suppositional reasoning: given the assumption that an arbitrary object is red, we can deduce that it is colored, discharge the assumption, and conclude by conditional proof that anything red is colored. However, we cannot use deduction to acquire knowledge of the conditional premise if this knowledge is presupposed in the deduction itself. So much the worse for the modus ponens model.

In summary, we need an account of deductive entailment that is broad enough to include materially valid as well as logically valid deductions. But is there any useful notion of deductive entailment that is broader than logical entailment but also narrower than metaphysical entailment?

9.3. Analytic Entailment

Paul Boghossian (2003) argues that a deduction is blind just in case there is an *analytic entailment* of the right kind from premises to conclusion. This extends the notion of analyticity from truth to entailment: if it is an analytic truth that anything red is colored, then there is an analytic entailment from the premise that x is red to the conclusion that x is colored. On Boghossian's epistemic conception of analyticity, the meaning of a sentence doesn't explain why it's true, but why you are justified in believing it. Similarly, the meaning of an argument doesn't explain why it is valid, but why it transmits justification from premises to conclusion. An entailment is analytic just in case it is a precondition for

understanding the argument that you are disposed to deduce the conclusion from the premises. Moreover, this is what explains your justification to perform the deduction.[27]

This proposal includes logical entailments within the class of analytic entailments, while excluding some metaphysical entailments. For example, Boghossian claims that *modus ponens* is analytic because understanding the concept of the conditional requires being disposed to make inferences in virtue of modus ponens form. In contrast, there is no analytic entailment from premises about water to conclusions about oxygen, since understanding the concept of water requires no disposition to make those inferences. Hence, the appeal to analytic entailment promises extensional adequacy as well as explanatory depth. Unfortunately, however, I'll argue that this promise cannot be sustained.

Aaron Zimmerman (2010, Ch. 5) defends the anti-Humean thesis by extending Boghossian's proposal to the moral domain. He argues that there are analytic entailments from non-moral premises to moral conclusions because understanding moral concepts requires some disposition to deduce moral conclusions from non-moral premises. He writes, "It is arguable that we must be disposed to reason from 'is' to 'ought' – or at least sometimes reason in this way – if we are to truly grasp moral concepts" (2010, 139).

The appeal to analytic entailment can be questioned on multiple fronts. First, it is disputable that there are any analytic entailments at all. Timothy Williamson (2007, 92-4) argues that not even modus ponens is analytic, since someone can understand the conditional without being disposed to infer by modus ponens. For example, Vann McGee (1985) understands the conditional well enough to publish papers on the topic, but his papers contain intuitive counterexamples to modus ponens that draw on his intuitive understanding of the conditional. Presumably, he is not disposed to make inferences in virtue of modus ponens form, rather than some more restricted form, since he denies

---

[27] Boghossian (2003, 239) remains neutral about whether entailment is necessary for transmission of justification by deduction, but I assume it is necessary for transmission of deductive knowledge.

that modus ponens is valid. Boghossian (2020, 190) replies that McGee's inferential dispositions are masked by his theoretical doubts. But, if so, we can imagine his fictional counterpart, "SchmcGee", whose inferential dispositions to use the restricted form are innate, rather than based on theoretical reflection. Given the similarity between their published views, it seems uncomfortable to deny that McGee and SchmcGee have any shared understanding of the conditional.

Second, Moore's (1903) open question argument provides special reasons to doubt that there any analytic entailments from non-moral premises to moral conclusions. Moore argues that we cannot define moral concepts in non-moral terms, since it's always an open question, conceptually speaking, whether an action described in non-moral terms is good or bad, right or wrong. According to Moore, such truths are never analytic, but always synthetic. For similar reasons, inferences from non-moral premises to moral conclusions are never analytic, but always synthetic.

Boghossian (2020, 196-9) bolsters Moore's argument by appealing to the normative character of moral concepts. When you judge that it's morally wrong to perform an action, you thereby judge that it's something you *shouldn't* do. It's always coherent for someone who knows the non-moral facts about an action to regard it as an open question whether they should perform the action. Moreover, we cannot close the question by stipulating that the action is wrong by definition, since we can coherently regard it as an open question whether we should perform actions that are wrong according to this definition. But this only goes to show that the proposed definition fails to capture our concept of moral wrongness.

Moore's open question argument threatens to undermine Zimmerman's proposal that there are analytic entailments from non-moral premises to moral conclusions. Someone who knows the non-moral premise that the teenagers are torturing a cat just for fun can coherently reject the conclusion that they are doing something wrong – or regard it as an open question – without thereby compromising their understanding of moral concepts. For example, a *moral nihilist* like John Mackie

(1977) can maintain his grasp of moral concepts while denying that the teenagers are doing anything wrong. As Zimmerman (2010, 138) surmises, Mackie may have been a "fine moral judge" who routinely made such inferences when he left his study. But we can imagine his fictional counterpart, "Schmackie", whose inferential dispositions cohere more systematically with his theoretical beliefs. Given the similarity between their published views, it seems uncomfortable to deny that Mackie and Schmackie have any shared understanding of moral concepts.

Similarly, a *moral deviant* can understand moral concepts while having systematically false moral beliefs and correspondingly unreliable dispositions to infer moral conclusions from non-moral premises. Our shared understanding of moral concepts provides plenty of latitude for substantive disagreement about the moral facts. We can disagree, for example, with a sadist who denies that there is anything wrong with torturing cats for fun. And there are no principled limits on how far our moral disagreements may extend. We can disagree extensively about the substance of morality without simply talking past each other. Our moral deviant may be substantively mistaken – perhaps dangerously so – but he need not be conceptually confused.

Third, there are more general reasons to doubt the prospects for an analytic explanation of the justification of deduction. As Boghossian observes, there are intuitive counterexamples to the following normative principle:

> *The Meaning-Entitlement Connection (MEC)*: Any inferential transitions built into the possession conditions for the concept are *eo ipso* entitling. (2003, 241)

For example, Boghossian imagines an ancient scientist who believes without evidence that water is composed of a substance called 'aqua' that is itself composed of H_2O molecules. His concept, *aqua*, is governed by the following introduction and elimination rules:

x is water. Therefore, *x* is composed of aqua.

x is composed of aqua. Therefore, *x* is composed of H_2O molecules.

MEC implies that the scientist is blindly justified in following these inference rules, since they are built into his concept of aqua. As we saw in §9.1, however, the scientist is not justified in deducing the conclusion that his glass contains H_2O molecules from the premise that it contains water, since he doesn't know that water contains H_2O molecules. The epistemic situation doesn't change just because he builds these inference rules into his understanding of the concept, *aqua.*

Boghossian argues that we need to restrict MEC to *non-defective* concepts in order to avoid counterexamples of this kind. He argues that a concept is non-defective just when either (i) the concept is conditional in form, or else (ii) no conditional form of the concept is available. The concept, *aqua*, is defective because it is unconditional in form: possessing the concept requires being disposed to make inferences that are truth-preserving only if aqua exists, which is a theoretical hypothesis that needs evidential support.[28] In contrast, there is nothing defective about the unconditional form of *the conditional*, since no conditional form of the concept is available without it. We cannot operate without an unconditional form of the conditional, since we need it in order to conditionalize on the existence of anything in the first place.

Boghossian's restriction on MEC threatens to undermine Zimmerman's proposal, since it implies that our moral concepts are defective in their unconditional form. There is no principled reason why we cannot operate with conditional versions of our moral concepts that are shared

---

[28] We cannot explain why aqua is defective by appeal to the thesis that the analytic inferences of any non-defective concept are necessarily truth-preserving (Peacocke, 1992, 21). After all, the analytic inferences rules for the concept, *aqua*, are necessarily truth-preserving so long as identity is a determinate form of composition and aqua is identical to water.

between moral realists and moral nihilists alike. When we see the teenagers torturing a cat, for example, we are not blindly justified in deducing the unconditional conclusion that they are doing something wrong, but only the conditional conclusion that that if there are moral facts, then they are doing something wrong. Moral realists and moral nihilists can agree on this conditional, while disagreeing about whether to use it in reasoning by modus ponens or modus tollens. This means we cannot defend the anti-Humean thesis in the way that Zimmerman proposes by extending Boghossian's account from logic to morality.

Rather than undermining the anti-Humean thesis, however, this only casts doubt on the prospects for an analytic explanation of justification. For example, Boghossian's restricted version of MEC has skeptical implications when we extend it from logical concepts to empirical concepts. Suppose, for example, that understanding the observational concept, *red*, requires being disposed to make inferences (or non-inferential transitions) of the following form:

x looks red in normal conditions. Therefore, x is red.[29]

Can we use this fact to give an analytic explanation of our defeasible justification to believe that things are red when they look red? Boghossian's version of MEC says no, since this unconditional version of the observational concept is defective. After all, there is no principled reason why we cannot operate instead with conditional versions of our observational concepts that can be shared between color realists and color nihilists alike. On this view, we are not blindly justified in believing that things are red when they look red, but only that they are red if anything is colored at all. For similar reasons, we

---

[29] Peacocke (1992, 7-8) proposes similar possession conditions for observational concepts. The threat of circularity can be avoided by treating the inputs to these transitions not as the belief that something looks red, but as a visual experience in which an object is non-conceptually represented as red.

should operate with conditional versions of all our empirical concepts. This means that we are not blindly justified in using memory to form beliefs about the past, induction to form beliefs about the future, or abduction to form beliefs about the unobserved. And so the analytic explanation of justification generates a more general form of skepticism.

In conclusion, there are general reasons to be suspicious of the program of giving an analytic explanation of our justification to use basic belief-forming methods, such as deduction, induction, and perception.[30] This reductive project is one of the legacies of logical positivism, which sought to reduce questions in metaphysics and epistemology to questions in the theory of meaning. In my view, this reductive project should be abandoned. A different approach is needed.

9.4. Evidential Entailment

The wrong-entailment problem is a problem for everyone. The problem is to explain what kind of entailment makes deduction possible. To do this, we need a concept of entailment that is intermediate between logical and metaphysical entailment. Analytic entailment cannot do the requisite explanatory work, since epistemic notions cannot be reductively explained in purely semantic terms. Instead, I'll now argue, we need to recognize a distinctively epistemic concept of *evidential entailment*. With this concept in hand, we can defend the anti-Humean thesis against the wrong-entailment objection.

Evidential entailment is one member of a family of concepts of entailment defined in modal terms. On the modal definition of entailment, a conclusion is *entailed* by some premises just in case it is *impossible* for the premises to be true when the conclusion is false. But there are many different ways in which the modal concepts of possibility and necessity can be understood. My contention is that we

---

[30] Compare Enoch and Schechter (2006) for general doubts about the prospects for an analytic explanation of our justification to use basic belief-forming methods. And see Williamson (2006) for more specific doubts about the analytic explanation of our justification for logical deduction.

cannot put the concept of entailment to work in explaining deductive knowledge unless it is understood in terms of an epistemic concept of modality. There should be no default assumption that we can reduce epistemic modality to logical, analytic, or metaphysical modality. Nevertheless, I'll suggest that this epistemic concept of modality can be explained in broadly evidentialist terms.[31]

We can begin with the idea that your total evidence supports various propositions to varying degrees. In a probabilistic framework, your degree of evidential support for a proposition is modeled by its evidential probability: that is, its probability conditional upon your total evidence.[32] The maximal degree of evidential support is *evidential certainty*: that is, evidential probability 1. With the concept of evidential certainty in hand, we can say that it is *evidentially necessary* that p if and only if, for any evidence e, it is evidentially certain that p given e. Moreover, we can say that p *evidentially entails* that q if and only if, for any evidence e, it is evidentially necessary that if p, then q. In this way, the concepts of certainty, necessity, and entailment can be inter-defined.

If evidential probabilities are defined over evidential possibilities in the usual way, then we can add the further principle below:

> *The Evidential Entailment Principle*: Necessarily, if p evidentially entails that q, then for any evidence e, the evidential probability that q given e is greater than or equal to the evidential probability that p given e.

---

[31] Compare Chalmers (2011) for a similar conception of epistemic modality. One theoretical difference is that Chalmers epistemic necessity in terms of a priori justification, whereas I define it in terms of evidential certainty. Otherwise, my account is very similar.

[32] See Williamson (2000: Ch. 10) for this notion of evidential probability, which he argues cannot be reduced to objective probability (chance) or subjective probability (credence).

More generally, whenever there is an evidential entailment from the premises of an argument to its conclusion, your degree of evidential support for the conclusion can be no less than your degree of evidential support for the conjunction of the premises. Thus, evidential entailment is a relation between premises and conclusion that preserves not only *truth*, but also *degrees of evidential support*.

The concept of evidential entailment captures an important epistemic distinction between deductive validity and inductive strength. In a deductively valid argument, the evidential entailment from premises to conclusion guarantees that the evidential probability of the conclusion can be no less than the evidential probability of the conjunction of the premises. In an argument that is merely inductively strong, in contrast, the evidential probability of the conclusion may be less – though not too much less – than the evidential probability of the premises. How much less depends on its degree of inductive strength, but only a deductively valid argument guarantees that the conclusion is no less probable than the conjunction of the premises.

At the same time, there is an important epistemic continuity between deductive validity and inductive strength. Arguments can be ordered on a continuum of inductive strength corresponding to the degree to which evidential support is transmitted from premises to conclusion. Brian Skyrms (2000, 20) proposes that the inductive strength of an argument can be modeled by its *inductive probability* – that is, the conditional probability of the conclusion given the premises. In a deductively valid argument, of course, the conditional probability of the conclusion given the premises is 1.[33] Hence, deductive validity is the upper limiting case of inductive strength in which the evidential entailment from premises to conclusion guarantees that there is no loss of evidential support from premises to

---

[33] This can be proved from the standard ratio formula for conditional probability, which says that $\Pr(q \mid p) = \Pr(p \& q) / P(p)$ where $\Pr(p) > 0$. If p entails q, then $\Pr(p) = \Pr(p \& q)$, since these propositions are logically equivalent. So, when $\Pr(p) > 0$, $\Pr(q \mid p) = 1$.

conclusion. As Skyrms writes, in a deductively valid argument, "the evidential link between premises and conclusion is as strong as possible" (2000, 15).

I contend that we need this concept of evidential entailment to give a satisfying response to the wrong-entailment problem. Evidential entailment, rather than logical or metaphysical entailment, is what explains the possibility of deduction. Quite plausibly, all logical entailments are evidential entailments, although not vice versa. The evidential entailments also include some material entailments that are either analytic or (as I think) synthetic, but knowable a priori. For example, there is an evidential entailment from the premise that x is red to the conclusion that x is colored, since your evidence cannot make it probable that something is red without thereby making it equally probable that it is colored. At the same time, not all metaphysical entailments are evidential entailments. For example, there is no evidential entailment from the premise that x contains water to the conclusion that x contains oxygen, since your evidence can make it probable that there is water in your glass without thereby making it equally probable that there is oxygen in your glass. This is exactly the predicament of an ancient scientist who has no evidence that water contains oxygen.

With the concept of evidential entailment in hand, we can now defend the anti-Humean thesis against the wrong-entailment objection. The anti-Humean thesis can be understood as the thesis that there are evidential entailments from non-moral facts to moral facts. More specifically, it says that all moral facts are evidentially entailed by non-moral facts on which they supervene. On this view, your evidence cannot make it probable that some teenagers are torturing a cat for fun without thereby making it equally probable that they are doing something wrong. Any evidence that someone is torturing a cat for fun is evidence that they are doing something wrong.

My goal in this section is not to argue for the anti-Humean thesis, but merely to defend it against the objection that the entailment from non-moral premises to moral conclusions is the wrong kind to explain deductive knowledge. My reply is that everyone needs the concept of evidential

entailment to solve the wrong-entailment problem and that proponents of the anti-Humean thesis can solve the problem in exactly the same way. This is not supposed to be an argument that there is an evidential entailment from non-moral premises to moral conclusions, but I gave my arguments for the anti-Humean thesis in §§1-6.

I cannot hope to defend the anti-Humean thesis by articulating a general criterion for evidential entailment and arguing that moral deductions satisfy this criterion. But there is no reasonable precedent for demanding such a criterion. Consider, by analogy, a Humean skeptic about induction who claims that all scientific knowledge must be based on deduction. We cannot assuage the inductive skeptic by appealing to some general criterion for justified inference that applies equally to deduction and induction. Of course, we can say that inductive arguments support their conclusions by making them probable rather than entailing them. But this, while true, is unlikely to convince a Humean skeptic who denies either that inductive arguments make their conclusions probable or that probabilistic support is enough to justify inference. A more promising strategy is to argue that we cannot explain our scientific knowledge unless we suppose that inductive arguments support their conclusions as deductive arguments do, although they do this by making their conclusions probable rather than entailing them.

I reply in much the same spirit to the Humean skeptic about moral deduction. We cannot hope to assuage him by appealing to any general criterion that applies equally to moral and logical deductions. Of course, we might say that all deductive arguments support their conclusions by entailing them. But while this is true, it is unlikely to convince a Humean skeptic who denies either that non-moral premises entail moral conclusions or that this is the right kind of entailment to justify deduction. A more promising strategy is to argue that we cannot explain our moral knowledge – or, more cautiously, we cannot explain it so well – unless we suppose that non-moral premises entail

moral conclusions in a way that explains the possibility of deductive knowledge. That was exactly my strategy in §§1-6.

10. The Problem of Hard Cases

This section defends the anti-Humean thesis against the objection that it makes implausible epistemic predictions about hard cases in which we encounter moral disagreement or uncertainty, such as whether it is morally permissible to eat meat, choose an abortion, or inflict the death penalty. I'll consider two different versions of this objection.[34]

The first objection is that we cannot acquire moral knowledge by deduction from non-moral premises in hard cases where the moral facts are disputed. Suppose we disagree about whether it was morally permissible to subject Saddam Hussein to capital punishment for crimes against humanity. The anti-Humean thesis seems to imply that we can resolve our disagreement by a simple deduction so long as we know all the relevant non-moral facts. But that is deeply implausible. As Joshua Schechter (2017) argues, we cannot acquire moral knowledge in hard cases by deduction from our knowledge of the non-moral facts unless we understand well enough how the non-moral facts support our moral conclusions.

In reply, the anti-Humean thesis implies that we're always in an epistemic position to know the moral facts by deduction when we know the non-moral facts that entail them. This follows from a plausible epistemic closure principle:

---

[34] The objection is pressed, in two different ways, by Schechter (2017) and McGrath (2018).

The Closure Principle for Epistemic Positions: Necessarily, if *p* evidentially entails that *q*, and you know that *p* without antecedent knowledge that *q*, then you're in an epistemic position to know that *q* by deduction.

However, this epistemic closure principle doesn't imply that we're always capable of converting our epistemic position into knowledge. In order to know a conclusion by deduction from known premises, you need to perform the deduction in a competent way. We all have some competence to acquire moral knowledge by deduction, but our deductive competence is imperfect. That is why we're competent to acquire deductive moral knowledge in easy cases, but not in hard cases. In hard cases, we need to rely on other ways of acquiring moral knowledge, such as reflecting on the moral relevance of non-moral facts or deferring to moral testimony from those with greater moral expertise.

What is it to make a competent deduction? This is a topic in general epistemology, rather moral epistemology, so I cannot defend a complete account here. Plausibly, however, you don't need to know facts about entailment in order to acquire knowledge by deduction: this would be to over-intellectualize the phenomenon. It is enough that you perform the deduction in a way that manifests the right kind of causal sensitivity to facts about entailment. To a first approximation, you know a conclusion by deduction from known premises just in case your deduction manifests a more general disposition that would not easily be manifested unless the premise entails the conclusion. This is an application of the more general idea that knowledge is always based on general dispositions that are safe from error.[35]

---

[35] Compare Setiya's proposal that doxastic justification requires "an epistemically reliable disposition: a disposition to form beliefs on the basis of evidence by which they are propositionally justified" (2012, 64).

Hard cases arise in every domain of enquiry, from science to detective work. In hard cases, we're not always capable of following our evidence where it leads. As a result, we don't always believe what our evidence supports. And, even when we do, our beliefs are not always properly based on the evidence. Our beliefs are sometimes based on unsafe dispositions that could easily yield beliefs in the absence of adequate evidential support. In such cases, our beliefs are justified in the *propositional* sense, but not the *doxastic* sense. We're not always capable of converting propositional justification into doxastic justification, since we cannot always form beliefs that are properly based on our evidence. This is a general feature of our epistemic predicament that explains why we're not always capable of converting our epistemic position into knowledge. It is not unique to moral epistemology.

The second objection is that the anti-Humean thesis implies that moral disagreement and uncertainty is irrational even in hard cases. After all, epistemic rationality requires proportioning your beliefs to the evidence: if you refrain from believing what is supported strongly enough by your evidence, you are thereby guilty of epistemic irrationality. So, if we know some non-moral fact N that evidentially entails some moral fact M, then we cannot rationally disagree or remain agnostic about M. Epistemic rationality permits only one doxastic stance: namely, believing the moral facts that are evidentially entailed by our non-moral knowledge. And yet this can seem intolerable. As Sarah McGrath writes, "surely you are not being irrational in failing to believe some moral claim simply in virtue of having strong evidence for lower-level claims on which its truth supervenes" (2018, 174).

In reply, the anti-Humean thesis implies that moral disagreement is never perfectly rational, but it doesn't imply that it's always irrational by ordinary standards. As Derek Parfit (1984, 119) notes, 'rational' is a context-sensitive term: to say that someone is irrational by ordinary standards is to say that they deviate so far from perfect rationality that they fall below some context-sensitive threshold. But we can deviate from perfect rationality without falling below this threshold. Our ordinary standards of rationality are somewhat forgiving because they take our human limitations into account.

We do not count ourselves "irrational" in hard cases precisely because these cases are hard for limited human agents. Indeed, we sometimes count ourselves as rational in hard cases even though we deviate from perfect rationality. It is enough that we can expect these deviations from any limited human agent no matter how rational they are. Again, this point arises everywhere in epistemology, not just in the moral domain.[36]

The problem of hard cases arises for any probabilistic theory of epistemic rationality, since it's an axiom of the probability calculus that all necessary truths have probability 1. It's a familiar idea that perfect rationality requires logical omniscience, since all logical truths have evidential probability 1. In the current framework, necessary truths about morality have the same epistemic status as logical truths. If some non-moral fact N evidentially entails the moral fact that M, then it is evidentially necessary that if N, then M; in other words, it is certain given any evidence that if N, then M. Hence, perfect rationality requires not only logical omniscience, but also moral omniscience: it requires being certain of all the necessary truths of morality.[37]

There is no implication that perfect rationality requires being certain of contingent moral truths. Your evidence for contingent moral truths is neither infallible nor certain. For example, if you have misleading evidence that some teenagers are torturing a cat, when in fact they are just acting in a play, then you have misleading evidence that they are acting wrongly. Similarly, if you cannot see exactly what the teenagers are doing, but it looks somewhat likely that they are torturing a cat, then it's at least somewhat likely that they are acting wrongly. Only the necessary truths of morality that are always certain given any body of evidence.

---

[36] For discussion of the distinction between perfect and imperfect rationality, or ideal and non-ideal rationality, see Smithies (forthcoming, b).

[37] See Christensen (2004, Ch. 6) and Smithies (2015) in defense of the thesis that perfect rationality requires logical omniscience.

Ralph Wedgwood (2019) defends a similar view that perfect rationality requires being certain of the ultimate principles of morality. He argues that we ought to comply with the requirements of perfect rationality and, moreover, since 'ought' implies 'can', we have no excuse for violating these requirements. In other words, any deviation from perfect rationality is inexcusable, including any uncertainty about ultimate moral principles. On Wedgwood's view, the difficulty of achieving perfection mitigates our blameworthiness to some extent, but it doesn't remove blame entirely.

This aspect of Wedgwood's view is hard to stomach. A more digestible alternative construes epistemic rationality in evaluative rather than deontic terms. There may be a sense in which we ought to meet the evaluative ideal of perfect rationality and perhaps there is even a sense in which we can achieve this ideal – it may be possible for us in principle even if not in practice. Even so, this is not the sense of 'ought' and 'can' that is relevant for ascriptions of praise and blame. We cannot be blamed for failing to achieve rational perfection when this is possible only given idealized versions of our psychological capacities, rather than the psychological capacities that we actually have.

The problem of hard cases arises throughout epistemology and not just in moral epistemology. In particular, it arises for any probabilistic theory of the evidential support relation, since probabilism implies that perfect rationality requires logical omniscience. As I've explained, there are familiar resources for addressing this problem, which can be applied in moral epistemology as elsewhere. I conclude that the problem of hard cases provides no special reason to reject the anti-Humean thesis.

11. Conclusions

I want to conclude by situating the position defended in this paper within the recent literature in moral epistemology. My position combines three central claims:

(1) *The Anti-Humean Thesis*: You can sometimes know moral conclusions by deduction from wholly non-moral premises.

(2) *Reductive Moral Epistemology*: Your evidence for any moral conclusion is reducible to the evidence for some non-moral premise that entails it.

(3) *Moral Rationalism*: You can sometimes acquire a priori knowledge of general moral principles with absolute certainty.

Each of these claims can be found within the literature, but only in isolation from the other two. Aaron Zimmerman (2010, Ch. 5) defends the anti-Humean thesis, but he rejects moral rationalism and doesn't explicitly endorse reductive moral epistemology. Kieran Setiya (2012) defends reductive moral epistemology, but he explicitly rejects the deductive version of the anti-Humean thesis and so cannot embrace the form of moral rationalism defended here. Meanwhile, Ralph Wedgwood (2019) defends moral rationalism, but he doesn't explain it by combining the anti-Humean thesis with reductive moral epistemology.

This paper does not just combine these three claims but integrates them within a unified moral epistemology. What unifies these three claims is a thesis about the nature of evidential entailment: namely, that every moral fact M is evidentially entailed by some non-moral fact N on which it supervenes. This fact about evidential entailment explains all three claims above. First, it explains the anti-Humean thesis that we can sometimes know the moral conclusion that M by deduction from the known non-moral premise that N. Second, it explains the reductionist thesis that evidence for any moral conclusion that M reduces to evidence for some non-moral premise N which entails that M.

And third, it explains the thesis of moral rationalism that we can sometimes acquire a priori knowledge of the general moral principle that if N, then M by means of suppositional reasoning.

Moral epistemology is a holistic enterprise: to decide between competing positions, we need to perform a complex cost-benefit analysis in which we weigh the strengths and weaknesses of each view against the alternatives. To keep this paper within reasonable bounds, I have focused primarily on the costs and benefits of the anti-Humean thesis. As far as I know, no competing position explains all the same data, but I postpone any further comparisons for future work. In the meantime, I conclude provisionally that there is a compelling case to be made for the anti-Humean thesis. It deserves more serious consideration than it has so far received.[38]

References

Audi, R. (2013). *Moral Perception*. Princeton: Princeton University Press.

Boghossian, P. (2003). Blind reasoning. *Aristotelian Society Supplementary Volume*, 77, 225–248. Doi: 10.1111/1467-8349.00110

———. (2020). Intuition, understanding, and the a priori. In P. Boghossian & T. Williamson, *Debating the A Priori* (186–207). Oxford: Oxford University Press.

Brady, M. (2013). *Emotional Insight: The Epistemic Role of Emotional Experience*. Oxford: Oxford University Press.

Byrne, A. (2005). Introspection. *Philosophical Topics*, 33, 79–104. Doi: 10.5840/philtopics20053312

Carroll, L. (1985). What the tortoise said to Achilles. *Mind*, 4, 278–280. Doi: 10.1093/mind/IV.14.278

Chalmers, D. (2011). The nature of epistemic space. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality* (60–107). Oxford: Oxford University Press.

Christensen, D. (2004). *Putting Logic in Its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.

Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46, 279–295. Doi: 10.1007/BF00372907

———. (2010). Bootstrapping, defeasible reasoning, and a priori justification. *Philosophical Perspectives*, 24, 141–159. Doi: 10.1111/j.1520-8583.2010.00188.x

Dancy, J. (2004). *Ethics Without Principles*. Oxford: Oxford University Press.

Döring, S. (2007). Seeing what to do: affective perception and rational motivation. *Dialectica*, 61, 363–394. Doi: 10.1111/j.1746-8361.2007.01105.x

Fodor, J. (1987). *Psychosemantics*. Cambridge: MIT Press.

Harman, G. (1973). *Thought*. Princeton: Princeton University Press.

———. (1977). *The Nature of Morality: An Introduction to Ethics*. Oxford: Oxford University Press.

Hawthorne, J. (2002). Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65, 247–269. Doi: 10.1111/j.1933-1592.2002.tb00201.x

———. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.

Huemer, M. (2005). *Ethical Intuitionism*. London: Palgrave Macmillan.

Hume, D. (1739). *A Treatise of Human Nature*. London: John Noon.

Jackson, F. (1998). *From Metaphysics to Ethics*. Oxford: Oxford University Press.

Jackson, F., Pettit, P., and Smith, M. (2000). Particularism and patterns. In B. Hooker & M. Little (Eds.), *Moral Particularism* (79–99). Oxford: Oxford University Press.

Kelly, T., and McGrath, S. (2010). Is reflective equilibrium enough? *Philosophical Perspectives*, 24, 325–359. Doi: 10.1111/j.1520-8583.2010.00195.x

Kripke, S. (1980). *Naming and Necessity*. Cambridge: Harvard University Press.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567. Doi: 10.1080/00048409612347521

Littlejohn, C. (2011). Defeating phenomenal conservatism. *Analytic Philosophy*, 52, 35–48. Doi: 10.1111/j.1468-0149.2011.00516.x

———. (2012). *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.

Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. London: Penguin.

McBrayer, J. (2010). A limited defense of moral perception. *Philosophical Studies*, 149, 305–320. Doi: 10.1007/s11098-009-9363-0

McGee, V. (1985). A counterexample to modus ponens. *Journal of Philosophy*, 82, 462–471. Doi: jphil198582937

McGrath, M. (2017). Knowing what things look like. *Philosophical Review*, 126, 1–41. Doi: 10.1215/00318108-3683602

McGrath, S. (2004). Moral knowledge by perception. *Philosophical Perspectives*, 18, 209–228. Doi:

10.1111/j.1520-8583.2004.00026.x

———. (2018). Moral perception and its rivals. In A. Bergqvist & R. Cowan (Eds.), *Evaluative Perception* (161–182). Oxford: Oxford University Press.

Moore, G. E. (1903). *Principia Ethica*. Mineola: Dover Publications.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Peacocke, C. (1992). *A Study of Concepts*. Cambridge: MIT Press.

Prior, A. N. (1960). The autonomy of ethics. *Australasian Journal of Philosophy*, 38, 199–206. Doi: 10.1080/00048406085200221

Putnam, H. (1975). The meaning of 'meaning.' *Minnesota Studies in the Philosophy of Science*, 7, 131–193.

Qu, H. (2019). Laying down Hume's Law. *Pacific Philosophical Quarterly*, 100, 24–46. Doi: 10.1111/papq.12231

Sinnott-Armstrong, W. (1996). Moral skepticism and justification. In W. Sinnott-Armstrong & M. Timmons (Eds.), *Moral Knowledge? New Readings in Moral Epistemology* (3–48). Oxford: Oxford University Press.

Schechter, J. (2017). Difficult cases and the epistemic justification of moral belief. *Oxford Studies in Metaethics*, 12, 28–50.

Schechter, J. & Enoch, D. (2006). Meaning and justification: the case of modus ponens. *Noûs*, 40, 687–715. Doi: 10.1111/j.1468-0068.2006.00629.x

Setiya, K. (2012). *Knowing Right From Wrong*. Oxford: Oxford University Press.

Skyrms, B. (2000). *Choice and Chance: An Introduction to Inductive Logic*. London: Wadsworth.

Smithies, D. (2015). Ideal rationality and logical omniscience. *Synthese*, 192, 2769–2793. Doi: 10.1007/s11229-015-0735-z

———. (Forthcoming, a). The problem of morally repugnant beliefs.

———. (Forthcoming, b). The epistemic function of higher-order evidence. In L. Oliveira & P. Silva (Eds.), *Propositional and Doxastic Justification: New Perspectives in Epistemology*. London: Routledge.

Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 141–153. Doi: 10.1111/0029-4624.33.s13.7

———. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1*. Oxford: Oxford University Press.

Tappolet, C. (2016). *Emotions, Values, and Agency*. Oxford: Oxford University Press.

Wedgwood, R. (2013). A priori bootstrapping. In A. Casullo & J. Thurow (Eds.), *The A Priori In Philosophy* (226–246). Oxford: Oxford University Press.

———. (2015). An inferentialist conception of the a priori. *Oxford Studies in Epistemology*, 5, 295–314.

———. (2019). Moral disagreement and inexcusable irrationality. *American Philosophical Quarterly*, 56, 97–108.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

———. (2007). *The Philosophy of Philosophy*. Oxford: Wiley-Blackwell.

———. (2013). Knowledge first. In M. Steup & J. Turri, *Contemporary Debates in Epistemology* (1–10). Oxford: Wiley-Blackwell.

———. (2019). Morally loaded cases in philosophy. *Proceedings and Addresses of the American Philosophical Association* 93: 159–172.

Zimmerman, A. (2010). *Moral Epistemology*. London: Routledge.