

TWO DILEMMAS FOR SETIYA'S REDUCTIVE EPISTEMOLOGY

Shlomit WYGODA COHEN

ABSTRACT: Kieran Setiya has developed an original and promising version of non-inferential moral epistemology which he calls Reductive Epistemology. While I find Reductive Epistemology to be an attempt in the right direction, it leaves some important questions unanswered. Tackling these questions brings to the surface the intuition underlying the view's plausibility, but also reveals a significant challenge in cashing out this underlying intuition.

KEYWORDS: non-inferential moral epistemology, reductive epistemology,
Kieran Setiya

1. Introduction

Last week, my sister asked me to help her with an errand. However, since I had more pressing things to do, I lied to her and said I was ill. Assuming it was wrong of me to lie to my sister in these circumstances, how is it that we come to know this moral truth? According to an initially plausible answer, we know moral principles (perhaps by intuition), and we can infer particular moral truths from one (or more) of these principles, supplemented by the relevant particular non-moral facts. In this instance, lying is wrong; I lied to my sister; therefore, what I did was wrong. Despite its initial appeal, this deductive model faces serious difficulties.¹ In view of these difficulties, many authors have lately engaged in defending the view that particular moral truths can be, and often are, known non-inferentially (that is without depending on any other belief for their justification);² One promising version of non-inferential moral epistemology is Reductive Epistemology, a view developed by Kieran Setiya.³ Although Reductive Epistemology is not widely held, its wider intuitive appeal makes the view particularly worthy of investigation. In a nutshell, I believe that this appeal is due to the underlying insight that evidence for that in virtue of which a moral fact obtains can be sufficient for justifiably believing that this moral fact

¹ See, for instance, (Harman 1977, 4), (McGrath 2018, 200-206), and (Setiya 2012, 44-46).

² (McGrath 2018, fn. 13). For a slightly different understanding of non-inferentiality see (Pryor 2005).

³ (Setiya 2012).

obtains. Any attempt to cash out this intuition has to address the challenge discussed in the paper.

I start by presenting Setiya's Reductive Epistemology raising a couple of questions it leaves unanswered (sec. 2). One of these questions pertains to an ambiguity of the view, and so I follow by discussing how to disambiguate it in the most charitable way (sec. 3). This puts me in a position to examine the other unanswered question which poses a more significant challenge to the view (sec. 4). I conclude by exploring alternative strategies available for moving forward given the gathered insight (sec. 5).

2. Setiya's Reductive Epistemology

Kieran Setiya has developed a version of non-inferential moral epistemology which he calls Reductive Epistemology.⁴ Setiya first introduces the following metaphysical principle:

ETHICAL SUPERVENIENCE: If an act or agent falls under ethical concept, E, it does so in virtue of falling under non-ethical concepts, N, such that necessarily, what falls under N falls under E.⁵

Setiya formulates the principle in terms of concepts rather than properties.⁶ However, it is sometimes more natural to speak in terms of property instantiation (i.e., facts) rather than concept application (i.e., propositions).⁷ Thus, according to Ethical Supervenience if an act or agent instantiates some ethical property E, it does so in virtue of instantiating a non-ethical (probably complex) property, N, such that necessarily, what instantiates property N instantiates property E. With this conditional in place, Setiya then asks:

What is the evidence by which I am justified in believing that an act is right or wrong, an agent generous or unjust? It is evidence that the act or agent falls under non-ethical concepts, N, where, necessarily, what falls under N is right or wrong,

⁴ Setiya's main motivation in developing his view is to address the sceptic about the possibility of being justified in moral beliefs. Particularly, his view is meant to have the advantage of providing a satisfactory answer to the sceptical argument from fundamental disagreement. This will not be my focus here.

⁵ (Setiya *ibid*, 8-11, 49).

⁶ He does so in order to accommodate the reductive naturalist who holds that any ethical property E is a non-ethical property N. (Setiya *ibid*, 9). See also (Sturgeon 2009).

⁷ I don't take this to be consequential. First, Setiya himself when presenting his view sometimes speaks in terms of facts rather than concepts. See for instance (Setiya *ibid*, 40, 54). Furthermore, this reformulation only serves for ease of presentation, and the subsequent argument could also be made without it.

generous or unjust. Call this Reductive Epistemology, since it reduces evidence in ethics to evidence elsewhere.⁸

According to Setiya, then, evidence that something falls under a non-moral concept *N* is evidence that it falls under a moral concept *E*. Hence, the proposition *x is E* is propositionally justified, that is, it is supported by the evidence for the proposition *x is N*.

However, that *x is E* is propositionally justified (i.e., that the evidence supports it), does not yet imply that the proposition is doxastically justified (i.e., that one who has the relevant evidence believes it with justification). When does this evidence play its doxastic justifying role? How are we ever justified in moving from this evidence to a moral belief? Here Setiya claims that in order to be justified in believing *x is E* on the basis of evidence that *x is N*, it is sufficient to non-accidentally “manifest an epistemically reliable disposition” – a disposition that non-accidentally “tracks, at least roughly, the conditionals involved in Ethical Supervenience”.⁹

Having the relevant disposition does not imply that given the evidence we always get things right. We may still on occasion make the wrong judgments or no judgment at all. Manifesting the disposition provides doxastic justification because of something about the way in which the belief is reached, not merely because it results in true belief. Thus, this disposition is a kind of epistemic virtue or competence, a disposition “[...] that would, in appropriately normal conditions, ensure (or make highly likely) the success of any relevant performance issued by it”.¹⁰

It is important to emphasize that Reductive Epistemology does not require that we *know or believe* the relevant Ethical Supervenience conditional but only that we *reliably track* it. In this manner, Reductive Epistemology eschews a problem Setiya presents in criticizing the deductive model:

Even detailed accounts of the circumstance of action, non-ethically described, leave room for countervailing reasons, facts consistent with the circumstance that would make an ethical difference. If I think you acted wrongly in hitting that child, there are countless ways the world could be, most of them unimagined by me, that would justify what you did. Perhaps you hit her in order to kill the lethal insect on her arm. Perhaps you acted under threat. Perhaps she will feel no pain. While I readily cite non-ethical facts as grounds for my belief—the child is young; you hit her hard—they fail to necessitate its truth. [...]. There are too many loopholes to

⁸ Setiya *ibid*, 49.

⁹ Setiya *ibid*, 65.

¹⁰ Sosa 2007, 29.

close.¹¹

Hence, in order to close all the possible loopholes and necessitate E, N would have to be an extremely long conjunction, too long to be psychologically graspable. This is a problem, however, only if our theory requires that we know or believe that N or the relevant Ethical Supervenience conditional obtain, which Reductive Epistemology does not.

However, there remains a difficulty. Suppose that some act is wrong, and suppose that the wrongness of the act holds in virtue of and is necessitated by a set of non-moral facts N. In order to be justified in believing the action is wrong, does Setiya's view require us to be in the relevant epistemic position for that full set or is it sufficient that we be in the relevant epistemic position only for a subset of N? If the former – is it plausible that we typically can be in the relevant epistemic position for the full set N? And if the latter – which subset and why is it sufficient?

Before we can address this challenge, there is a preliminary issue that must be addressed. Namely, what exactly is included in the set N? As I argue in the following section, there are two possible interpretations of the Ethical Supervenience conditional each of which yields a somewhat different answer on this issue.

3. The ambiguity of the Ethical Supervenience conditional

There are two possible readings of the Ethical Supervenience conditional, and throughout his book it is not altogether clear which one Setiya has in mind: supervenience or grounding. Roughly speaking, supervenience is necessary covariation, while grounding is the notion used to pick out the non-causal because or in-virtue-of relation or operation.¹² On the one hand, the name of the principle, and the insistence on its being uncontroversial point to the supervenience interpretation. On the other hand, the locution 'in virtue of' seems to indicate that what Setiya actually has in mind is grounding. So, which of these is the more charitable interpretation of Setiya's view?

Setiya understands evidence "as a set of propositions or attitudes to which one bears some special relation, and to which one's degree of belief should be proportioned".¹³ In the non-moral case, it seems plausible that evidence for what grounds something is at least some evidence for it. Evidence that C-fibers are firing in the patient's brain, for instance, is at least some evidence that they are in pain.

¹¹ (Setiya *ibid*, 44)

¹² Dancy (1981) was the first to notice this important distinction. See (Berker 2017) for a thorough history of this distinction.

¹³ (Setiya *ibid*, 41).

But it is far less plausible that when A supervenes on B, evidence for B is evidence for A. Evidence that the sky is blue is not evidence (i.e., does not provide propositional justification) for or against Goldbach's conjecture even though the mathematical supervenes on the physical. Setiya's Reductive Epistemology is therefore far more attractive if we understand the reduced evidence for the moral claim as evidence for that in virtue of which the moral fact holds, i.e., its grounds.¹⁴

Grounding theorists typically distinguish between full and partial grounds. Roughly, A full ground of p is a set of facts that collectively (non-causally) make it the case that p. Any subset of this set is a partial ground of p.¹⁵ So, we might ask, is the grounding claim of the supervenience conditional a full or partial grounding claim. However, in this respect the Ethical Supervenience conditional is hardly ambiguous since the grounding involved in it is one that necessitates and this is certainly not the case for partial grounds of particular moral facts.¹⁶ Hence, Setiya's Ethical Supervenience conditional, is to be understood not merely as a partial grounding claim, but as a *full* grounding claim. There is some intuitive appeal to the idea that sufficient evidence for the full ground of E is sufficient for justification of E: given sufficient evidence for the full ground of E, there isn't any additional evidence out there that provides better insight into that in virtue of which E obtains.

It is worth noticing that nobody denies that evidence for at least certain parts of the grounds of a particular moral fact provides some evidence for the proposition that this fact holds. For example, even on the deductive model evidence that the action caused unnecessary pain is *some* evidence that the action was wrong. So, if Setiya's Reductive Epistemology is non-trivial and original at all, it should be understood not merely as the claim that evidence for what fully grounds a moral fact provides *some* evidence for the proposition that this fact holds, but rather as the claim that such evidence can be *sufficient* to (non-inferentially) propositionally justify belief. Notice, that Setiya's claim then is not merely that such evidence is sufficient qua *evidence* but another sort of justification is also necessary (this is compatible with the deductive model). Setiya's claim is stronger, it is that such evidence can in itself be sufficient for *propositional justification*.

Thus, according to the more charitable interpretation of Ethical Supervenience N *fully* grounds and necessitates E. Still, it is worth noting that on this interpretation one supposed advantage of the view gets lost. If understood as a

¹⁴ This relative advantage remains although it is not the case that evidence for the grounds of a moral facts *always* seems to be evidence for that moral fact.

¹⁵ (Rosen 2010), (Fine 2012), (Correia and Schnieder 2012, 20–21).

¹⁶ A partial ground can, in some case, necessitate. For example, the fact that it's raining in London and π is irrational is partially grounded, and necessitated, by the fact that it's raining in London.

supervenience claim, the Ethical Supervenience conditional is uncontroversial and unproblematic. But if understood as a grounding claim the conditional is far from being undisputed. Though it is quite uncontroversial that particular moral facts are not free floating—in the sense that they must be grounded in or hold in virtue of something—this Ethical Supervenience conditional rests on a couple of controversial assumptions. First, the assumption that ethical properties are fully grounded in non-ethical properties, an assumption typically denied by non-naturalists.¹⁷ And second the assumption that what fully grounds an ethical fact also necessitates it, an assumption rejected by contingentists.¹⁸

In sum, the interpretative question can be presented as a dilemma. If we understand the Ethical Supervenience conditional as a supervenience claim, then it is indeed uncontroversial as desired, but then reductive epistemology does not seem very plausible. On the other hand, if we understand the Ethical Supervenience conditional as a grounding claim then reductive epistemology seems plausible but the conditional itself is contested. The best avenue for friends of Reductive Epistemology is to accept the second horn and admit that this kind of moral epistemology is not available to all metaethical positions. Reductive Epistemology is more limited than some may have hoped. That being said, the challenge discussed in the next section can also be raised (more forcibly even) against those who prefer the first horn.

4. The Challenge

Equipped with an answer to the question of what the set N refers to, we can now consider whether Setiya's view requires us to be in the relevant epistemic position for that full set. This difficulty too can be stated in the form of a dilemma. Evidence for how much and which part of the non-moral base (i.e. of the full ground) is required for there to be sufficient evidence for the moral proposition E? On the one hand, if the answer is that we need evidence for all of it, this leads to moral skepticism. On the other hand, if the answer is that evidence for some of it is sufficient, the proponent of Reductive Epistemology has yet to explain which part and why. As far as I can see, there is no middle ground.

On the first horn of the dilemma, then, in order to justifiably believe E on the basis of N what is required is evidence for all of N. However, I will argue that this is implausible. First, it seems that we don't need to have evidence for each and every feature of the situation that circuitously bears on some possible loophole, in order to

¹⁷ (Enoch 2019)

¹⁸ (Leuenberger 2014), (Skiles 2015).

justifiably believe some moral fact. But more to the point, it seems that this requirement is so epistemically demanding that embracing it would lead to moral skepticism.

Suppose that what is required is evidence for all of the non-moral base. In the example of the child cited above, there seems to be a straightforward way to argue that there is nothing too epistemically demanding in the requirement that we have evidence for N. When I have evidence that the adult struck the child, I am in a position to justifiably believe the negation of various loopholes; I'd be justified, for instance, in taking certain bets that it was not meant to kill a lethal insect. This is presumably because, usually or under normal conditions, when someone hits a child, it is not in order to kill a lethal insect. In other words, in most cases or under normal conditions, where this evidence obtains, N obtains. If this is right, then it is plausible that I must either have evidence for these propositions or some sort of default warrant for believing them. Either way, I have a set of propositions or attitudes to which my degree of belief that N should be proportioned, i.e., evidence that N.

The reason we are justified in believing N in this example, is that (supposedly) regarding all of the aspects of the situation that could bear on the moral value of hitting the child, there is a way the world normally is. However, not all cases are like the child hitting case. Consider the following case:

If I think you acted wrongly in lying to your sister, there are countless ways the world could be, most of them unimagined by me, that would justify what you did. Perhaps you lied to her in order to protect her. Perhaps you have a running agreement that you should lie to each other in these circumstances, perhaps you were playing a lying game, perhaps you know the truth would hurt her more than anything, etc., etc. There are too many loopholes to close.

Here, it doesn't seem that when I have evidence that you lied, I am in a position to justifiably believe a single set of facts N that necessitate that you acted wrongly. That is to say, there is no such set N that is also such that given the evidence that you lied I either have evidence for it or some sort of default warrant for believing it obtains. Evidence that you told a lie cannot do this because there is no single set of facts that always obtain under normal conditions of lying: Sometimes people lie in order to protect loved ones and sometimes they do not. Sometimes people lie because of previous agreements and sometimes there is no such agreement. And usually, the situation is more complex involving various features particular to the relationship its history and significance that bear in a unique way on the propriety of lying. Evidence that you told a lie cannot settle anything regarding those other relevant facts.

Similarly, it doesn't seem that when I have evidence that you lied, it is more likely than not that none of the possible disabling conditions obtains. There is no reason to suppose that the various possible ways the world could be such that your lying is not wrong-making, are any less frequent than those in which it is wrong-making (which is not to say that it is more frequent), and there is no morally neutral reason (although there might be a moral one) to suppose that the circumstances under which lying is not wrong-making are not normal.

One might respond that where the evidence doesn't settle whether defeaters obtain, so that we are not justified in believing the full and necessitating non-moral ground, we are likewise not justified in believing that the action was wrong on this evidence (at least not in the way Reductive Epistemology suggests). My response is that according to our practices of moral reasoning, we are at least on occasion justified (albeit defeasibly) in our moral judgments about certain instances of lying even though it is doubtful that we can ever be in a position to justifiably believe the negation of the various loophole. It seems that I can be justified (though defeasibly) in believing that you acted wrongly in lying to your sister without being privy to the intricacies of your relationship and thus without being able to settle whether certain defeaters obtain. And, again, there is no reason to think that these defeaters are rare or unlikely to occur. Moreover, this is not just a matter of our practices; upon reflection, this aspect of our practices seems appropriate. Even though I can go wrong precisely because I am not aware of every relevant detail of the situation, I am justified in making a moral judgment regarding the propriety of your lying. And generally speaking, it is a weakness for a moral epistemology to require a radical revision of our appropriate practices. The challenge for Setiya's moral epistemology is therefore to explain how we can be justified in such cases. If, as seems to me to be the case, most of our moral judgments are more like the lying case than the child hitting case, this poses an even more significant challenge for Reductive Epistemology. Though it is plausible that evidence for that in virtue of which you did something wrong in lying to your sister (N), can be sufficient for justifiably believing that you did something wrong, Setiya's picture as it stands needs to explain how we can have evidence for N in such cases. I believe the prospects for meeting this skeptical challenge are bleak.

This brings us to the second horn of the dilemma. Perhaps what is required according to Reductive Epistemology for sufficient justification is not evidence for N but evidence for some subset of N. Here, however, a different kind of challenge lurks. While it seems appropriate to form the (defeasible) belief that your action was wrong solely on the evidence that you hit the child, it seems misguided to form this belief solely on the basis of evidence that there are no lethal insects around. To drive

the point home, even if, statistically speaking, or under normal conditions, you tend to act wrongly when it snows outside, evidence that it is snowing outside is just not of the right kind to count on its own as even some evidence that you acted wrongly. Evidence for N is evidence for E because this evidence gives insight into that in virtue of which E holds, not merely because E is more likely given this evidence. This is what makes Setiya's Reductive Epistemology plausible in the first place. Evidence that there are no poisonous insects around does not seem to be, by itself (without evidence that you hit the child or evidence for relevant statistics), even *some* evidence that you acted wrongly. What Setiya is missing here is a way of distinguishing among the non-normative facts that necessitate some normative fact E, those that are such that evidence for them can constitute sufficient evidence for E, and those that are such that evidence for them cannot constitute sufficient evidence for E.¹⁹

Admittedly, there are many possible ways to specify a subset of the relevant non-moral facts and there is nothing in Setiya's view to rule them out. However, within Reductive Epistemology there is no way to justify appeal to any such subset. Supposedly, N grounds and necessitates E and so evidence for N is evidence for E. But a subset of N, even a relevant subset of N, does not ground and necessitate E and so some alternative rationale needs to be presented for why evidence for this relevant subset is sufficient evidence for E. It is not clear how Setiya can – given his theoretical apparatus – maintain the initial plausibility of his position while also arguing that evidence for a special subset plays a special epistemic role.

In sum, Reductive Epistemology faces the following dilemma: on the one hand if the evidence required is evidence for N as a whole, Reductive Epistemology turns out to be too demanding in many important cases, and thus leads to moral skepticism. If, on the other hand, the evidence required is evidence for some subset of N, Reductive Epistemology loses its initial plausibility. The view as it is presented by Setiya, lacks the tools to point to a privileged subset of N that is such that evidence for it is manageable and could plausibly constitute sufficient evidence for E.

5. Concluding Remarks

From what I can gather, there are two possible approaches or strategies for addressing this challenge. The first strategy is to narrow the scope of Reductive Epistemology and claim that the kind of moral facts that Reductive Epistemology

¹⁹ The structure of my argument here is somewhat analogous to Dancy's (2004) objection to Jackson's Supervenience argument in that it builds on the observation that the Supervenience tool is too rough to pick out the relevant evidence ("constitutive reasons" in Dancy's terminology).

discusses are those with merely *prima facie*²⁰ moral status. Lying, for instance, fully grounds and necessitates *prima facie* moral wrongness. Thus, evidence that you lied is sufficient evidence that the act is *prima facie* wrong, i.e., that it is wrong unless some "loophole" is present. And so, there is no need for evidence for the absence of all possible loopholes after all. However, this renders Reductive Epistemology quite underwhelming since it leaves unanswered one of the most important parts of moral epistemology – how we are justified in believing particular *unqualified* moral truths. This would require some sort of weighing procedure of *prima facie* duties, that is some sort of way of telling based on the *prima facie* duties what one ought to do overall. So long as this piece of the puzzle is missing, our moral epistemology is lacking. Importantly, this lack is not a marginal aspect of moral epistemology. Although moral epistemology is conducted by philosophers, one of its most important goals is to explain how normal people can be justified in their beliefs about how they ought to behave. Thus, even though this strategy technically avoids the dilemma, it leads to skepticism about the possibility of significant moral knowledge which is precisely the trouble with the first horn of the dilemma. I therefore don't think this is a very promising avenue.

A more ambitious strategy requires significantly amending Reductive Epistemology. This can be done for example by revising Ethical Supervenience so as not to require necessitation. This would point to a subset of N that is privileged in a relevant way, and that is such that there is no difficulty in explaining how we have evidence for it. The challenge would then be to explain why we should think that evidence for this privileged subset can be sufficient to propositionally justify that an action is wrong without first settling whether the loopholes obtain.²¹ Even if such a view would no longer count as a version of Reductive Epistemology, it would nonetheless retain important elements of Setiya's original approach.²²

References

Berker, S. (2017). The unity of grounding. *Mind*, 127(507), 729-777.
doi:10.1093/mind/fzw069

²⁰ In the sense in which W.D Ross uses this term with regard to moral duties.

²¹ For a view along these lines see (Wygoda Cohen 2025).

²² For helpful conversations and comments on this and previous versions, I thank Dan Baras, Aaron Elliott, David Enoch, Noga Gratvol, Avi Kenan, Bar Luzon, Oded Naaman, Peter Shiu-Hwa Tsu, and Preston Werner. The ideas discussed in this paper were initially discussed in a presentation I made at the 2019 annual conference of The Centre for Moral and Political Philosophy at the Hebrew University of Jerusalem. I thank the conference participants for their helpful comments.

- Correia, F. and Schnieder, B., 2012. Grounding: An opinionated introduction. In: F. Correia and B. Schnieder eds., *Metaphysical grounding: Understanding the structure of reality*. Cambridge University Press, pp.1–36.
- Dancy, J. (1981). On Moral Properties. *Mind*, 90(359), 367–385.
- Dancy, J. (2004). Discussion on the importance of making things right. *Ratio*, 17(2), 229–237. doi: 10.1111/j.1467-9329.2004.00250.x
- Enoch, D. (2019). How Principles Ground. In *Oxford Studies in Metaethics Volume 14*. Oxford University Press.
- Fine, K., 2012. Guide to ground. In: F. Correia and B. Schnieder eds., *Metaphysical grounding: Understanding the structure of reality*. Cambridge University Press, pp.37–80.
- Harman, G. (1977). *The Nature of Morality*. Oxford: Oxford University Press.
- Leuenberger, S. (2014). Grounding and necessity. *Inquiry*, 57(2), 151–174.
- McGrath, S. (2018). Moral Perception and Its Rivals. In: A. Bergqvist and R. Cowan (Eds.), *Evaluative Perception*. Oxford University Press, 161–182
- Pryor, J. (2005). There is immediate justification (pp. 181–201) in M. Steup, J. Turri, and E. Sosa, (Eds). *Contemporary debates in epistemology*. John Wiley & Sons.
- Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale, and A. Hoffmann (Eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press, 36–109.
- Ross, W. D. (1930) *The Right and the Good*. Oxford: Clarendon Press.
- Setiya, K. (2012). *Knowing right from wrong*. Oxford University Press.
- Sosa, E. (2007). *A virtue Epistemology: Apt belief and reflective knowledge*, (Vol.1). Oxford University Press.
- Skiles, A. (2015). Against grounding necessitarianism. *Erkenntnis*, 80(4), 717–751.
- Sturgeon, N. (2009). Doubts about the Supervenience of the Evaluative. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics*, 4, 53–90.
- Wygod Cohen, S. (2025). Making Moral Judgments on Adequate Grounds: From Transparent Evidence to Justified Moral Belief. *Ratio*, 38(3), 147–156.