Is, Ought, and the Regress Argument

Abstract: Many take the claim that you can't 'get' an 'ought' from an 'is' to imply that non-moral beliefs are by themselves incapable of justifying moral beliefs. I argue that this is a mistake and that the position that moral beliefs are justified exclusively by non-moral beliefs– a view I call *moral inferentialism* – presents an attractive non-skeptical moral epistemology.

> *"Why should I believe this set of propositions?" is quite different from the logical question: "What is the smallest and logically simplest group of propositions from which this set of propositions can be deduced?"*
> –Bertrand Russell, *The Philosophy of Logical Atomism p. 129*

§1   Is to Ought

Of the many theses that have been held about the relationship between thoughts about what *is* and thoughts about what *is good* or what *ought to be*, one of the most straightforward often gets overlooked. The following is, I imagine, a common experience: a thought crosses your mind about how things are, were, or might be. Next, you have an evaluative or a deontic thought – a thought about goodness, badness or about what ought to be done. Not only does the one thought follow the other, but it seems the first thought – the thought about what 'is' – *justifies* the second thought about value or obligation. For instance, you think: 'The old man collapsed and is in pain,' and that seems to support the thought that something bad has happened and you ought to help. As Frankena puts it, in a passage discussed at length below, 'we do sometimes seem to justify an ethical judgment by an appeal to fact.'

But Frankena adds, 'it becomes clear on a moment's thought that our [ethical] conclusion does not rest on our factual premise alone.' Most who write on such matters agree that thoughts about *what is* are themselves always insufficient to justify thoughts about what is good or ought to be done. I attack this consensus in §§3-4 by showing that the doctrine that you can't 'get' an 'ought' from an 'is' rests on certain obvious mistakes. There are notable exceptions to the general consensus, but even those who grant

that an 'is' may justify an 'ought', usually think such justifications are *rare.* In §§5-6, I argue that the justification of ethical thoughts by exclusively non-ethical thoughts is *pervasive.*

In the next section, to show that the view I defend is a neglected part of the conceptual space, I introduce the moral regress argument – a skeptical argument that distinguishes the ways a belief might be justified and claims that no moral beliefs are justified in those ways. Those who discuss this argument typically dismiss the idea that moral thoughts are justified by exclusively non-moral thoughts. But, as I'll argue, this view represents a plausible solution to the skeptical challenge and has independently attractive implications for moral investigation and argumentation. I discuss these implications, including a way the view might address another central skeptical problem – the accidentality challenge – in §§7-8.

§2   The Regress Argument

The regress argument begins with a distinction between inferential and immediate justification. Inferentially justified beliefs depend on other beliefs for their justification. If my belief is justified inferentially, it's because I can cite other things I believe that make the belief in question reasonable.[1] Immediately justified beliefs, in contrast, do not depend on other beliefs for their justification. They are justified by some non-doxastic state. Traditional candidates include the experience of the agent or the self-justifying character of the belief's content. To explain how an immediately justified belief is justified, you might say, 'I can just see it!' or 'If you reflect on the proposition, you'll understand that it must be true!'

---

[1] There may be special reasons why, in some instance, I can't cite the justifying beliefs. I may be distracted or incapacitated. We can ignore these complications here.

This distinction is usually understood to be exhaustive.[2] Thus the regress argument begins:

(1) If a moral belief is justified, it is either justified immediately or inferentially.

The argument's next step is to list the two ways a belief can be justified immediately and then to deny that any moral beliefs are justified in either of those ways.

(2) If a moral belief is justified immediately, it is either justified by the experiences of the agent or by the self-justifying character of the belief's content.

(3) No moral belief is justified by the experiences of an agent.

(4) No moral belief has a self-justifying content.

Externalists deny (2), holding that a belief might be immediately justified by, e.g., the law-like character of the connection between the belief and the state of affairs that make it true [Star 2008]. Others deny (3), and claim that some distinctively moral perceptions constitute experiences that justify some moral beliefs [Mcgrath 2004; McBrayer 2010a, 2010b; Audi 2013]. Usually such justification is thought only to apply to beliefs about the moral properties of some *particular* person, action or state of affairs. Some deny (4), claiming that some moral beliefs are *self-evident* [Sidgwick 1879; Moore 1903; Ross 1930; Shafer-landau 2003; Audi 2005]. Usually this kind of justification is thought to apply to general moral beliefs, either principles like 'lying is wrong' or highly abstract axioms like the principle of utility. All three of these positions are forms of *moral foundationalism.* This family of views claims both that some moral beliefs

---

[2] But it need not be understood as exclusive. Some moral belief might be justified either by other beliefs or by some non-doxastic state. Understanding this disjunction as inclusive is one mark of the "new intuitionism" championed by [Audi 2005].

are justified immediately and that everything else we know about morality can be justified inferentially from these immediately justified beliefs.

Once moral foundationalism has been dispatched, the argument considers the possibility that our moral beliefs are all justified inferentially. Here the skeptic sets the stage for a dilemma:

(5) If a moral belief is justified inferentially, there are either moral beliefs among the justifiers or not.

No one denies this tautology. It is also rare for philosophers to deny the next premise, which is the focus of the sections to come.

(6) No moral belief is justified inferentially with only non-moral beliefs among the justifiers.

(1)-(6) entail that if a moral belief is justified, it is justified inferentially with other moral beliefs among the justifiers. The final step claims that this is impossible.

(7) It is not the case that every justified moral belief has moral beliefs among its justifiers.

If all moral beliefs *were* justified by other moral beliefs, then either (a) these justifying moral beliefs would themselves fail to be justified, (b) our chains of justification would be objectionably circular, or (c) our chains of justification would be objectionably infinite. *Contextualists* allow that unjustified beliefs can justify other beliefs [Wellman 1971; Annis 1978; Timmons 1998]. C*oherentists* do not take the circularity to be objectionable [Brink 1989; Sayre-McCord 1996]. I*nfinitists* embrace the conclusion that our chains

of justification go on indefinitely.[3] Though these positions have able defenders, they also face difficult challenges.

(1)-(7) imply that there are no justified moral beliefs. To avoid skepticism we must find a premise to deny.

Most philosophers who address the regress argument, including those cited above, refer to the doctrine that you can't 'get' and 'ought' from an 'is' and then accept (6) too quickly and, as I hope to show, for bad reasons. Since almost everyone accepts (6) and settles on some version of foundationalism, coherentism, contextualism or skepticism, there is no standardly accepted name for the position that the denial of (6) represents.[4] We might call the position *moral inferentialism*.[5]

There are well known problems for other non-skeptical solutions to the regress problem, so it's worth exploring the idea that there is widespread justification of moral beliefs by non-moral beliefs. A full defense of this claim, however, is beyond the scope of this paper. What I will argue is simply that certain standard objections to moral inferentialism are not as strong as many think and that the view presents an attractive non-skeptical moral epistemology.

§3   Justification, Entailment and Moral Principles

---

[3] There are no specifically *moral* infinitists. But the general position is defended in [Klein 2007].
[4] Some *have* defended the position. It's arguably the view of [Anscombe 1958a, 1958b]. [Thomson 1996; Nelson 2003] come closest to an explicit endorsement. [Zimmerman 2010] discusses this possibility sympathetically. [Setiya 2012: 49] also suggests that there are some non-moral to moral justificatory relations, though he implies that such relations are rare.
[5] The term 'inferentialism' is often used for a particular theory of *meaning* that holds that the meaning of some term is a matter of the inferences that can be made to or from sentences involving that term. See [Brandom 2000] and, for a related point about moral terms, [Wedgewood 2001]. Though nothing here rules out the possibility that the meaning of moral terms might be explicated along these lines, my thesis is about *justification,* not meaning.

There are four mistakes that explain why many accept (6): (i) They confuse justification with entailment. (ii) They think that when some belief or set of beliefs justifies another belief, the agent must have in mind some principle connecting the two. (iii) In the context of the regress argument, they take a cavalier attitude towards the truth of certain moral principles. Finally, (iv) they fail to appreciate the differences between inferences and corresponding principles, neglecting important features of inferential justification. In the remainder of this section and the next, I'll use a classic passage by Frankena to illustrate these mistakes. As references below and in the previous section show, this way of thinking remains widespread.

> [W]e do sometimes seem to justify an ethical judgment by an appeal to fact. Thus, we say that a certain act is wrong because it injures someone, or that a certain painting is good because it has symmetry. However, it becomes clear on a moment's thought that our conclusion does not rest on our factual premise alone. In the first case, we are tacitly assuming that injurious acts are wrong, which is a moral principle; and in the second, that paintings with symmetry are good, which is a value judgment. In such cases, then, we are not justifying our original ethical judgment by reference to fact alone but also by reference to a more basic ethical premise. [Frankena 1973: 97]

The claim that 'our conclusion does not rest on our factual premise alone,' is ambiguous. A conclusion 'resting' on a premise can refer to the *truth* of a conclusion being *guaranteed* by the *truth* of some premise(s). It can also refer to *belief* in a conclusion being *justified* by belief in some premise(s). Premise (6) of the regress argument says that no moral belief can be *justified* by exclusively non-moral beliefs. Though we may think that non-moral propositions by themselves are incapable of *guaranteeing* the truth

of a moral proposition, it does not follow that non-moral beliefs are by themselves incapable of justifying a moral belief.[6]

Frankena later suggests that he is talking about justification, but he may have illicitly slipped between the two senses of 'rests.' Many thinkers confuse inferences that preserve truth with inferences that preserve justification. For example, Huemer gives this argument for foundationalism:

> Undoubtedly some moral beliefs are accounted for by inference from other moral beliefs. But since no moral belief can be derived from wholly non-moral premises, we must start with some moral beliefs that are not inferred from any other beliefs. [Huemer 2005: 103]

This is the same mistake. It's true that no moral belief is entailed by wholly non-moral premises, but it is still possible that non-moral premises are themselves sufficient to *justify* a moral conclusion.

Here is another example of the same illicit move in David Brink's discussion of Sidgwick:

> Sidgwick defends the need for some moral intuitions by appeal to … a logical gap between is and ought and a familiar regress argument (cf. ME: 98). His argument has … the following form.
>
> 1. No moral statement can be derived from purely nonmoral statements.
>
> 2. Hence, if a moral statement is to be derived from further claims, these claims must contain moral statements.

---

[6] There are classic counterexamples to the claim that no moral proposition can be *deduced from* (or*, is entailed by)* exclusively non-moral premises [Prior 1960; Searle 1964; Kurtzman 1970]. But see [Jackson 1974; Brown 2014] for statements of the doctrine that avoid these objections.

3. Hence, if a moral belief is to be inferentially justified, its justification must appeal to other moral beliefs.

4. But, on pain of infinite regress or circularity, the process of inferential justification must come to a stop somewhere.

5. Hence, there must be some self-evident moral beliefs. [Brink 1999: 203]

If we remember that 'derived from' can mean either 'entailed by' or 'inferentially justified by,' we can see that an equivocation leads Brink's Sidgwick from (1) to (3).[7] Perhaps these authors hold the thesis that for one belief to justify another, the first must entail the second. But they do not give an argument for that assumption. If we were to accept it, very few of the beliefs we normally take to be justified would, in fact, be justified. It is not often that we can produce valid arguments from known premises to justify our inferentially justified beliefs.

If we interpret Frankena's 'rests' as referring throughout to justification and not entailment, then he makes a different mistake: thinking that for P to justify Q, one must *believe* something along the lines of 'if P, then Q.'[8] He suggests this when he says we are 'tacitly assuming …' This step in the argument is startling because it requires us to believe, e.g., that ANY injurious act is wrong. This is the cavalier acceptance of a false moral principle. I almost never believe a simple conditional connecting the reasons I think something is wrong with its being wrong. Ethical thinking is shot through with exceptions, defeated reasons, holistic considerations, etc. [Dancy 2004, Horty 2012].

---

[7] Sidgwick shares with Frankena, Huemer and Brink the view that moral premises are needed to justify moral conclusions.

> The premises of our reasoning [to ethical conclusions] … must contain, implicitly or explicitly, the elementary notion signified by the term "ought"; otherwise, there is no rational transition possible to a proposition that does affirm "what ought to be." [Sidgwick 1879: 107]

[8] Zangwill [2006: 276] and Sinnott-Armstrong [2007: 76] make similar assumptions.

To save Frankena's point about our tacit assumptions, we need to interpret 'injurious acts are wrong,' in another way. There are two strategies: weaken the consequent of the conditional or add some *hedging* terms to the antecedent. An example of the first strategy would yield a Rossian-style principle, 'If an act is injurious, then that act is *pro-tanto* wrong.' This principle claims that the fact that an act is injurious always contributes to the wrongness of the act, though there may be other considerations that make the act permissible overall. But this ignores the possibility that some act may be permissible *because* it is injurious (e.g. on a retributivist view of punishment). It would again be cavalier acceptance of a doubtful moral principle to insist that we need a pro-tanto principle to 'connect' the premise that an act is injurious to the conclusion that it is wrong. And even believing some true pro-tanto principle in this case cannot save Frankena's point. What his position requires is the claim that there is *always* some true pro-tanto principle that can 'complete' our inferences from non-moral to moral thoughts.

The other interpretation of Frankena involves adding some hedging term to the antecedent of 'If an act is injurious then that act is wrong.' Candidates include 'prima facie,' 'ceteris paribus,' or 'if nothing else is morally relevant.' This strategy is more promising; many such principles *are* things we can accept even after some scrutiny. However the trouble here is twofold. (i) It isn't clear how we are justified in believing these kinds of hedged principles. Unlike the simple conditional and the pro-tanto principle, hedged principles are not good candidates for self-evident truths.[9] (ii) Even if we did justifiably believe such principles, it isn't clear why believing 'ceteris paribus, injurious acts are wrong,' in addition to believing 'this act is injurious,' makes 'this act is wrong' justified. What does the principle add, aside from the complication that we must now wonder if the ceteris paribus clause holds?

§4   Beliefs and Inference Rules

---

[9] A defense of this claim is beyond the scope of this paper, but one quick way to see why it is plausible is to ask: Where, outside of ethics, would anyone claim that a ceteris paribus or prima facie principle was self-evident?

After the passage quoted above, Frankena says that, if our moral beliefs *were* justified by non-moral beliefs, then …

> This would mean that conclusions with terms like "ought" and "good" in them can be logically inferred from premises, none of which contain these terms; this simply cannot be done by the rules of ordinary inductive or deductive logic. To try to do so is essentially to argue that A is B, ∴ A is C, without introducing any premise connecting B and C. In this sense, those who insist that we cannot go from Is to Ought or from Fact to Value are correct. Such an inference is logically invalid unless there is a special third logic permitting us to do so … But the theory and the rules of such a logic have not yet been satisfactorily worked out, and until they are we can hardly take this possibility seriously. In any case, it is hard to see how such a "rule of inference" differs in substance from the "premise" that injurious acts are wrong, or how its justification will be different. [Frankena 1973: 97]

It is strange to insist that we need a *logic* to permit us to make certain inferences – as if the reasonableness of e.g., our inductive inferences, depends on the antecedent development of an adequate inductive logic. The more plausible picture is that we develop various logics to capture and systematize what we already accept as reasonable inferences. As noted above, it is also strange to suggest that whenever we make a justified inference from P to Q we must justifiably believe some corresponding principle to 'connect' P to Q. Lewis Carroll [1895] illustrated the absurdity of this general claim. In his dialogue, Achilles and the Tortoise consider the following argument:

> (A) Things that are equal to the same are equal to each other.
>
> (B) The two sides of this triangle are things that are equal to the same.

∴ (Z) The two sides of this triangle are equal to each other.

The Tortoise claims that he believes (A) and (B), but does not accept (Z) and he asks Achilles to 'force [him] logically' to accept (Z). Achilles suggests the Tortoise accept:

(C) If (A) and (B) are true, then (Z) is true.

The Tortoise complies, but claims that, even though he believes (A), (B) and (C), he might still fail to believe (Z). To remedy this, Achilles now asks him to accept a new conditional:

(D) If (A), (B) and (C) are true, then (Z) is true.

When the Tortoise comes to believe (D), he says that, believing (A) - (D), he might still fail to believe (Z). The dialogue ends with a deflated Achilles proposing increasingly complicated conditionals which the Tortoise accepts, getting no closer to (Z).

If we always needed to justifiably believe a principle relating premises to a conclusion in order to justifiably infer that conclusion from those premises, then we, like the Tortoise, would require an ever-expanding set of beliefs to make the simplest justification-preserving inference. It is perfectly natural for me to believe Q on the basis of P without thinking 'If P then Q' or any variation thereof.

Frankena and company might agree that we don't always need to believe some principle to have one belief inferentially justify another, while still insisting that we do need a principle in the non-moral to moral case because in this case there is a new predicate in the conclusion that doesn't appear in the

premises. Against this suggestion, notice that we *do* often reason from premises like 'A is B' to conclusions like 'A is C' without precise beliefs connecting B and C. When playing games, for instance, I make inferences from the behavior of my opponents to conclusions about their intentions, about what they'll do next, and about what I should do in response. But I can't articulate principles connecting my beliefs about their behavior to conclusions about these other things. While this suggestion from Frankena wouldn't require *as many* beliefs as the tortoise required to account for intuitively obvious cases of knowledge, it requires many more than seems reasonable.

Suppose Frankena admits that we don't need principles for *all* cases where new predicates are introduced, but still insists we need them in the non-moral to moral case. Without some additional rationale to think that non-moral beliefs are insufficient to justify moral ones in the absence of a moral principle, we can answer this insistence with inspection of cases. As examples in the next two sections will show, we often justifiably believe moral claims on the basis of non-moral claims even when we lack any clear belief in a connecting principle.

Finally, consider the assertion that there is no 'substantive difference' between a non-moral to moral 'rule of inference' and a corresponding moral principle. Frankena might grant that one does not need to justifiably believe a moral principle to justifiably infer a moral conclusion from non-moral premises. But, by emphasizing the distinction between *justifiably believing a principle* and *being justified in believing a principle*, he can claim that if our inference is really justified, then we would *be justified in believing some principle*. Whatever it takes to justifiably infer Q from P would justify belief in a principle relating P to Q.

If we were simply arguing that premise (6) of the regress argument is false, then we could accept this suggestion. Recall that (6) claimed that no moral belief is justified solely by non-moral beliefs. The claim that justifying an inference would *eo ipso* justify some moral principle is compatible with the rejection of (6).

But before we accept this suggestion, it's worth noting that the justification of an inference and the justification of a corresponding belief are not equivalent issues. Inferences are something we *do,* while principles are something we *believe.* In general, the considerations that show a belief to be justified must speak to the *truth* of that belief. But there is no question of the *truth* of an inference. We can only ask about its usefulness for furthering our epistemic or practical goals. Consider a case where you are out to lunch, deciding whether to order steak or salad. You might think, 'ordering steak causes harm to animals.' Suppose, after considering the evidence, you aren't sure if it's wrong to cause harm to animals. Still, it may be reasonable to infer that you shouldn't order the steak. The inference may be justified because it's reasonable to hedge your moral bets. Because an inference is something we *do,* these kinds of pragmatic considerations make sense. They're more problematic when considering the justification of a corresponding belief.

Nothing said here precludes the possibility that *sometimes* the question of what justifies an inference is related to the question of what justifies a corresponding belief. Sometimes, we may need belief in some principle – or at least justification for believing some principle – to allow for inferential justification. But what the rejection of moral inferentialism would require is that you *always* need a justified belief in a moral principle – or the justification for believing a moral principle – to make a justified inference from non-moral beliefs to moral ones. In general, claims about the impossibility of 'deriving an ought from an

is' are insufficient for that result. Premise (6) of the regress argument is more difficult to establish than many have thought.

§5   Can All Moral Beliefs Be Justified By Inferences from Non-moral Beliefs?

I've argued that the reasons philosophers accept (6) are inadequate, but what can be said positively for the denial of (6)? And even if it's possible that *some* moral beliefs are justified inferentially from non-moral beliefs, is it plausible that *all* our moral beliefs are justified in this way?

I think it is. When I reflect on my moral beliefs, I find that I hold each of them as the result of holding some non-moral belief. If I gave up the non-moral belief then I would also give up the moral one. This is most obvious when it comes to moral beliefs about particular actions, persons or states of affairs: I believe what she did was wrong because I believe she dropped a bowling ball out the window; I believe he is courageous because I believe he often overcomes fear to do what he thinks right; and I believe that things in Syria are bad because I believe that many people there are suffering and afraid. Doubting the non-moral beliefs in these examples would cause me to doubt the moral ones.

Each of these cases involves an ascription of a deontic, evaluative or virtue term to a particular action, state or person. My thesis may not seem controversial for claims of this type. When an action, state or person has some moral property, it always has that property in virtue of having other, non-moral properties. But the claim also holds for more general moral beliefs. For instance, I believe it's wrong to hit someone over the head with a frying pan because I believe doing so causes pain and injury. I believe clean drinking water is good because I believe it promotes health. Doubting the non-moral beliefs would cause me to doubt the moral ones.

Admittedly, I've only noted that the justification of some moral beliefs depends, in part, on some non-moral beliefs. While many will admit that non-moral beliefs play an important role in the justification of *most* moral beliefs, they will also insist: (A) that the justification for these moral beliefs *also* depends on my belief in some moral principle, and (B) that my justification for believing some moral principles *does not depend* on any non-moral beliefs.

The reason for doubting (A) has already been mentioned: it's difficult to frame the moral principle that both 'completes' the inference and is believable. I almost never believe simple conditionals connecting the reasons I think something is wrong to its being wrong. Returning to our first example, though I believe that she did something wrong *because* I believe she dropped the bowling ball out the window, I don't believe that if someone drops a bowling ball out the window, then that person has done something wrong. There may be circumstances where dropping a bowling ball is permissible. Likewise, I don't believe that if something (in this case, frying pans to the head) causes pain and injury, then that act is wrong. There are cases where causing pain and injury is the right thing to do.

Even if we don't believe simple conditionals connecting the reasons we think something is wrong with its being wrong, it is natural to suspect that the kinds of justifying beliefs I have cited are elliptical and that more is needed to 'complete' the inference. But unless we require entailment for inferential justification, why should we think we need a *moral* principle?

Against (B), notice that our justification for believing even very abstract moral principles may depend on our non-moral beliefs. I believe, for example, that pain is bad. I believe this because I believe that pain causes anxiety and discomfort and that people desire not to be in pain. If I were to give up these non-moral beliefs, I would doubt my belief that pain is bad. Likewise, I believe that if it's right for me to

do some act, it would be right for similarly situated persons to do similar acts. I believe this because I think that, in certain respects, people are the same. If I were to give up this non-moral belief, I might give up my moral belief. As above, I don't believe simple conditionals like: if something (e.g. pain) causes anxiety and discomfort, then that thing is bad. Some causes of anxiety and discomfort (weddings, competition) are good things.

Agreeing with me that the belief 'pain is bad' needs to be justified by some non-moral belief is compatible with the view that pain is intrinsically bad. In general, it is possible to think that the belief, 'X is bad' is justified by the belief 'X is P' without thinking P makes X bad. We shouldn't conflate reasonable epistemological claims like, 'the fact that people desire not to be in pain justifies the belief that pain is bad,' with controversial metaphysical ones like, 'pain is bad because people desire to avoid it.'

Failure to appreciate this point lies behind Moore's famous criticism of Mill. Moore says,

> Mill has made as naïve and artless a use of the naturalistic fallacy as anybody could desire. 'Good,' he tells us, means 'desirable,' and you can only find out what is desirable by seeking to find out what is actually desired. [Moore 1903: §40]

In Mill's defense, it is possible to take the fact that something is desired to be evidence – perhaps the sole evidence – for its desirability without identifying the desirable with the desired. Like the writers above who conflated entailment and justification, Moore conflates the claim that the instantiation of one property is evidence for the instantiation of another with the claim that the properties are identical.[10]

---

[10] In the same section, Moore says that Mill's analogy of desired/desirable with seen/visible shows that Mill is confused since the analogy requires understanding desirable as 'able to be desired,' instead of 'ought to be desired.' Whatever might be said against Mill's analogy, it shows he *wasn't* committing the naturalistic fallacy by identifying

§6 More Examples and What You Really Know

We justifiably believe many kinds of moral propositions. I've discussed *verdictive* moral beliefs like 'she was wrong to drop the bowling ball,' that ascribe moral properties to particular persons, actions or states of affairs. When persons, actions or states have some moral property, they do so in virtue of having other, non-moral properties. So, we can say that such beliefs are justified by beliefs about the non-moral properties on which the moral ones are consequent.

Instead of offering verdicts, other kinds of moral beliefs about particular persons, actions, or states, ascribe dependence relations between non-moral and moral properties. Not only might we think that she did something wrong by causing injury but we might think that it was the injuriousness of her act that made it wrong. Such beliefs are also, I hold, justified by non-moral beliefs. To say which non-moral beliefs I'll need to introduce some terminology. The relationship between a moral property and the non-moral property on which it is usually consequent is, under certain conditions, broken. Call these conditions *defeating conditions*. Defeating conditions in this case might include the consent of the injured or the fact that the injury was given as punishment. To justify the belief that the injuriousness of her act made it wrong, we would need to infer it from the fact that the injury was not consented to and was not given as punishment. In general, the belief that some particular person, act, or state has some moral property *because* it has some non-moral property can be justified by the belief that defeating conditions do not obtain.

Some moral principles are generalizations with simple non-moral antecedents. Such principles (e.g. clean drinking water is good) can be justified by beliefs about the non-moral properties of the subject of the

---

'desirable' with 'desired.' No one would identify 'visible' with 'seen,' even if it's reasonable to think that what has been seen is evidence – perhaps the sole evidence – for what is visible.

antecedent (e.g. clean drinking water promotes health). Principles which are not simply generalizations, but assert some dependence between a non-moral property and a moral one (e.g. injurious acts are wrong) can be justified by general beliefs about defeating conditions (e.g. people do not consent to be injured). If there are no defeating conditions (e.g. pain is bad) the justifiers will be purely evidential (e.g. people desire not to be in pain).

Sometimes a moral principle is best understood as a universal generalization, sometimes as a statistical one, sometimes as a generic,[11] and sometimes as containing implicit hedging terms. In each case, the kind of non-moral belief that can justify adjusts accordingly. For instance, to justify 'For all acts, if the act is injurious, then it is wrong,' we would need a strong non-moral belief to the effect that the defeating conditions never obtain. If we were only trying to justify the statistical generalization that 'most injurious acts are wrong, 'we would need a non-moral belief like 'most of the time, people don't consent to be injured.' Likewise for generic and hedged principles.

Generally, the difficulty of justifying a moral proposition is a function of the difficulty of coming to know the non-moral facts that do the justifying and the strength of the connection between these non-moral and moral propositions. Universal and highly general moral beliefs are more difficult to justify, as are principles with no defeating conditions. Verdictive beliefs and various kinds of qualified, statistical or generic principles are more easily justified.

I've given examples of how various moral beliefs might be justified by various kinds of non-moral belief. One might hope to say something more systematic. In particular, one might think that inferences from

---

[11] Generics are generalizations with no explicit quantification. Examples include 'ducks lay eggs' and 'mosquitoes carry west nile virus.' There is disagreement about how to understand the logical form of generics so it's difficult to give an uncontroversial account of their truth conditions [Leslie 2007]. There hasn't been much work interpreting moral principles as generics, but it seems like a promising hypothesis [Lerner and Leslie 2013].

non-moral beliefs provide justification for moral ones only if the pattern of inference is a good one. What, then, constitutes a good non-moral to moral inference pattern?

Answering that question is, on the moral inferentialist's view, one of the central aims of moral philosophy. No answer can be given in advance of extensive moral theorizing. To expect otherwise is like asking 'how do we predict the future from observations about the past?' and thinking that an answer can be given in advance of some form of scientific theorizing. If you want to know what inferences from beliefs about the past to beliefs about the future are justification preserving, you'll have to do some science. Likewise, if you want to know what inferences from non-moral beliefs to moral ones are justification preserving, you'll have to do some moral philosophy.

§7 The Accidentality Challenge

I've argued that non-moral beliefs by themselves are capable of justifying moral beliefs, that this thesis is dismissed for insufficient reasons, and that it presents an attractive solution to the regress problem. In these final two sections, I'll explore some implications of this view and some ways it might be extended to deal with another family of skeptical problems.

The skeptical problems I have in mind begin from claims about the origin of our moral beliefs. One version notes that our moral beliefs are dependent on the time and place of our upbringing. Another asserts that evolutionary forces influenced our moral beliefs. Others say that our moral views are a product of a particular history. Finally, some assert that our moral beliefs simply have *some* naturalistic explanation. These etiological claims paired with certain views about the nature of moral truth (that it is not dependent on the time and place of our upbringing, evolutionary forces, historical contingencies, or that it is *nonnatural*) leads to a skeptical conclusion. If the explanation of our moral beliefs cannot show

how those beliefs are connected to moral truths, then it would be a massive coincidence if our moral beliefs turned out true. Even if, by some miracle, they *were* true, it would be an *accident*. Such beliefs would be *gettierized,* at best.[12]

Here is how a moral inferentialist would approach this problem. First, notice that our non-moral beliefs can explain our moral beliefs. My belief that she dropped the bowling ball out the window explains my belief that she did something wrong. Next, unless all explanans must *entail* their explanandum, it is natural to think that moral truths can be explained by non-moral facts. If a child asked 'why is what she did wrong?' part of the answer would be, 'because she dropped the bowling ball out the window.' Finally, it is uncontroversial that non-moral facts explain our non-moral beliefs. In this case, the fact that she dropped the ball explains my *belief* that she dropped the ball. So that fact – that she dropped the ball – explains both the moral fact that she did something wrong and, via my non-moral belief that she dropped the ball, my moral belief that she did something wrong. The moral truth and my moral belief have a common explanation; that is how moral beliefs can be non-accidentally true.

This account of the non-accidentality of our moral beliefs is rarely taken seriously for reasons that are, by now, familiar. Since you can't 'get' an 'ought' from an 'is,' it's assumed that non-moral beliefs can't explain moral beliefs and that non-moral facts can't explain moral truths. Of course, showing that non-moral beliefs *can* sufficiently explain moral beliefs and that non-moral facts *do* sufficiently explain moral truths involves additional issues. First, what kind of explanatory connections can account for the non-accidentality of our beliefs? Second, how do we integrate different explanations? My belief that she dropped the ball is *part* of the explanation for my belief that she did something wrong. But my moral

---

[12] The term 'gettierized' refers to cases made famous by [Gettier 1963]. See [Yamada 2011] for an attempt to analyze the relevant notion of accidentality. There is a dispute about how to understand non-accidentality. Many think it requires a modal condition. Others understand the issue as I have, in terms of explanation. It should be clear how to modify my claims for modal accounts.

belief might also be partially explained by the evolution of reciprocal altruism, the contingencies of my upbringing, etc. When the whole picture is considered, is my moral belief explained in a way that 'connects' it to the moral truth? Finally, we may wonder *how* non-moral facts explain moral facts and whether accepting these explanatory relations is compatible with our favored view about moral metaphysics. E.g., can a non-naturalist admit that non-moral facts sufficiently explain moral truths?

However these questions are answered, there is something intuitively appealing about the inferentialist approach. Instead of *all* our moral knowledge being threatened by e.g., its evolutionary origin, *particular* beliefs or sets of belief will stand or fall with our success in identifying some non-moral fact that could explain both our moral belief and the moral truth. If someone challenged your knowledge that she did something wrong on the grounds that your belief has evolutionary origins, you could successfully confront this challenge by citing the fact that she dropped a bowling ball out the window. But if you can't find such a fact – if the facts that explain your belief can't explain her wrongdoing (e.g. she looks suspicious) or if the facts that explain the moral truth can't explain your belief (e.g. she had malevolent intentions about which you were ignorant) – then the skeptical challenge succeeds.

§8   Moral Doubt, Moral Investigation and Moral Philosophy

Sometimes we doubt our moral beliefs. A theory of justification tells us what considerations could rightfully put our doubts to rest. The theories I've mentioned in discussing the regress argument make different claims about how to render our beliefs defensible. Moral Foundationalists claim that, when you doubt a moral belief, you should look to see if that belief is immediately justified, either by some feature of the content of the belief (its self-evidence) or by some non-doxastic mental state (an appearance or seeming). If it isn't defensible in one of these ways, it must be inferable from moral beliefs that *are* defensible in one of these ways and the non-moral facts. Coherentists tell us that, when we doubt a moral

belief, we should see if it coheres with the rest of our moral beliefs and the non-moral facts. Because of a misunderstanding of the claim that you can't 'get' an 'ought' from an 'is,' these have often been seen as the only viable positions.[13] Moral inferentialism presents a third alternative. To check if a moral belief – whether principle or verdict – is defensible, we should look to the non-moral beliefs on which our moral belief is based and ask ourselves if they support the moral belief in question. Importantly, this is not a matter of finding some moral principle to connect the non-moral facts to the moral ones. To combat skepticism, the best we can do is to remind ourselves of the important non-moral facts: people really do suffer and want various things. We might wonder if these facts are sufficient to allay our skeptical doubts, but trying to convince ourselves that some moral belief is self-evident, or that we have moral-perceptual experiences, or that our moral beliefs form some coherent whole, work no better as guards against skepticism than reminding ourselves of the most important non-moral facts.

Theories of justification also tell us something about how knowledge-oriented interlocutors should investigate their beliefs together. On the view sketched here, moral discussion is primarily an exchange of non-moral information. If you want to convince me that some moral proposition is true, tell me the non-moral propositions on which it is based. My other beliefs might prevent those particular non-moral propositions from justifying the moral belief under discussion. Or I might fail to 'see' the relevance of those non-moral propositions. In either case, disputes about principles which would 'complete' the inference, about the self-evidence of abstract principle, about moral 'seemings,' or about the overall coherence of our moral beliefs will not advance the discussion. Anyone who reflects on their own productive moral discussions or on a time when they improved their own moral beliefs should find such a picture of moral inquiry attractive.

---

[13] I'm leaving infinitism, externalism and contextualism out of this part of the discussion. They are less popular positions and, more importantly, they don't provide a clear answer to the question of how a belief may be defensible from one's own point of view.

For example, you may have wondered if there is an obligation to donate some of your income to alleviate poverty. It might seem possible to gather all the relevant non-moral facts and still be confused about your obligation.[14] But it's worth considering just what non-moral facts are relevant and what it takes to know them. Do you know exactly how aid organizations utilize your money, precisely how those in need would be affected and how your own life would be impacted? Do you fully appreciate the relevance of these facts? Getting a clear and comprehensive understanding of the non-moral facts is a large and difficult task – and though doubt and disagreement may remain even after fully grasping (or simply stipulating) the non-moral facts, the most we can do to resolve these doubts is to redescribe, reexamine, reassess and otherwise improve upon our understanding of the non-moral facts.

Finally, moral inferentialism has implications for the aims of moral philosophy. Moral philosophy is not about finding self-evident moral principles that could explain (or be explained by) the rest of the moral facts. Nor is it about finding principles which systematize the moral propositions we would believe if our thinking were ideally coherent. Instead, moral philosophy is an investigation into the ways in which non-moral facts explain moral ones and the ways non-moral beliefs justify moral ones. It is an attempt to understand the inferential patterns characteristic of good moral thinking. And that component of moral philosophy called moral theory – the part which aims to justify our most general moral beliefs – requires highly general and comprehensive non-moral knowledge. Moral theory is hard – much more difficult than everyday moral reasoning.

If these implications seem plausible or intriguing then the view that non-moral beliefs by themselves are capable of justifying moral beliefs should receive more attention than it has traditionally enjoyed. At the

---

[14] Thanks to an anonymous referee for suggesting this example.

very least, it should not be summarily dismissed by remarks about the impossibility of 'deriving' an 'ought' from an 'is.'[15]

References

Annis, D. B. 1978. A Contextualist Theory of Epistemic Justification, *American Philosophical Quarterly* 15/3: 213–219.

Anscombe, G. E. M. 1958a. Modern Moral Philosophy, *Philosophy* 33/124: 1–19.

Anscombe, G. E. M. 1958b. On Brute Facts, *Analysis* 18/3: 69–72.

Audi, R. 2005. *The Good in the Right: A Theory of Intuition and Intrinsic Value*, Princeton: Princeton University Press.

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

Brandom, R. 2000. *Articulating Reasons: An Introduction to Inferentialism*, Cambridge: Harvard University Press.

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

Brink, D. O. 1999. Objectivity and Dialectical Methods in Ethics, *Inquiry* 42/2: 195–212.

Brown, C. 2014. Minding the Is-ought Gap, *Journal of Philosophical Logic* 43/1: 53–69.

Carroll, L. 1895. What the Tortoise Said to Achilles, *Mind* 4/14: 278–280.

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

Frankena, W. K. 1973. *Ethics*, Englewood Cliffs, N.J.: Prentice-Hall.

Gettier, E. L. 1963. Is Justified True Belief Knowledge?, *Analysis 23*/6: 121-123.

Horty, J. F. 2012. *Reasons as Defaults*, New York: Oxford University Press.

Huemer, M. 2005. *Ethical Intuitionism*, Houndmills: Palgrave Macmillan.

Jackson, F. 1974. Defining the Autonomy of Ethics, *The Philosophical Review* 83/1: 88–96.

Klein, P. 2007. Human Knowledge and The Infinite Progress of Reasoning, *Philosophical Studies* 134/1: 1–17.

Kurtzman, D. R. 1970. " Is," "Ought," and the Autonomy of Ethics, *The Philosophical Review* 79/4: 493–509.

Lerner, A. and Leslie, S. J. 2013. Generics, generalism, and reflective equilibrium: Implications for moral theorizing from the study of language, *Philosophical Perspectives* 27/1: 366-403.

Leslie, S. J. 2007. Generics and the Structure of the Mind, *Philosophical Perspectives* 21/1: 375-403.

McBrayer, J. P. 2010a. A Limited Defense of Moral Perception, *Philosophical Studies*, 149/3: 305–320.

McBrayer, J. P. 2010b. Moral Perception and the Causal Objection, *Ratio*, 23/3: 291–307.

McGrath, S. 2004. Moral Knowledge by Perception, *Philosophical Perspectives* 18/1: 209–228.

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

Nelson, M. 2003. Who Needs Valid Moral Arguments? *Argumentation* 17/1: 35-42.

Prior, A. N. 1960. The Autonomy of Ethics, *Australasian Journal of Philosophy* 38/3: 199–206.

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

Sayre-McCord, G. 1996. Coherentist Epistemology and Moral Theory, in *Moral Knowledge? New Readings in Moral Epistemology*, ed. Walter Sinnott-Armstrong and Mark Timmons, New York: Oxford University Press: 137-189

Searle, J. R. 1964. How to Derive "Ought" From "Is", *The Philosophical Review* 73/1: 43–58.

Setiya, K. 2012. *Knowing Right from Wrong*, Oxford: Oxford University Press.

Shafer-Landau, R. 2003. *Moral Realism: A Defence*, New York: Oxford University Press.

Sidgwick, H. 1879. The Establishment of Ethical First Principles, *Mind* 4/13: 106–111.

Sinnott-Armstrong, W. 2007. *Moral Skepticism*, New York: Oxford University Press.

Thomson, J. J. and Harman, G. 1996. *Moral Relativism and Moral Objectivity*, Malden, MA: Blackwell.

Timmons, M. 1998. *Morality Without Foundations: A Defense of Ethical Contextualism*, New York: Oxford University Press.

Yamada, M. 2011. Getting It Right By Accident, *Philosophy and Phenomenological Research* 83/1: 72-105.

Wedgwood, R. 2001. Conceptual Role Semantics for Moral Terms, *The Philosophical Review* 110/1: 1-30.

Wellman, C. 1971. *Challenge and Response*, Carbondale: Southern Illinois University Press.

Zangwill, N. 2006. Moral Epistemology and the Because Constraint, in *Contemporary Debates in Moral Theory*, ed. Jamie Dreier, Malden, MA: Blackwell.

Zimmerman, A. 2010. *Moral Epistemology*, New York: Routledge.