

Moral perception, inference, and intuition

Daniel Wodak¹

Published online: 23 January 2019
© Springer Nature B.V. 2019

Abstract Sarah McGrath argues that moral perception has an advantage over its rivals in its ability to explain ordinary moral knowledge. I disagree. After clarifying what the moral perceptualist is and is not committed to, I argue that rival views are both more numerous and more plausible than McGrath suggests: specifically, I argue that (a) inferentialism can be defended against McGrath’s objections; (b) if her arguments against inferentialism succeed, we should accept a different rival that she neglects, intuitionism; and (c), reductive epistemologists can appeal to non-naturalist commitments to avoid McGrath’s counterexamples.

Keywords Moral epistemology · Moral perception · Inferentialism · Intuition · Reductive epistemology

1 Introduction

Sarah McGrath’s ‘Moral Perception and Its Rivals’ argues for the modest but controversial claim that some of our moral knowledge is perceptual.¹ An important and welcome feature of McGrath’s approach is that it explores a familiar issue in moral epistemology via considering simple, mundane cases. Her central example involves no esoteric moral propositions or improbable stipulations about extraordinary circumstances. It involves ordinary agents witnessing ordinary instances of moral wrongdoing: ‘While walking through a grocery store parking lot, you witness

¹ See McGrath (2018). All parenthetical page references in the main text are to this manuscript unless specified otherwise.

✉ Daniel Wodak
dwodak@vt.edu

¹ Virginia Tech, Blacksburg, USA

a frustrated adult lash out and strike a child across the face, causing the child to draw back in pain and surprise' (1). Let's assume, following McGrath, that you know that the adult acted wrongly. The central question is: *how* do you know this? And McGrath's thesis is that your knowledge that the adult acted wrongly is perceptual knowledge.

Another important feature of McGrath's approach is that she defends her thesis by eliminating rival hypotheses about how we have moral knowledge in such ordinary cases. This flips the typical dialectic (2). Instead of playing defense against critics of moral perception, McGrath goes on the offensive, arguing that alternative views are implausible.

I am sympathetic to the view that some of our moral knowledge is perceptual. But I contend that rival hypotheses can plausibly explain how we have moral knowledge in ordinary cases. To show why, I will take on three main tasks in this paper. The first will be to defend inferentialism against McGrath's objections: her arguments against toy versions of inferentialism overgeneralize, and leave sophisticated versions of the view unscathed. The second will be to argue that if her arguments against inferentialism succeed, we should simply accept intuitionism: this rival hypothesis, which McGrath ignores, has a much easier time explaining moral knowledge about unperceived thought experiments. The third will be to argue that reductive epistemologists can draw on recent work on metaethical non-naturalism to avoid McGrath's counterexamples.

2 The perceptualist

Before we consider rival hypotheses, let's start by outlining the perceptualist explanation of ordinary moral knowledge. In particular, let's ask: Must perceptualists about morality take on implausible commitments about the nature and epistemological import of perception generally?

According to McGrath, the answer is clearly no. To show why, she offers what she calls a 'relaxed' account of perceptualist moral knowledge (21 ff.), which allows the perceptualist to say that we can have perceptual knowledge that the adult's act is wrong without necessarily having a perceptual experience with the content 'the adult's act is wrong' (24).

To appreciate the significance of this, consider *austere* views about the contents of perceptual experience, according to which we see colors and shapes and not much else. If we adopt such a view, our visual experience of the parking lot scene will not represent moral wrongness; but nor will our visual experience represent various non-moral objects, properties and relations: the adult, the child, the frustration, the causation, or the pain. Proponents of austere views can and should still accept that we can know, via perception, that the frustrated adult struck the child, causing pain. Likewise, they can accept that we know that the adult's act was wrong via perception, even though 'wrongness' was not represented in perception. Instead, as McGrath writes, it might be that 'the features of the scene that you do take in in your visual experience *trigger* or *prompt* you to take up the immediate, non-inferential belief' that the adult's act was wrong (24).

As McGrath notes, this way of defending moral perception is not novel.² And in fact, the same point has been made in Margot Strohminger's recent defense of the highly counter-orthodox view that we can have perceptual knowledge of metaphysical possibilities.³ I will assume that one's views about moral and metaphysical perception should be closely aligned. (The contrary assumption seems implausible: why think that the commitments and merits of perceptualism will differ so wildly between these domains?) *Prima facie*, this should be welcome: if metaphysical perception is defensible, and defensible on similar terms, this is good news for moral perception. Though it shall also place some constraints on McGrath's options in response to one problem we will encounter in Sect. 4.

But let's not get ahead of ourselves. One curious feature of McGrath and Strohminger's views is that while they argue that perceptualists need not take on contentious commitments about the *contents* of perceptions, both also take on a strong commitment about the nature of perception: namely, both commit to the view that perception is non-inferential.⁴ As we saw above, on McGrath's view 'your visual experience *trigger[s]* or *prompt[s]* you to take up the immediate, non-inferential belief'. This is a common albeit contentious view: many, such as Jerry Fodor, reject it.⁵ Shouldn't we prefer a form of moral perceptualism that is more ecumenical?

It may seem that this commitment is unavoidable for perceptualists, so they must inevitably take up arms against views like Fodor's. After all, if perception is inferential, whither the disagreement with inferentialism?

However, the perceptualist need not adopt this commitment. To see why, consider what we should say about the epistemology of perception if we adopt the view that perception is an inferential process. Perhaps your perceptual belief that the adult intentionally struck the child rests on sub-conscious inferences from complex conditional propositions. This is a claim about the *etiology* of perception. From this claim about its etiology, nothing follows whatsoever about the *epistemology* of perception.

An inferentialist epistemology of perception is one on which justification or knowledge is transmitted from premises to the concluding perceptual belief.⁶ This view is not forced on us if we believe that perception is inferential. You may think that even if your perceptual belief that the adult intentionally struck the child was

² See McGrath (2018, footnote 16), and references therein.

³ Strohminger (2015, especially at p. 321) ('I am not claiming that sensory experiences themselves ever have modal contents. ... Arguably, perceptual knowledge is possible even when the content known is not represented in sensory experience').

⁴ *Ibid.*, p. 365: citing Miller (perceptual knowledge 'arises immediately from current perception, that is, without inference from prior assumptions').

⁵ See, e.g., Fodor (1983), and for very helpful discussion see Firestone and Scholl (2016). This etiological point is distinct from the epistemological claim [raised, most pertinently here, by Faraci (2015)], that moral knowledge by perception must rest on non-perceptual moral knowledge, so perceptualism cannot explain how we have moral knowledge at all.

⁶ This may not be the happiest way of drawing the distinction between perceptualism and inferentialism. Some philosophers argue that you can gain knowledge via inferences from premises that are false and hence unknown. See Warfield (2005) and Fitelson (2010).

caused by a sub-conscious inference from complex premises, you do not know that these premises are true, or at least, you do not know these premises as well as you know the conclusion. That is, we can have a perceptualist epistemology of perception while accepting an inferentialist account of the nature of perception. So, we can preserve a meaningful disagreement between perceptualist and inferentialist epistemologies about morality or modality.

This helps us clarify what perceptualism is and is not a view about. It is first and foremost a view about epistemology: about justification, knowledge, and so on. And it is not a contentious view about the nature of perception or perceptual contents. If this is right, what could or should settle the disagreement between perceptualists and inferentialists?

3 The inferentialist

McGrath's main argument for perceptualism is presented as follows (22):

(P1) You know that the adult acted wrongly.

(P2) If you know that the adult acted wrongly, then you know this either on the basis of perception or on the basis of inference.

(P3) You do not know that the adult acted wrongly on the basis of inference.

(C) Therefore, you know that the adult acted wrongly on the basis of perception.⁷

Moral error theorists, moral skeptics, and (other) moral monsters might deny (P1), but let's ignore them here. We will consider (P2) in the next section. Let's focus on (P3) here. Why does McGrath reject inferentialism?

Following McGrath, let's start with a toy version of inferentialism. According to 'the Deductive Model', you know that what the adult did was wrong because you infer this from two known premises: (i) she ϕ d; (ii) if she ϕ d, she acted wrongly; so, you infer that she acted wrongly. Crucially, on this view the first premise is known via perception, but the second conditional moral proposition is not known via perception.

McGrath considers three challenges to the Deductive Model. The first is phenomenological: it doesn't seem as though your belief is the upshot of reasoning (4). This is not necessarily probative, however, as proponents of inferentialism understand 'inference' broadly to include unconscious, immediate reasoning to which we lack phenomenological access.⁸

The second is epistemological: the Deductive Model requires that you know a sufficient condition for acting wrongly, and indeed requires that you know this condition as well as you know that what the adult did was wrong. Of course, there may be many sufficient conditions for acting wrongly that you can know as well as

⁷ I have changed McGrath's numbering of these premises for the sake of clarity.

⁸ See, for instance, Väyrynen (2008).

you know that what the adult did was wrong: you may know that if what the adult did was evil then the adult acted wrongly. But that does not help the inferentialist, because now the antecedent of this conditional—which she grants is known perceptually—has moral content. To avoid embracing moral knowledge by perception, the inferentialist must insist that this sufficient condition for acting wrongly (“she ϕ d”) has exclusively non-moral content. And, of course, it must be knowable via perception. Dialectically, it will be hard for the inferentialist to identify a good candidate here. Say she takes the view that we specify the first premise of the inference in sparse vocabulary: in terms of arrangements of colors and shapes that are the contents of perception on austere views. This looks hopeless. What arrangements of shapes and colors form a known sufficient condition for acting wrongly? Things don’t look hopeless if we specify the first premise of the inference—‘she ϕ d’—in the rich terminology (including ‘adult’, ‘child’, ‘causation’, ‘intention’, ‘pain’, ‘frustration’). But McGrath asks, ‘once we accept a view on which the possible content of our perceptual knowledge is rich enough to include things like the fact that someone is in pain ... what principled reason is there to exclude the possibility that some perceptual knowledge could have moral content?’ (6).

McGrath’s third and final objection to the inferentialist is a ‘hybrid’ of the first two challenges (8). The second objection turned on *propositional justification*: what justification you have for believing that the adult acted wrongly. But you can have propositional justification that p and yet your belief that p may be unjustified. Say that the adult in question is black and you are biased against blacks: you might have good reasons to believe that the adult acted wrongly, but you don’t believe that she acted wrongly on the basis of these reasons; you believe she acted wrongly because she’s black. In this case, you would have propositional justification but lack *doxastic justification*, because your belief is not based on what would justify it. With this familiar distinction in mind, consider the Deductive Model once more: for your belief that the adult acted wrongly to be justified, or for it to constitute knowledge, not only must you be justified in believing the premises (i) and (ii), but ‘your justification for believing the premises’ must play ‘the right *psychological role* in your coming to believe that the adult acted wrongly’ (9). Given the considerations about the phenomenology and epistemology described above, it is implausible that this condition will be met. That’s why this third objection is a ‘hybrid’.

To see why it will be hard for proponents of the Deductive Model to explain how we are doxastically justified in believing that what the adult did was wrong, consider what they might say about our justification for believing premise (ii) of the inference sketched above. They might say that as ‘you are a competent user of the concept WRONG, you could at least in principle work out the relevant sufficient condition that is satisfied by the adult’; but unless your belief that what the adult did is wrong is based on the justification afforded by your conceptual competence, which seems unlikely (9), this will secure propositional but not doxastic justification.

This final objection strikes me as the most powerful blow to inferentialism. But it is also a fairly subtle objection. It might help to compare this point to an argument suggested by intuitionists against rivals to their view. The content of this argument might vary, but its structure is roughly this:

P1: S believes that p because it seems to S that p [empirical claim].

P2: If P1 is true, then S's belief is justified only if the fact that it seems to S that p provides justification for S to believe that p .

C: S's belief is justified only if the fact that it seems to S that p provides justification for S to believe that p .

To be clear, this argument is often only suggested by intuitionists.⁹ And for it to be complete we would want to substitute values for the variables in P1, which is meant to be an empirical claim that our beliefs about such-and-such (e.g., that the adult acted wrongly) are based on intuitions. If we want to explain how such beliefs can constitute knowledge, the upshot of this would be that we have to grant that intuitionism is true: the fact that it seems to you that p is a reason to believe that p . The point McGrath is pressing is similar. If your belief that the adult did something wrong is based on your perceptions in the relevant sense, then it can constitute knowledge only if there is perceptual moral knowledge.

We've now considered McGrath's three objections to inferentialism, in order from least to most damning. Do they rule out this rival view?

I don't think so. I will first argue that McGrath's three objections to the Deductive Model overgeneralize, and then I will argue that these objections leave sophisticated versions of inferentialism unscathed. If either of these responses succeed, inferentialism remains defensible.

Why think that these objections to the Deductive Model overgeneralize? Let's consider a case where a simple inferentialist view seems plausible. Consider your knowledge that this sentence are ungrammatical. Let's grant that you know this on the basis of either perception or inference. It is hard to see how you could know that sentences are ungrammatical via perception, in any interesting sense, without taking on very strong commitments about linguistics.¹⁰ So it seems more plausible that you know the sentence is ungrammatical *inferentially*. This means you have to know, and base your belief on, a sufficient condition for a sentence being ungrammatical. But you would probably struggle to specify it accurately.

If you doubt this final claim, consider how hard it is to specify sufficient conditions for subject-verb agreement for sentences involving either conjunctive and disjunctive subjects or subjects that are syntactically plural and semantically singular (e.g., 'You *are* my best friend'). Or just switch to a better example where we know whether particular sentences are grammatical without being able to state the relevant grammatical rule.

Suppose the above is the right thing to say about our knowledge of particular grammatical truths: such knowledge is gained inferentially from applications of

⁹ In particular, by Huemer (2001 pp. 55–57, 2007, pp. 39–41). Huemer is explicitly concerned with the more narrow and defensive argument that objections to intuitionism that depend on intuitions are self-defeating. But he suggests a more general and interesting point in that context when he writes, for instance, that 'I have argued that if [intuitionism] is false, we have no justified beliefs even if we have available justification for various propositions', p. 45. This is the upshot of the schematic argument above.

¹⁰ Why is that? Briefly, because our grammatical knowledge about, say, subject/verb agreement, seems to depend entirely on our a posteriori knowledge of contingent rules.

complex rules of grammar that we cannot specify. If this is right, or even plausible, then there should be a good story about how we can know and base beliefs upon complex rules of grammar that we cannot specify. And whatever that story is, we should be able to wield it out to defend the view that we can gain knowledge of particular moral facts by deductively applying complex moral rules that we cannot specify. That is, we should be able to defend even the toy Deductive Model.

Note that the foregoing appeal to companions in guilt need not involve any contentious commitments about grammatical knowledge. We need not say that our knowledge of the rules of English grammar is innate, for instance. We just need to be able tell a good story about how we can know, apply, and base other beliefs upon, complex rules that we cannot state.

Of course, this appeal to companions in guilt only goes so far. For one, the inferentialist account of grammatical knowledge would need to be motivated. In part, more would need to be said about a widely accepted rival to both inferentialism and perceptualism in this context: that we know particular facts about grammaticality on the basis of linguistic intuitions about (un)acceptability. If that's a viable option, we should not accept the counterpart to McGrath's (P2): If you know that 'This sentence are ungrammatical', you know this either on the basis of perception or on the basis of inference. As we'll see in the next section, the dialectical situation here is very similar with respect to moral perception.

But before we get to moral intuitionism, let's consider another way of defending inferentialist explanations of how you know that what the adult did was wrong. Say that the Deductive Model fails. That would not be surprising, since as Gil Harman (1984) famously argued, deductive syllogisms are a poor model of human reasoning and inferences. So do McGrath's objections target more sophisticated forms of inferentialism?

Let's consider how they might appeal to a view that plausibly matches the actual structure of human inferences: John Horty's 'default reasoning'. The idea behind Horty's view is best conveyed with his main example:

If an agent is told only that Tweety is a bird, it would be natural for that agent to conclude that Tweety is able to fly. Our everyday reasoning seems to be governed by a general default according to which birds, as a rule, are able to fly (Horty 2007, p. 3).

That this rule is a 'general default' means that it admits of exceptions. Tweety might be a penguin or an emu or any other flightless bird. But it is still a good rule all the same. Inferring that Tweety can fly from this rule plus the information that Tweety is a bird seems to be reasonable.¹¹

How might inferentialists apply defaults to explain moral inferences? They might say that 'ϕing is wrong' is a known defeasible generalization, on which your conclusion that what the adult did is wrong is based. Given the description of the case above, there are plenty of good candidates for what that generalization might be: hitting children is wrong, causing others pain out of sheer frustration is wrong,

¹¹ For more on this idea, see Horty (2007, 2012).

and so on. None of these rules, remember, must be exception-less in order for this model to work.

Admittedly, McGrath considers a view that sounds a lot like the above. But she insists that the relevant generalization must be *statistical*: the generalization would not be that ‘birds fly’, but that ‘most birds fly’. This raises two problems for the inferentialist. The first, which she does not mention, is that as McGrath reconstructs the relevant inference, we would no longer get knowledge that what the adult did was wrong. We would get knowledge that ‘(probably) what the adult did was wrong’ (10). The inferentialist should explain how we know more than this hedged claim. The second and more pressing difficulty is that if the generalization is statistical, then we would require an induction from observed instances:

[The premise that ‘most cases in which a child is struck are cases of wrongful action’] is not a potential piece of a priori knowledge: given that at least some strikings of children are not cases of wrongful action, it is an empirical question whether most are. A claim such as “Most strikings of children are cases of wrongful action” is a statistical generalization, and it is known, if it is known at all, in the way that statistical generalizations are (11).

But if we do not know this statistical generalization a priori—if we only know it based on observed instances wrongful strikings of children—this version of inferentialism turns out to depend on there being prior perceptual moral knowledge. In which case, the perceptualist wins.

Inferentialists should avoid this problem by denying that the relevant generalizations must be statistical. The relevant type of generic generalization, they should say, is *characteristic* rather than *majority*: it is in the nature of birds to fly, and in the nature of strikings of children to be wrong.¹² This does not entail that all birds fly or that all strikings of children are wrong. Characteristic property generics admit of exceptions. Nor does it *entail* that the generalization can be known a priori. We might only be able to know that birds fly on the basis of observation. But plausibly, we can know some such *moral* generic generalizations on the basis of a priori reasoning. After all, a great deal of moral philosophy involves arguments for conclusions like ‘Eating meat is wrong’ or ‘Abortion is wrong’. These conclusions are best interpreted as characteristic property generics (it is, allegedly, in the nature of eating meat or having an abortion to be wrong). And yet the relevant arguments are not inductive inferences from observed correlations. They are a priori.

Notice that in the foregoing, I have defended inferentialism by assuming that the first premise of the relevant inference can have rich content (e.g., ‘The adult struck the child’) and yet be known via perception alone. This means that the versions of inferentialism that I have defended are still impaled on the second horn of the dilemma in McGrath’s epistemological objection: if the content of our perceptual knowledge is rich enough to include facts like that, the inferentialist still needs to provide a principled reason to exclude there being perceptual knowledge with moral

¹² See Wodak et al. (2015) and Wodak and Leslie (2017). Here ‘in the nature of’ tracks the psychologist’s (rather than the metaphysician’s) notion of essence: see Leslie (2013).

content. This might undermine inferentialist arguments against perceptualism. But that does not make inferentialism itself an unattractive rival. So it does not help McGrath execute her offensive strategy for defending her thesis.

4 The intuitionist

So much for the inferentialist. Now let's consider a rival to perceptualism that McGrath ignores: moral intuitionism. As I noted, once we recognize that intuitionism is a rival hypothesis for how you know that what the adult did was wrong, we should not accept that you know this either on the basis of perception or on the basis of inference; at least, not without some further argument against the intuitionist explanation.

Intuitionism and perceptualism have a great deal in common, in large part because the intuitionist mantra has long been that whatever we say about the epistemology of *perceptual seemings*, the intuitionist will say the same about *intellectual seemings* (i.e., intuitions). Some intuitionists even think that there is just one foundationalist story about the epistemology of seemings or appearances, and that if one believes that (say) perceptual seemings confer justification on their contents but intellectual seemings don't, then one must identify an epistemically relevant difference between the two.¹³ As I noted above, intuitionists also motivate their views with similar considerations to McGrath's, by appealing to the basing relation.

This will make it much harder for the perceptualist to offer a simple objection to the intuitionist explanation of how you know that what the adult did was wrong. Whatever the perceptualist says about how you know this based on perceptual seemings, the intuitionist can say something similar about how you know this on the basis of intellectual seemings. Indeed, the explanation can be far more direct, because plausibly we can and do have intellectual seemings with moral content. If you see someone strike a child, that will probably *seem* morally wrong.

That's a minor advantage for the intuitionist. But they also have a more significant advantage over perceptualist explanations of ordinary moral knowledge. Consider how much of our moral inquiry is a priori. We do not push fat men in front of trains to see with our eyes whether it was wrong. Instead, we run thought experiments. We consider imagined cases.¹⁴

To dramatize this point, let's return to McGrath's main example. She uses the example in the following way: Suppose you witness the adult strike the child in a parking lot; in that scenario, you would know that this act was wrong. Instead, let's now *imagine* a frustrated adult striking a child in a parking lot. Plausibly, you can actually know right now that what the imagined adult did was wrong. Notice two things about this claim. First, it concerns the actual you reading this paper, not the

¹³ See, e.g., Huemer (2007, p. 32).

¹⁴ This is leaving aside Kagan (2001)'s concern that the cases we imagine may be case types rather than tokens. See also McGrath (2018)'s response to Kagan.

supposed you in the parking lot. Second, it concerns you actually knowing that a hypothetical act is wrong; the wrongness is not actually instantiated.

Now we can ask: what explains how the actual you can know that what the imagined adult did was wrong? The intuitionist can explain this very easily: our intuitions can concern actual and hypothetical cases. It seems to the actual you that what the adult did was wrong, and that's the starting point for explaining how you can know that what the adult did wrong.

What can the perceptualist say in answer to this question? One option is to say that *moral perception* does not explain how you can actually know that what the imagined adult did was wrong. McGrath only claims that there is some perceptual moral knowledge; she does not embrace an outlandish universally quantified claim: that all moral knowledge is perceptual.

If McGrath takes this first option, she will be in good company. In defending modal knowledge by perception, Strohming (2015) writes:

There are many instances of knowledge of nonactual possibilities that cannot be based on sense perception. Thought experiments seem to be capable of yielding modal knowledge. However they do so, it is clear that the knowledge is not perceptual (p. 369).

If this is the right view about modal knowledge in thought experiments, and we want a general view about perceptual *modal and moral* knowledge, then we are under a great deal of pressure to take this first option.

But the first option comes with some costs. We're considering rival explanatory hypotheses. Hypotheses that explain more of the data are, *ceteris paribus*, superior. If the intuitionist's resources can explain more than the perceptualist's—if they explain our moral knowledge of actual and imagined particular cases—intuitionism has at least one theoretical advantage. Moreover, the perceptualist will have to grant that there is some good explanation for how we know that what the imagined adult did was wrong. What would that explanation be? If it is inferential, it would seem that we could run all of McGrath's objections to inferentialism once more and they will be just as good. Ditto for any objections to intuitionist explanations of our moral knowledge in actual particular cases. That makes it hard to see how the perceptualist can allege that intuitionism comes with serious theoretical disadvantages. And finally, there is at least something to be said for the thought that our moral knowledge of actual and imagined cases seem similar enough that we should explain both the same way. Granted that appeals to phenomenology are not always probative in epistemology, as we saw above, but the perceptualist at least should bear some burden in explaining why the explanation of moral knowledge in actual cases and imagined cases is utterly different when the two seem so very similar.

So much for the first option. Now let's consider the second: the perceptualist could say that her view also explains how you can actually know that what the imagined adult did was wrong. This is no easy task. Since wrongness is not instantiated, there is no wrongness for the actual you to perceive in this imagined scenario. Perceptualists could deny this. But in explaining how moral perception works 'offline' in relation to imagined cases, they will have to take on many further

commitments.¹⁵ Even if this explanation works, the intuitionist will still have a theoretical advantage. Intuitionism explains more with less. And *ceteris paribus*, we should prefer hypotheses that manage to explain more with less.

Perhaps, at this point, the perceptualist will push back that all things are not equal. Perhaps her theory becomes more complex once it expands to allow for moral knowledge in actual and imagined particular cases, but the intuitionist explanation is worse along some other dimension. The perceptualist might press, as McGrath did,¹⁶ that intuitionism is worse as it leads to skepticism in cases involving interpersonal disagreement.

Kieran Setiya recently made this objection to intuitionism. First, he says, we should be agnostic in cases involving interpersonal perceptual disagreement, such as in the following example from Adam Elga:

You and a friend are to judge the same contest, a race between Horse A and Horse B. Initially, you think that your friend is as good as you at judging such races. In other words, you think that in case of disagreement about the race, the two of you are equally likely to be mistaken. The race is run, and the two of you form independent judgments. As it happens, you become confident that Horse A won, and your friend becomes equally confident that Horse B won (Setiya 2012, pp. 28-29, citing Elga).

Second, he continues, if the epistemology of intellectual seemings mirrors the epistemology of perceptual seemings, then we should also be agnostic in cases involving interpersonal moral disagreements like the following:

Suppose ... that you belong to a homogeneous community whose ethical beliefs are true and who are non-ethically well-informed. Let us add that your beliefs are proportioned to your intuitions, finding the perfect balance of simplicity, power, and explanatory depth, weighed against fidelity to how things seem. For the first time, you meet a stranger. He agrees with you outside of ethics, but when it comes to practical reason, his intuitions are shocking. Fill in the details as you like. Perhaps it seems to him that we should act on our final desires, whatever they are, that we should be utterly selfish, that we should maximize aggregate happiness no matter who is trampled on the way. Despite this, his ethical beliefs are as well-proportioned to his intuitions as your beliefs are to yours. It turns out that he, too, belongs to a homogenous community, exactly as numerous as your own. What should you now believe? If intuitions play the role of perceptual appearances and provide us with evidence in a similar way, the horse race argument applies. Before you meet the stranger, you have justified beliefs. But now you know how things seem to him, your confidence should fade. Apart from the intuitions in dispute, you have no basis for your ethical beliefs, and no more evidence of your reliability than his. You have no independent reason to discount the stranger's intuitions. This being so, you should give how things seem to him as much evidential

¹⁵ See, e.g., Johnston (2004) and Lord (forthcoming).

¹⁶ This was McGrath's central response in personal correspondence.

weight as how they seem to you. You should thus become agnostic about the ethical questions on which you disagree (*ibid.*, p. 30).

Setiya says that agnosticism ‘is not the right response’—‘we should not defer to moral monsters but condemn them, no matter how coherent or numerous they are’—so he concludes that intuitionism is false.¹⁷

Does this objection succeed? And if so, does it help the perceptualist? Let’s take these questions in reverse order, as the second is easier to answer. No, it can’t help the perceptualist. The objection starts with a claim about the epistemology of perception in general: in cases of disagreement involving perceptual seemings, become agnostic. If that’s true for perception in general, it’s true for moral perception. And if it’s not true for moral perception, the intuitionist will say it’s not true of moral intuition. The intuitionist’s mantra, recall, is that whatever we say about the epistemology of *perceptual seemings*, the intuitionist will say the same about *intellectual seemings*. If perceptual seemings that yield moral knowledge allow us to stand our ground in the face of disagreement, why would intellectual seemings that yield moral knowledge be different?

We just saw that if the objection works at all, it works equally well against the perceptualist; intuitionism has no distinctive theoretical vice here. Thankfully for both theories, the objection does not work at all. The explanation for why turns on a key phrase from Setiya above: ‘You have no independent reason to discount the stranger’s intuitions’, so ‘you should give how things seem to him as much evidential weight as how they seem to you’. This claim should be rejected, as it gets the wrong results about the epistemology of perception in general.¹⁸ As Katia Vavova has argued, we need to distinguish between *not having a reason for being confident that we’re more reliable* and *having a reason for being confident that we’re not more reliable*.¹⁹ It matters a great deal where we put the ‘not’. If we should be agnostic when *we have no reason* for being confident that we’re more reliable, then we get the result that we should be agnostic when we disagree with someone about everything (or at least, we get this result so long as we accept a plausible commitment called Independence).²⁰ But we don’t get this result if we set the bar for agnosticism higher by requiring that we have a reason to think that we’re not more reliable. In the horse race case, this higher bar is met—reasonably, ‘you think that in case of disagreement..., the two of you are equally likely to be mistaken’. In Setiya’s example, this higher bar is not met: ‘you have no independent reason to discount the stranger’s intuitions’, but you do not have an independent reason to think that you and the stranger are equally reliable!

¹⁷ *Id.* The dialectic in Setiya’s chapter is actually more complex: he concludes that intuitionism either gets the wrong results about disagreement or is committed to an implausible form of egoism; either way, he concludes, it turns out to be false.

¹⁸ See Christensen (2011, especially at p. 15).

¹⁹ Vavova (2014, especially at pp. 316–318). Cf. Setiya (2012, p. 20, n. 14).

²⁰ Independence is formulated by Vavova (2014, p. 309) as follows: ‘In evaluating the epistemic credentials of another’s expressed belief that *p*, in order to determine how (or whether) to modify my own belief about *p*, I shouldn’t rely on my initial belief that [not] *p*, nor on the reasoning behind that belief’. (The negation in the square brackets has been added to amend a typographical error in the text.)

So, it looks like the objection that intuitionism leads to skepticism relies on a false premise about the epistemology of perceptual disagreement, and even if it doesn't the objection would also target perceptualism. This suggests that if perceptualists want to motivate their view by ruling out rival hypotheses, they need to say much more to rule out intuitionism.

5 The reductive epistemologist

Finally, we come to the reductive epistemologist. This view was advanced a few years ago by Setiya as a response to the skeptical challenge allegedly posed by disagreement. Reductive Epistemology has received little critical attention; McGrath, however, offers an interesting objection to the view, with the aim of showing that it is not a plausible rival to perceptualism.

Let's start by getting the view into focus. Setiya assumes the following:

Supervenience: If x 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 (Setiya 2012, p. 10)

To illustrate this with a toy example: if x is bad, it is bad in virtue of being painful, such that necessarily if something is painful then it is also bad.

This looks a lot like traditional formulations of the supervenience of the ethical on the non-ethical, but with two twists. One is that this claim is formulated in terms of objects falling under concepts; that's intended to make the thesis compatible with certain naturalist theories. The other is that **Supervenience** picks out *two* relations that hold between, in the illustrative example above, painfulness and badness: painfulness *necessitates* badness, and badness obtains *in virtue of* painfulness. In other words, **Supervenience** captures both supervenience and grounding.²¹

Assume, for the moment, that **Supervenience** obtains for some x , N and E . Where that assumption is true, Setiya's epistemological view is as follows:

Reductive Epistemology: What justifies S in believing that x falls under E just is what justifies S in believing that x falls under N .²²

So, in relation to our toy example, what justifies you in believing that something is bad just is what justifies you in believing that its painful. This view *reduces* moral evidence to non-moral evidence, hence the label.

An important feature of this view is that it is *non-inferential*. It does not require ordinary agents to believe, or have any reason to believe, the facts about supervenience and grounding. Say that if x is bad, it is bad in virtue of being painful, such that necessarily if something is painful then it is also bad. Call this a normative law. In the parking lot case, you don't need to know that this normative law is true,

²¹ Depending on how these claims are construed, the claim about supervenience could be redundant: it is widely thought that if the fact that p is grounded in the fact that q , then q entails p . [See especially 'the entailment principle' in Rosen (2010, p. 118.).]

²² Setiya (2012, p. 49).

or anything of the like. Your evidence that striking the child is painful is sufficient for you to have justification for believing that striking the child is bad. *Mutatis mutandis* for any other ethical concept (e.g., WRONGNESS), and any more complex normative law about that ethical concept. That's the strategy the reductive epistemologist will use to explain ordinary moral knowledge, and it seems well-designed to avoid McGrath's objections to inferentialism.

However, McGrath raises a further objection to Reductive Epistemology that turns on transitivity. Assume that pain supervenes on microphysical facts, namely C-Fiber firings. The supervenience relation is transitive, so it follows from this and our earlier assumption that badness supervenes on C-Fiber firings: if the moral supervenes on the mental, and the mental supervenes on the material, the moral supervenes on the material. Now, McGrath presses, the Reductive Epistemologist must say that evidence of how things are microphysically *just is* evidence of how things are morally.

This commitment, McGrath thinks, would be problematic. Consider an ordinary agent in the following extraordinary circumstances: she is a subject in an experiment where neuroscientists show her a video of the microphysical changes in a child's neural system after a certain stimulus. What stimulus? A slap from a frustrated adult in a parking lot. But our experimental subject does not know this. Nor does she know the normative law about badness and pain. And nor does she know the corresponding mental law about pain and C-Fibers. All she knows from the scientists' testimony and the video is that C-Fibers were firing. Question: does the ordinary agent thereby have evidence of badness?

McGrath believes, quite rightly, that the answer is 'No'. She also believes that the Reductive Epistemologist must say 'Yes'. The Reductive Epistemologist could try to say that this result is not so bad, I but think the Reductive Epistemologist need not even answer 'Yes' in the first place.

Why is that? It's because **Supervenience** is not just supervenience. It's supervenience plus grounding. The former is undeniably transitive (if there's no A difference without a B difference, and no B difference without a C difference, there's no A difference without a C difference). But whether the latter relation is transitive is controversial. If the moral is grounded in the mental, and the mental is grounded in the material, does it follow that the moral is grounded in the material? Maybe. But maybe not. If grounding is one relation, it is a type of explanatory relation, and not all explanatory relations are obviously transitive. Consider causation: that causation is a non-transitive explanatory relation is at least tenable. Moreover, it is not even clear that grounding is just one relation. Many treat it as a term for a family of different asymmetric and irreflexive necessitation relations (some of which are arguably non-transitive).²³ If p grounds₁ q and q grounds₂ r , must p grounds _{n} r ? Not obviously. Moreover, many in metaethics have argued that non-naturalists should think that there is a relation of *normative grounding* that is distinct from the relation of *metaphysical grounding*, such that if the material metaphysically grounds the mental and the mental normatively grounds the moral it

²³ See, among others, Wilson (2014) and Bennett (2017).

does not follow that the material grounds the moral in any interesting way.²⁴ So, there are options for the Reductive Epistemologist to avoid being saddled with the committed to the transitivity of **Supervenience**, and thereby avoid committing to the claim (in our toy example above) that how things are materially *just is* evidence of how things are morally.

It is worth pausing here to note that the lattermost option (positing a distinct species of normative grounding) is closely associated with non-naturalism about normativity. This may not be in keeping with Setiya's own metaethical sympathies. But it should be an option for Reductive Epistemologists more generally. Setiya motivates his formulation of **Supervenience** in terms of concepts rather than properties to make the view friendly to naturalists who hold that the property of badness is identical to a natural property (e.g., painfulness). But as Gideon Rosen (2017) has argued at length, claiming that the normative is metaphysically grounded in the non-normative is unfriendly to non-naturalism. This makes **Supervenience** less ecumenical than it was perhaps intended to be. So the Reductive Epistemologist should be open to a plurality of positions on how we understand 'in virtue of' talk in these contexts, to ensure that Reductive Epistemology is not tethered to a reductive metaphysics.

Could the Reductive Epistemologist just modify their view by dropping the commitment to grounding? Not without generating a litany of counterexamples. These will include the counterexample above, which involves transitivity. But it will also include cheap supervenience theses. If there can be no A difference, then trivially there can be no A difference without a B difference. So if we cannot vary the necessary normative facts, then they trivially supervene on anything and everything. This means that evidence of literally anything and everything will be evidence of all of the necessary normative truths, for a modified Reductive Epistemology.

So, if the Reductive Epistemologist is stuck with a commitment to some sort of grounding relation holding between the normative and the non-normative, but wants their view to be ecumenical, then they should be open to there being a plurality of grounding relations. That makes the view friendly to a least one kind of metaethical non-naturalism. And it allows Reductive Epistemologists to avoid McGrath's counterexample.

Interestingly, this isn't the route Setiya prefers to take in response to such concerns. Instead, he emphasizes that Reductive Epistemology only explains propositional justification, rather than doxastic justification. In the case where scientists give you evidence of C-Fiber firings, you are only propositionally justified in believing that what's happening is bad. To be doxastically justified, Setiya adds, you need (at least) to exercise a reliable disposition to form beliefs about badness on the basis of C-Fiber firings.²⁵

This makes the Reductive Epistemologist's explanation for why the ordinary agent knows that the adult's act was wrong more complex. Knowledge requires doxastic justification. So the agent not only needs evidence that the adult's act

²⁴ Cf. Bader (2017) and Berker (2018).

²⁵ See Setiya (2012, pp. 64–5); see also McGrath (17–18).

causes pain, or whatever; she also needs to be exercising a reliable disposition to form beliefs about wrongness.

McGrath thinks that this appeal to ‘the machinery of epistemically reliable dispositions *does* provide a satisfying answer’ to the question of why if you only had evidence that some act makes C-Fibers fire, your belief that the act is bad would not be justified; but she contends that Setiya still has not explained why it would be uniquely rational in this case for you to suspend judgment about that matter (19). I agree with the latter point, but I am not sure about the former. Is the answer satisfying? What if you just started believing that acts are wrong whenever you believed that they make C-Fibers fire, without having any reason whatsoever to believe the relevant normative laws or mental laws that link C-Fibers to badness? This would give you a reliable belief-forming disposition. Is it enough to give you doxastic justification? I think the answer is ‘No’. And Setiya agrees: his view captures this by pointing out that it is an *accident* that your belief-forming disposition is reliable here. Setiya has much to say about accidentality and how it relates to knowledge and justification.²⁶ I will not take issue with views on it, but rather suggest that accidentality is not the most profitable place to look to explain why your belief is unjustified. You are not disposed to form such beliefs because they are true; the reason why is that (a) you have no knowledge of the supervenience of the moral on the mental and of the mental on the physical, and (b) the latter relation (between pain and C-Fibers firing) is a posteriori and outside your ken.

This suggests a diagnosis for the, or at least a, problem for Reductive Epistemology. Setiya rightly wants to avoid an overly intellectualized pictures like the Deductive Model: ‘it is not a condition of evidentially justified belief that one believe a conditional whose antecedent is one’s evidence and whose consequent the content of one’s belief’.²⁷ So he embraces a ‘natural’ alternative: that doxastic justification requires one to ‘manifest a disposition’ that tracks propositional justification, and knowledge requires one to ‘manifest a disposition’ that tracks truth (where the reliability of these dispositions is ‘no accident’). This alternative may look ‘natural’ when we restrict our focus to cases in which the truth of the relevant conditionals is knowable a priori (which Setiya does, in focusing on normative laws). But things are different when we turn to conditionals that are only knowable a posteriori: in many such cases not only will we not believe the relevant conditionals, but we will have no reason to believe that they are true: in that sense, they’re outside our ken. This does not mean that C-Fibers firing cannot be evidence of badness. But it does suggest that more will be required for you to *possess* this evidence, such that beliefs (reliably) formed on this basis can be justified.²⁸

²⁶ Setiya (2012, pp. 65 ff). The relevant principle governing accidental reliability, *K*, explicitly applies to knowledge (p. 96), but extends to doxastic justification since doxastic justification depends on the capacity for knowledge (ch. 4). I am grateful to Kieran Setiya for clarifying and confirming these details in personal correspondence.

²⁷ *Ibid.*, p. 63.

²⁸ What more is required? This turns on the long-standing debates about what it takes to have a reason: see Schroeder (2008), Whiting (2014), Sylvan (2015) and Wodak (2017).

The dialectic here, admittedly, has become quite messy, in large part because Reductive Epistemology on its own does not provide clear enough verdicts about cases: to do so, it must be coupled with orthogonal theses about metaphysics (is grounding transitive?) and epistemology (what additional feature is required for doxastic justification on this view?). It is worth registering, however, that it would be nice to be able to test Reductive Epistemology against simpler, more mundane cases. If Reductive Epistemology goes awry, it seems to go awry in artificial cases involving neuroscientists and given esoteric claims about transitivity. It does not seem to stumble in explaining ordinary moral knowledge. So it is not an unattractive rival to perceptualism with respect to our original *explanandum*: how do you know that what the adult did was wrong?

6 Conclusion

The two most important features of McGrath's approach to defending perceptualism are (a) her focus on the importance of explaining ordinary moral knowledge, and (b) her argumentative strategy of eliminating rival explanations to moral perception. I am very sympathetic to the focus on ordinary moral knowledge, but I am skeptical that we are yet in a position to settle on moral perception by ruling out its rivals. This is partly because with some rivals to moral perception (like moral inference and the reduction of moral evidence to non-moral evidence) come in far more varieties than we might have thought, and partly because there are other alternative *explanans*, like moral intuition, that provide formidable rivals.

References

- Bader, R. (2017). The grounding argument against non-reductive moral realism. *Oxford Studies in Metaethics*, 12, 106–134.
- Bennett, K. (2017). *Making things up*. Oxford: Oxford University Press.
- Berker, S. (2018). The unity of grounding. *Mind*, 127(507), 729–777.
- Christensen, D. (2011). Disagreement, question-begging, and epistemic self-criticism. *Philosophers' Imprint*, 11(6), 1–22.
- Faraci, D. (2015). A hard look at moral perception. *Philosophical Studies*, 172(8), 2055–2072.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: evaluating the evidence for 'top-down' effects. *Behavioral and Brain Sciences*, 39, 1–72.
- Fitelson, B. (2010). Strengthening the case for knowledge from falsehood. *Analysis*, 70(4), 666–669.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge: MIT Press.
- Harman, G. (1984). Logic and reasoning. *Synthese*, 60(1), 107–127.
- Horty, J. (2007). Reasons as defaults. *Philosophers' Imprint*, 7, 1–28.
- Horty, J. (2012). *Reasons as defaults*. Oxford: Oxford University Press.
- Huemer, M. (2001). *Skepticism and the veil of perception*. Lanham: Rowman & Littlefield.
- Huemer, M. (2007). Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research*, 74(1), 30–55.
- Johnston, M. (2004). The obscure object of hallucination. *Philosophical Studies*, 120(1–3), 113–183.
- Kagan, S. (2001). Thinking about cases. *Social Philosophy and Policy*, 18(2), 44–63.
- Leslie, S. J. (2013). Essence and natural kinds: When science meets preschooler intuition. *Oxford Studies in Epistemology*, 4, 108–165.

- Lord, E. (forthcoming). How to learn about aesthetics and morality through acquaintance and testimony. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics*.
- McGrath, S. (2018). Moral perception and its rivals. In R. Cowan, A. Bergqvist (Eds.), *Evaluative perception*. Oxford: Oxford University Press.
- Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffmann (Eds.), *Modality: Metaphysics, logic, and epistemology* (pp. 109–136). Oxford: Oxford University Press.
- Rosen, G. (2017). Metaphysical relations in metaethics. In T. McPherson & D. Plunkett (Eds.), *Routledge companion to metaethics*. Abingdon-on-Thames: Routledge.
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139(1), 57–71.
- Setiya, K. (2012). *Knowing right from wrong*. Oxford: Oxford University Press.
- Strohming, M. (2015). Perceptual knowledge of nonactual possibilities. *Philosophical Perspectives*, 29(1), 363–375.
- Sylvan, K. (2015). What apparent reasons appear to be. *Philosophical Studies*, 172(3), 587–606.
- Vavova, K. (2014). Moral disagreement and moral skepticism. *Philosophical Perspectives*, 28(1), 302–333.
- Väyrynen, P. (2008). Some good and bad news for ethical intuitionism. *The Philosophical Quarterly*, 58(232), 489–511.
- Warfield, T. (2005). Knowledge from falsehood. *Philosophical Perspectives*, 19, 405–416.
- Whiting, D. (2014). Keep things in perspective: Reasons, rationality and the a priori. *Journal of Ethics and Social Philosophy*, 8(1), 1–22.
- Wilson, J. (2014). No work for a theory of ground. *Inquiry*, 57, 535–579.
- Wodak, D., Leslie, S.-J., & Rhodes, M. (2015). What a loaded generalization: Generics and social cognition. *Philosophy Compass*, 10(9), 625–635.
- Wodak, D. (2017). Can objectivists account for subjective reasons? *Journal of Ethics and Social Philosophy*, 12(3), 259–279.
- Wodak, D., & Sarah-Jane, L. (2017). The mark of the plural: Generic generalizations and race. In P. C. Taylor, L. M. Alcoff, & L. Anderson (Eds.), *Routledge companion to philosophy of race* (pp. 277–289). Abingdon-on-Thames: Routledge.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.