White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of Leeds.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/114536/

Version: Accepted Version

Book Section:
Vayrynen, P (2018) Doubts about Moral Perception. In: Bergqvist, A and Cowan, R, (eds.) Evaluative Perception. Mind Association Occasional Series. Oxford University Press. ISBN: 9780198786054.

https://doi.org/10.1093/oso/9780198786054.003.0006

(c) 2018, the several contributors. This is an author produced version of a chapter published in Evaluative Perception, edited by A. Bergqvist & R. Cowan and published by Oxford University Press. Reproduced by permission of Oxford University Press.

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

UNIVERSITY OF LEEDS
University of Sheffield
UNIVERSITY of York

Doubts about Moral Perception

Pekka Väyrynen
University of Leeds

Abstract: I defend doubts about the existence of genuine moral perception, understood as the claim that at least some moral properties figure in the contents of perceptual experience. Standard examples of moral perception are better explained as habitual implicit inferences or transitions in thought whose degree of psychological immediacy varies with how readily non-evaluative perceptual inputs, jointly with the subject's background moral beliefs, training, and habituation, trigger the kinds of phenomenological responses that moral agents are disposed to have when they represent things as being morally a certain way.

1. Introduction

Can moral properties – such as being wrong, or good, or unjust – be *perceived*, in any robust sense? Many philosophers claim so.[1] This paper will press doubts.

It is tempting to count perception as one mode of access to at least some instances of moral right and wrong, goodness and badness, and the like.[2] We often

1 See Iris Murdoch, *The Sovereignty of the Good* (London: Routledge & Kegan Paul, 1970); David McNaughton, *Moral Vision* (Oxford: Blackwell, 1988), 57–59; Michael DePaul, *Balance and Refinement* (London: Routledge, 1993), 202–9; Lawrence Blum, *Moral Perception and Particularity* (Cambridge: Cambridge University Press, 1994); John Greco, *Putting Skeptics in Their Place* (Cambridge: Cambridge University Press, 2000), 231–48; Michael Watkins and Kelly Dean Jolley, 'Pollyanna Realism: Moral Perception and Moral Properties', *Australasian Journal of Philosophy* 80 (2002), 75–85; Terence Cuneo, 'Reidian Moral Perception', *Canadian Journal of Philosophy* 33 (2003), 229–58; Sarah McGrath, 'Moral Knowledge by Perception', *Philosophical Perspectives* 18 (2004), 209–28; Jennifer Wright, 'The Role of Perception in Mature Moral Agency', in *Moral Perception*, ed. Jeremy J. Wisnewski (Newcastle: Cambridge Scholars Publishing, 2007), 1–24; Timothy Chappell, 'Moral Perception', *Philosophy* 83 (2008), 421–37; Andrew Cullison, 'Moral Perception', *European Journal of Philosophy* 18 (2010), 159–75; Jonathan Dancy, 'Moral Perception', *Proceedings of the Aristotelian Society Supplementary Volume* 84 (2010), 99–117; Justin P. McBrayer, 'A Limited Defense of Moral Perception', *Philosophical Studies* 149 (2010), 305–20; Sabine Roeser, *Moral Emotions and Intuitions* (Basingstoke: Palgrave Macmillan, 2011); Robert Audi, *Moral Perception* (Princeton: Princeton University Press, 2013); Robert Cowan, 'Perceptual Intuitionism', *Philosophy and Phenomenological Research* 90 (2015), 164–93; Preston J. Werner, 'Moral Perception and the Contents of Experience', *Journal of Moral Philosophy* 13 (2016), 294–317.

2 I'll focus on descriptively 'thin' evaluative properties, such as being good/bad, right/wrong, and just/unjust. Some defenders of moral perception pin their case on 'thick' evaluative properties such as kindness, cruelty, and selfishness. Dancy, for instance, writes: 'I can perceive its being her turn

describe our moral experiences – experiences in which things strike us as being a certain way morally – in apparently perceptual terminology. Consider:

> *Cat:* 'If you round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it, you do not need to *conclude* that what they are doing is wrong; you do not need to figure anything out; you can *see* that it is wrong.'[3]

> *Tiananmen Square:* 'Could one not see terrible injustice by viewing soldiers shooting citizens who are peaceably criticizing their government?'[4]

But care is due. Sometimes 'see' designates a kind of intellectual understanding, as when we say things like 'I see that P doesn't entail Q, but surely the truth of P would be an unexplained accident unless Q were true'. And not all experiences that have a phenomenal character which 'see' may be used to designate are clearly perceptual; talk of 'seeing spots' might be an example. Even when 'see' is used in a perceptual context, the mental state to which it is applied may not be a case of perception. We can talk of a scientist 'seeing' a proton by seeing a vapor trail in a cloud chamber, but this doesn't mean we can perceive protons. So although examples like *Cat* and *Tiananmen Square* can be felicitously described as cases of 'seeing' something to be wrong, it doesn't follow that these moral experiences involve perceptual representations of moral properties; they may merely be prompted by perceptual input. But as I'll discuss below, developments in the philosophy of perception provide tools for assessing whether moral perception occurs in any robust sense.[5]

(my favourite example of a thick moral concept), the caring nature of a response, the courtesy or the rudeness of a gesture'; Dancy, op. cit. note 1, 114. This is an appealing move: thick properties are more closely moored to properties that are agreed on all hands to be perceptible. But it is also a risky move: it is controversial whether thick terms and concepts are relevantly evaluative in content, and if they aren't, they won't determine evaluative properties as their referents. Perceiving something to be cruel, or kind, or selfish, wouldn't then be evaluative perception in the relevant sense. See Pekka Väyrynen, *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics* (New York: Oxford University Press, 2013). I also suspect that the doubts I'll raise against the perception of thin evaluative properties can also be run against thick evaluative properties; see note 35 below.

3 Gilbert Harman, *The Nature of Morality* (Oxford: Oxford University Press, 1977), 4.

4 Audi, op. cit. note 1, 31.

5 Other discussions of moral perception that are significantly informed by recent work in the philosophy of perception include Cowan, op. cit. note 1, and Werner, op. cit. note 1.

The claim that there is moral perception might be found attractive on different grounds. One is epistemological. Assuming we have some moral knowledge, *how* do we have it? Very many people think that if anything gives us knowledge of the external world, perception does. So if moral properties were something things could be perceived to have, this might help with the difficult project of developing a plausible moral epistemology and integrating it with a plausible overall epistemology. The other is phenomenological: certain moral experiences might be best explained as perceptions of moral qualities. The two motivations are logically distinct. Even if some representations of moral properties were best explained as genuinely perceptual, it would be a further question whether this makes an independent contribution to a plausible moral epistemology. (Conversely, the desired epistemological benefits might be secured by some other modes of sensitivity to situations.[6]) But some moral experiences must be perceptual for moral perception to provide epistemological benefits.

This paper aims to undermine the phenomenological motivation for moral perception. I'll argue that cases like *Cat* and *Tiananmen Square* aren't best explained as cases of moral perception even if we assume that the contents of perceptual experience aren't in general restricted to 'lower level' properties like shape, motion, and color. Positing distinctively perceptual representations of moral properties would add no explanatory power because a simpler and a more unified account treats their representations in the relevant cases as resulting instead from implicit transitions in thought. In closing I'll briefly discuss the negative implications of my argument for the epistemological role of moral perception. (I'll provide no separate argument that moral perception wouldn't be an epistemologically independent source of knowledge or justification.)

---

6 A virtuous person has a better epistemic access to moral facts than someone faced with the same situation who lacks virtue. This sensitivity might be treated as a perceptual capacity: the virtuous person “sees” things that the non-virtuous doesn't; see e.g. Murdoch, op. cit. note 1, and John McDowell, 'Virtue and Reason', *The Monist* 62 (1979), 331–50. Separate discussions of whether it is a robustly perceptual capacity or how exactly a virtuous character interacts with the contents of perceptual experience are beyond this paper's scope; the bearing of the discussion to follow on these issues will remain largely implicit. The same goes for our experiences of certain actions as afforded or even mandated by our environment; for discussion, see Susanna Siegel, 'Affordances and the Contents of Perception', in *Does Perception Have Content?*, ed. Berit Brogaard (Oxford: Oxford University Press, 2014), 51–75.

2. Moral Perception: Focusing the Issue

My argument will concern the following literal notion of moral perception:

> (MP) At least some moral properties can figure in the contents of (veridical) perceptual experience.

Perceptual experience has phenomenal character: there is 'something it's like' to have a perceptual experience. As stated, (MP) follows most recent discussions of moral perception in assuming a representational theory of perception: roughly, to have a perceptual experience of an object O as having a property F is to be in a perceptual mental state (a state with phenomenal character) which has the representational content that O is F. For simplicity, I'll assume that the relevant kind of content can be understood as a kind of condition under which the experience which has such a content is accurate.[7]

To focus discussion, I'll also grant two general claims about perceptual experience which have recently attracted considerable attention. The first is that perceptual experience is *cognitively penetrable*. This is to say that what one perceives can be influenced, in a direct and non-trivial way, by the states of the subject's cognitive system, such as moods, beliefs, and desires.[8]

The second claim is that perceptual content can be *rich*. Perception isn't restricted to representing 'low level' properties, such as spatial properties, color, shape, pitch, odor, motion, and illumination, which are perceived if anything is. Some other properties or relations can also be represented in perceptual experience.[9] This thesis has been defended for natural kind properties (such as being a pine tree), artifactual kind properties (such as being a table), causal relations, semantic properties, and dispositional properties (such as being

7 This is how Siegel describes what she calls the 'Content View' of perceptual experience; see Susanna Siegel, *The Contents of Visual Experience* (Oxford: Oxford University Press, 2010), 4, 28–30. Below I'll grant what Siegel calls the 'Rich Content View' (ibid., 7, 97).

8 The cognitive penetrability of perceptual experience is usefully surveyed in Dustin Stokes, 'Cognitive Penetrability of Perception', *Philosophy Compass* 8 (2013), 646–63. See also Susanna Siegel, 'Cognitive Penetrability and Perceptual Justification', *Noûs* 46 (2012), 201–22, and Jona Vance, 'Emotion and the New Epistemic Challenge from Cognitive Penetrability', *Philosophical Studies* 169 (2014), 257–83.

9 The cognitive penetrability of perceptual experience and the richness of perceptual content are distinct issues; see Siegel, op. cit. note 7, 10. Perceptual content could be rich (or even feature moral properties in particular) even if the contents of perceptual experience weren't cognitive penetrable. And perceptual content might be cognitive penetrable but not rich if cognitive states only influenced *which* low level properties figure in perceptual content.

edible), among others. These are often grouped together as 'high level' properties. It is unclear whether this is to be understood as a claim about metaphysical hierarchy: odors aren't obviously in a different (and more fundamental) metaphysical boat than natural kinds, and semantic properties and causal relations not obviously in the same. So perhaps a high level property is supposed to be such that if it is perceived at all, perceptions of it depend on the perception of other (low or low*er* level) properties. Moral properties would qualify as high level properties in this sense. (In fact, they would be high*er* level relative to some other high level properties. The non-moral properties on which moral properties depend will often include other high level properties, only some of which will plausibly be perceptible if any are.)

These two assumptions will be helpful in two respects. First, they help to distinguish doubts about (MP) which are relatively local to the moral case from more general doubts about whether high level properties can be perceived. My doubts will be of the former type. Second, the two assumptions help to distinguish (MP) from neighboring views which accord perception some weaker role in moral experience than (MP) does.[10] When it comes to experiences like *Cat*, it is one thing to say that an overall mental state that has a perceptual experience as a component can also involve a representation of a moral property as another component, quite another to say that the moral property figures in the content of that perceptual experience. It seems to make sense to speak of 'the overall experience I'm having' when, surrounded by strange noises in a dark forest, I feel frightened, and to say my overall experience would be different if I felt frightened and the forest were silent.[11] The issue is *how* moral properties are represented in certain overall experiences, not *whether* they are.

Those who accept (MP) and those who reject it can agree that we often perceive non-normative and non-evaluative properties on which moral and other normative and evaluative properties depend for their instantiation. (I'll bracket a host of complex issues concerning normative dependence.) The hoodlums can be perceived to ignite the cat, irrespective of whether the wrongness that is grounded in their doing so also figures in the contents of our perceptions. Both sides can also agree that perception can assist us in getting a clearer moral gauge

10 Many discussions of moral perception elide one or another of the distinctions that I draw in what follows, including the works cited in note 1 by McNaughton, DePaul, Greco, Watkins and Jolley, Cuneo, McGrath, Wright, Chappell, Dancy, and Cullison.

11 I hope an example is enough. I don't have an account of experience individuation to deal with such fascinating but complicated phenomena as cross-modal integration and influence.

of things. It can, for instance, play a role in developing discriminative abilities that make us see behaviors that in no way alter in a different light – as refreshingly simple rather than unrefined, or spontaneous rather than lacking in dignity.[12]

Both sides can agree that in cases where perception plays some important role in informing us about events or states that bear moral properties, our overall experience can include representation of moral properties. In *Cat*, for instance, seeing the hoodlums ignite the cat and its striking you as wrong can belong to the same overall experience irrespective of whether wrongness figures in the contents of the perceptual experience that is part of the overall experience. Both sides can further agree that the ability to represent moral properties on the basis of perception is an acquired skill whose possession and exercise require training and various sorts of background beliefs and other cognitive states. And both sides can agree that if representing something as being morally a certain way on the basis of perceptual experience were a result of an inference or some other transition in thought, such transitions needn't be conscious, but can be implicit.

Those who accept (MP) and those who reject it can also agree that moral properties can make a phenomenological difference to overall mental states. If what the hoodlums are doing in *Cat* didn't strike me as bad, then (all else equal) my overall experience would feel different to me. I might flinch, but I wouldn't have the kind of emotional or affective responses which psychologically normal moral subjects tend to have in scenarios like *Cat* – responses such as indignation, revulsion, disturbance, a felt disapproval, and affective empathy, as well as other phenomenal responses like a sense of 'unfittingness'.[13] No doubt many of these responses can also co-occur. In this way moral properties can make a difference to the phenomenology of the experience. In general, a property F can make a phenomenological difference to an overall mental state S irrespective of whether F figures in the contents of a perceptual experience that S has as a part.

At least in psychologically normal subjects, moral experiences (things striking a subject as being morally a certain way) seem closely bound up with emotional or affective responses. We shouldn't assume that emotions themselves

12 Compare Murdoch, op. cit. note 1, 16–17.

13 Audi, op. cit. note 1, mentions all of these states in discussing moral perception. Audi doesn't endorse (MP), however. He only claims that we perceive moral properties by perceiving the non-moral properties that ground their instantiations. For discussion, see Antti Kauppinen's review of Audi, *Notre Dame Philosophical Reviews 2013.06.29* (https://ndpr.nd.edu/news/40724-moral-perception/).

are a kind of perception.[14] But both the proponents and the opponents of (MP) can allow two weaker claims about how emotional and affective responses relate to the overall mental states in question.

The first claim is that the relevant emotional and affective dispositions (for indignation, revulsion, disturbance, felt disapproval, affective empathy, and the like) are normally shaped in part by the same background moral beliefs which presumably also play a role in representations of moral properties. It is in part because I believe that causing seemingly gratuitous suffering to sentient creatures is normally wrong that I'm disposed both to be disturbed and repulsed when I see what the hoodlums are doing and to represent their behavior as wrong. (One simple model would portray the relevant background moral belief as a moral principle which tends to get engaged in cognitive processing when non-moral perceptual input matches the non-moral condition laid down by the principle.[15]) If my background moral views were different, my dispositions to be disturbed and repulsed would also tend to be different, at least in the conditions that trigger them. So even if my having *some or other* relevant emotional dispositions doesn't depend on my particular background moral beliefs, just *which* such dispositions I have depends in part on the influence of those beliefs.[16]

The second claim is that our phenomenological responses tend to be 'phenomenally integrated' with our perceptions of the non-evaluative features of the situation.[17] Once I take in that the hoodlums are igniting a cat, my perception of what they are doing, my repulsion by their act, and its striking me as wrong won't normally seem like distinct co-occurring experiences. (It won't be like seeing Rio de Janeiro and visually imagining myself hang gliding above it.) The various responses we tend to have in cases like *Cat* normally seem like one fairly seamless overall experience.

14 For criticisms of the 'perceptual model' of emotional experience, see Michael S. Brady, *Emotional Insight* (Oxford: Oxford University Press, 2013).

15 I mention this model just for illustration; options where the background beliefs aren't general principles remain less clear. Note that a background moral principle can play a role in my coming to represent what the hoodlums are doing as wrong even if it isn't exceptionless (*contra* McGrath, op. cit. note 1, 217–20). A moral principle can also influence one's recognitions of moral properties in other ways than by functioning as a premise in inferences from non-moral inputs to moral beliefs; compare Pekka Väyrynen, 'Usable Moral Principles', in *Challenging Moral Particularism*, ed. Mark Norris Lance, Matjaž Potrč, and Vojko Strahovnik (London: Routledge, 2008), 75–106.

16 Thanks to Preston Werner for comments that helped to make this paragraph more precise.

17 I borrow the phrase 'phenomenal integration' from Audi, op. cit. note 1, 38–9. It isn't fully clear how Audi understands the phrase. The main text provides the sense in which I'll use it.

The extent of the common ground between those who accept (MP) and those who reject it helps us to identify what is at issue between them. The distinction between overall experiences and the perceptual experiences they may have as parts, for instance, enables each side to allow that 'such ordinary perceptions as seeing a violent seizure of an old man's wallet or hearing an abusive vulgarity screamed at a conference speaker' can have some such moral element as 'a phenomenal sense of wrongdoing integrated with our perceptual representation of the wrong-making facts'.[18] (MP) advances a claim about where the representation of moral properties is properly located in the relevant overall experiences – namely, in the contents of perceptual experiences they have as parts. The version of (MP) on which I'll focus here says that there are some mechanism(s) of cognitive penetration (CP) which enrich the contents of perceptual experience with moral properties. There may be other mechanisms capable of making (MP) true.[19] But this version of (MP) is of double interest: it promises to contribute to a plausible moral epistemology and it provides a test case for how widely the most prominent current argument for high level perceptual content generalizes.

The foregoing suggests one possible mechanism of the requisite kind. Background moral beliefs might help to enrich perceptual content with moral properties owing to CP facilitated by the emotional and affective dispositions they shape.[20] This would explain not only why certain emotional or affective responses go reliably with certain moral representations, but also why both seem

18 Audi, op. cit. note 1, 44–5.

19 Perhaps moral representations are hard-wired into our perceptual system. Or perhaps (MP) is made true by perceptual learning without a robust CP-type influence of background states with moral content. Neither proposal strikes me as particularly promising.

20 Both Cowan, op. cit. note 1, and Werner, op. cit. note 1, suggest this kind of possibility. Cowan (77–79) also describes various possible mechanisms of CP for the moral case. One direct model is that a cognitive state like an emotional experience is triggered directly by perceptual input and then cognitively penetrates perceptual experience, adding moral properties to its content. (I'll bracket the possibility that background moral beliefs themselves somehow directly penetrate perceptual content. It is unclear how this model is supposed to work or to ground the kind of 'phenomenal contrast' argument for (MP) that I'll discuss in section 3 below.) One indirect model is that perceptual input first 'matches' memory representations of badness and this matching process leads to the formation of a cognitive state like an emotional experience that cognitively penetrates perceptual experience, adding moral properties to its content. Note that rival views to (MP) can adopt similar models short of adding moral properties into perceptual content. (For instance, a matching process can make a phenomenological difference to one's overall experience either way.)

seamlessly integrated with the perceptual inputs that activate them.

The appeal to CP also helps this interpretation of (MP) not to portray moral perception as a function of a special 'faculty' or sensory modality. All sides can agree that if moral properties figured in the contents of perceptual experience, they would have to do so in a different way from colors, shapes, and other low level properties information about which is carried by a distinctive sensory modality. Moral perception would be 'amodal' in the same kind of way as perception of many other higher level properties.[21] In the moral case, then, CP would be better regarded as adding content to a framework provided by non-moral perceptual representations than as altering those representations.[22]

This, at any rate, is the kind of account of moral perception against which I'll press doubts below.[23] I'll argue that certain things which any account of the role of perception in moral cognition should be able to explain don't seem to be best explained by (MP). To make the case concrete, I'll sketch a rival explanation that is at least as powerful as (MP) but simpler and more unified. The view I'll sketch treats the representation of moral properties in the standard examples of moral experience as resulting from transitions in thought whose degree of psychological explicitness or implicitness can vary with how habituated they are and how tightly bound up they are with the relevant phenomenological responses. But we won't have reason to take these transitions to enrich the contents of perceptual experience. Positing distinctively perceptual

21 This kind of parallel between moral properties and some other higher level properties is drawn in the works cited in note 1 by DePaul, Greco, Watkins and Jolley, McGrath, Chappell, Cullison, McBrayer, Cowan, and Werner.

22 Most proponents of (MP) suppose that perception of moral properties somehow depends on perception of non-moral properties. It is less clear how exactly to cash out this assumption. The constellations of non-moral properties on which moral properties depend are typically very complex and some of the relevant base properties (including various complex social and historical phenomena) plausibly don't themselves figure in the contents of perceptual experience. If we perceive moral properties by perceiving their non-moral grounds but not all parts of such grounds are themselves perceivable, what would this imply regarding (MP)? Worries might arise if, for instance, instantiations of moral properties were perceivable by perceiving their non-moral grounds only if the latter metaphysically necessitated or determined the former. Audi supposes something like this, but he underestimates the complexity of non-moral grounds that will follow; see Audi, op. cit. note 1, 58-60, 108. The non-moral grounds of moral properties might thus be too complex for moral properties to be perceived on the basis of perceiving the perceivable parts of their grounds.

23 I detect this kind of account, or at least significant strands of it, in Cowan, op. cit. note 1, and Werner, op. cit. note 1. Audi, op. cit. note 1, is a less clear case; his view allows but doesn't require the possibility of rich perceptual contents effected by mechanisms of CP.

representations of moral properties would add no explanatory power.[24]

3. The Method of Phenomenal Contrast

How might one try to argue for (MP)? Philosophers who accept (MP) often just defend it against objections instead of giving a positive argument.[25] But we might adapt positive arguments from philosophy of perception to the moral case.

Arguments for high level perceptual content often deploy the *method of phenomenal contrast*.[26] I'm happy to grant the general method.[27] It begins with a description of two overall mental states each of which has a perceptual experience as a part. The description is supposed to elicit the intuition that the overall mental states differ in their phenomenology but not in the lower level properties represented by the perceptual experiences. The claim that one but not the other of these perceptual experiences involves a particular high level property (being a pine tree, one event causing another, or whatever) as part of its content will be warranted if it best explains why the mental states differ in their phenomenology. For instance, the claim that we can visually experience the property of being a pine tree might be defended as the best explanation of the kind of phenomenological difference that characterizes your experiences of pine trees before and after you develop a disposition to recognize pine trees.

To assess whether the method of phenomenal contrast can be used to provide a positive argument for (MP), we require a suitable target moral experience and a phenomenologically contrasting experience. A contrast case that involves no moral representation at all would be too dissimilar from the target experience to provide a good test case for (MP). Further, subjects with

---

24 More issues arise here than I can address in what follows, many of them concerning perception in general. For instance, perhaps there is no determinate fact as to whether some higher level property figures in the contents of perceptual experience; compare Heather Logue, 'Visual Experience of Natural Kind Properties: Is There Any Fact of the Matter?', *Philosophical Studies* 162 (2013), 1–12. The discussion to follow concerns what to say about (MP) presuming the issue isn't indeterminate.

25 Examples of the defensive strategy include Cullison, op. cit. note 1, and McBrayer, op. cit. note 1. Werner, op. cit. note 1, offers a positive argument for (MP); see note 28 below.

26 For this strategy, see especially Siegel, op. cit. note 7, chs. 4–5, and 'Which Properties are Represented in Perception?', in *Perceptual Experience*, ed. Tamar Szabó Gendler and John Hawthorne (Oxford: Oxford University Press, 2006), 481–503. My presentation of the method of phenomenal contrast has benefited also from Logue, op. cit. note 24.

27 For various complications that I cannot discuss regarding the relationship between the phenomenal contents of experiences and putative higher level contents, see Nicholas Silins, 'The Significance of High Level Content', *Philosophical Studies* 162 (2013), 13–33.

notably different background moral beliefs will respond differently in many particular cases, but for reasons that are orthogonal to (MP). Recall, however, that both sides can agree that representing something as wrong, or bad, on the basis of non-moral perceptual input tends to be intimately bound up with certain emotional or affective responses at least in psychologically normal subjects. A contrast argument for (MP) might then appeal to phenomenological differences between normal subjects and subjects who share their moral perspective but suffer from affective deficit disorder, such as inability to feel affective empathy for a subject of distress while understanding from behavioral and contextual clues that the subject is in distress.[28]

Thus consider Norma and Alex who both witness a scenario like *Cat*. Their non-moral perceptual inputs (whether low or high level) and spatial attention are the same. What the hoodlums are doing strikes Norma as bad and she also has the sorts of phenomenological responses to the scenario which psychologically normal subjects (with a capacity for affective empathy and so on) tend to have. Alex also represents what the hoodlums are doing as bad. For he is capable of recognizing the cat's distress on the basis of behavioral and auditory cues, and he can learn and apply moral norms. But Alex lacks the sort of phenomenological responses (affective empathy, revulsion, and so on) which Norma experiences. Nor is his representation of what the hoodlums are doing as bad psychologically immediate; he needs to 'figure it out'. So Norma's and Alex's overall experiences differ in their phenomenology. (I'll sometimes refer to this contrast as 'Norma/Alex'.) The question for those who advance a phenomenal contrast argument for (MP) is whether this contrast is best explained by the hypothesis that some such property as being (morally) bad figures in the content of Norma's percepetual experience but not Alex's.

(MP) offers a *prima facie* plausible explanation of the Norma/Alex contrast. (A) Norma's disposition to have the type of experience she has in response to *Cat* can more or less reliably track badness.[29] (B) Nor would Norma easily have the

28 After first devising the sort of contrast pair that follows in the text, I discovered that Werner, op. cit. note 1, gives a similar contrast argument for (MP) from phenomenological differences between 'emotionally empathic dysfunctional individuals' and normally functioning human adults. My presentation here adapts some features of Werner's argument. His argument is more subtle than mine – but not, I think, in ways that affect what I want to say. Finally, we both appeal to cases of interpersonal contrast. I agree with Werner that clean intrapersonal cases of phenomenologically differing moral experiences are hard to come by.

29 I'll pass over the complication that different moral properties can prompt the same response, such as indignation or outrage in the case of both wrongness and injustice.

phenomenological responses in question if what the hoodlums are doing weren't bad; this counterfactual connection can hold at least locally, across the sorts of circumstances Norma is likely to encounter. (C) Norma is also disposed to form moral beliefs (such as that the cat's suffering is bad) based on experiences of this type, and much more spontaneously so than Alex. States that meet conditions like (A)-(C) seem plausible candidates for states that represent some property F, in this case (moral) badness.[30] Although the phenomenological difference between Norma and Alex is emotional or affective in some sense, it might nonetheless make a difference to the contents of perceptual experience owing to some suitable mechanism of CP.

But is this the best explanation of the phenomenal contrast between Norma and Alex? I want to press doubts by pointing to some respects in which this explanation is problematic or inferior to a rival explanation which I'll sketch.

One immediate worry about the above contrast argument for (MP) is that it threatens to show too much. The distinction between an overall mental state and a perceptual experience that it has as a part implies that Norma's overall mental state can meet conditions like (A)-(C) even if her representation of what the hoodlums are doing as bad doesn't figure in the contents of the perceptual experience that is a part of it. Consider a case from physics that mimics *Cat*:

> *Proton:* When Marie, a trained physicist, sees a trail of vapor bubbles in the cloud chamber, she doesn't need to figure anything out, she can just *see* that a proton is going by.

Its striking Marie that a proton just passed can meet conditions (A)-(C). Her disposition to have this experience can reliably track the presence of protons. It is a disposition she has developed in the course of her training and practice as an experimental physicist. This disposition can have a relevantly associated phenomenology. Marie's overall mental state can include aspects such as a feeling of conviction, and it is psychologically immediate in a way that it wouldn't be if Marie were an untrained physicist who needs to reason her way from a visual experience of a trail of vapor bubbles and the physical theory to the presence of a proton. Owing to how tightly these phenomenological responses can (given suitable scientific training and practice) become bound up with observing vapor trails in a cloud chamber, the responses can also be

30 Werner, op. cit. note 1, uses reasoning like this in taking Norma's case to support (MP).

counterfactually correlated with the presence of protons, at least locally. Marie is also disposed to form beliefs about protons passing based on experiences of this type, and more spontaneously so than if she were untrained.[31]

It doesn't follow from (A)-(C) that the phenomenological difference between the trained and untrained response is better explained by the hypothesis that protons figure in the contents of the trained Marie's visual experiences than by the hypothesis that her response is a kind of trained judgment or transition in thought. The point isn't that you cannot see unobservables. Seeing an F is different from having a visual experience with the content that there is an F. Even so, there had better be some limits on what properties can figure in true instances of 'It visually seems to S that there is an F'. Marie's reports aren't a reliable guide to these limits, since what one is trained to say upon having certain experiences ('I see a proton'; 'That's a proton passing') is one thing, the content of those experiences is another. So granting (A)-(C) in *Proton* leaves it open whether Marie *visually* represents a proton passing. It gives no compelling reason to deny that Marie's psychologically immediate proton representation is anything other than trained scientific judgment operating habitually.[32] By parity, then, it remains far from clear that method of phenomenal contrast supports (MP) in cases like Norma/Alex. At the very least, no explanatory schema that we apply to Norma/Alex should straight away imply that we can visually experience protons; that should remain a further issue.

So the question remains: where in the relevant moral experiences is the representation of moral properties best located – the perceptual experiences they have as parts or some other part of the overall mental state? And the question remains: what sorts of considerations would break the issue one way or another?

4. Three Points Against (MP)

I'll now raise three related points against the claim that the phenomenological differences between subjects like Norma and Alex are best explained by (MP). They'll count as points of criticism given the standard theoretical virtues of simplicity, unity, and explanatory power. (One explanation is better than

31 These points chime well with the view of scientific 'intuitions' as a species of trained judgment. See Richard Boyd, How To Be a Moral Realist', in *Essays on Moral Realism*, ed. Geoffrey Sayre-McCord (Ithaca, NY: Cornell University Press, 1988), 181–228.

32 I confess to not being moved by the tentative defense of 'the possibility of the phenomenal representation of protons in perceptual experience' in Cowan, op. cit. note 1. Delicate issues arise here concerning the relation between perceptual experience and perception.

another, all else equal, if it is simpler, and likewise for the other virtues.) The cumulative upshot of these points is that there is a rival model that seems to be able to explain everything that (MP) explains but is simpler and, at least by one relevant measure, more unified.

Comparing rival explanations against multiple criteria can get tricky.[33] In the present context, for instance, positing distinctively perceptual representations of moral properties doesn't automatically make the account less simple. Perhaps each rival must posit some *other* distinctive psychological mechanism. But if a rival explanation can instead appeal to some general mechanism that we need to recognize anyway, independently of the role of perception in moral experience, this is a reason to consider the rival explanation simpler, and possibly more unified. This would make it superior, at least if it also explains at least as much as (MP).

My first critical point builds on points made earlier. When what the hoodlums are doing in *Cat* strikes me as bad, this representation depends somehow on my having certain background cognitive states with moral content. I'm taking (MP) to say that these background states cognitively penetrate the contents of my perceptual experience. But cases like *Proton* suggest an alternative model. The alternative I'll adopt in order to facilitate concrete comparisons is that when Norma sees what the hoodlums are doing in *Cat* and represents it as bad, this representation results from an implicit habitual inference or some other type of transition in thought which can be reliably prompted by the non-moral perceptual inputs jointly with the relevant background moral beliefs.[34] (Perhaps the badness of what the hoodlums are doing fits Norma's perceptual evidence and has a high prior probability given her background moral beliefs.) The transition can also be psychologically immediate and bound up with the relevant emotional or affective responses, explaining how things can 'strike' us as being morally a certain way.[35] Sometimes those who

33 Thanks to Susanna Siegel for pressing me to clarify my methodological assumptions.

34 Robert Cowan suggested to me that this account might overgeneralize to exclude many high level properties whose relation to perceptual content I want to leave open. But whether some higher level property can be perceived needs to be assessed on a case-by-case basis. The rival explanation that I run for moral properties carries no automatic commitment regarding causal relations, kind properties, or semantic properties. (Apparent differences include the way in which emotional or affective responses are involved.) If a rival explanation along those lines is the best (or not, for that matter) for some other higher level properties, so be it.

35 This account of Norma's experience seems structurally parallel to how Marie's experience in *Proton* is bound up with phenomenological responses shaped by her training. Note that a

defend (MP) tend to contrast it with views that involve a very narrow notion of inference. It is therefore worth noting that we can reasonably speak of inference in a broader sense even when someone doesn't explicitly represent the premises or cannot articulate them.[36] Alternatively the relevant transition might be more like a recognition based on taking in a pattern one isn't able to articulate, where what pattern gets recognized may be influenced by prior training and background cognitive states.[37]

The foregoing is a sketch, but the basic explanatory schema is hopefully clear enough. There are many ways to model the relevant kind of transitions more precisely. One picture is that non-moral perceptual input triggers (directly or indirectly) certain emotional or affective dispositions which have been shaped by background moral beliefs and whose manifestations are intimately causally bound up with representing things as being morally a certain way.[38]

My critical point is twofold. First, there is a principled model for explaining how Norma's representation of what the hoodlums are doing as bad can be phenomenally integrated (owing to related emotional and affective dispositions) with her non-moral perceptions without cognitively penetrating the contents of perceptual experience. Second, the model appeals to no psychological states or mechanisms that we don't need to recognize anyway. My outline only appeals to inferences and other reliable transitions in thought and to emotional and affective dispositions. This explanation of cases like Norma/Alex thus promises to be simpler than (MP).

My second point builds on the first. Cases like *Proton* aren't the only instructive parallel for *Cat*. Representations of some non-moral *evaluative* properties seem to parallel moral properties but are more plausibly located outside the perceptual components of overall experiences. Consider:

---

similar account can be run for variants of Norma/Alex where the target property isn't a thin evaluative property such as being bad but a thick property such as being cruel; recall note 2.

36 Nicholas Sturgeon offers epistemological motivations for a broad notion of inference which applies to 'cases in which someone is unable to articulate the premises, and also to cases in which someone is unconscious of making an inference and perhaps even of accepting the premises'; Sturgeon, 'Ethical Intuitionism and Ethical Naturalism', in *Ethical Intuitionism: Re-evaluations*, ed. Philip Stratton-Lake (Oxford: Clarendon Press, 2002), 184–211, at 209. Note that inference so understood remains a process at the personal rather than subpersonal level.

37 Emotions and affects can work like this. For instance, one might be afraid (or find the used car salesman untrustworthy) without being able to put one's finger on what is wrong. Pattern recognition is discussed in the context of moral perception also by Chappell, op. cit. note 1.

38 Compare the various direct and indirect models of CP described in note 20 above.

Fine Wine: Greg, an experienced wine maker, reports that when he samples wine he perceives it as having various non-evaluative qualities which form his basis for classifying it as fine or not. Michael, a wine connoisseur, says that he can taste also fineness in wine.[39]

Greg's and Michael's experiences can also differ phenomenologically, such as in the immediacy of their felt satisfaction in tasting a fine wine and its phenomenal integration with the relevant non-evaluative inputs. Given the reported differences, what would warrant attributing the property of being a fine wine to the content of some of Michael's perceptual experiences but not those of Greg's?

Greg and Michael both can recognize fine wine. But I see no good reason to suppose that Michael has some training, background knowledge, or discriminative ability which Greg lacks. According to the rival explanation sketched above, their overall mental states involve the same kinds of prior expectations, beliefs, training, and dispositions to form certain evaluative classifications in response to certain non-evaluative perceptual inputs. A realistic example of this sort will add that sometimes Greg's response is more immediate than his typical response, and Michael's less so. This suggests more forcefully still that their responses differ primarily in degree: in how psychologically immediate their responses tend to be and how integrated the non-evaluative inputs tend to be with the feelings of satisfaction that tasting a fine wine tends to produce.[40] I take this to increase the plausibility of supposing that Greg and Michael share the same discriminative abilities. On this rival account of *Fine Wine*, such representations are even in Michael's case at least as plausibly an upshot of inferences or some other reliable transitions in thought as they are distinctively perceptual. This parallels my suggested account of the Norma/Alex contrast.

My third point is that moral properties can get represented in moral experience on the basis of diverse kinds of input. In vivid moral experiences the

39 I first used this (real life) example in Pekka Väyrynen, 'Some Good and Bad News for Ethical Intuitionism', *Philosophical Quarterly* 58 (2008), 489–511, at 499.

40 Two further points. First, as with 'see', so with 'taste': even if Michael doesn't need to stop to figure out that a wine is fine, it doesn't follow that the property of being a fine wine figures in the contents of his perceptual experience. Second, parallel reasoning in the intrapersonal variant of *Fine Wine* suggests that the phenomenological differences between Michael now and in the previous stage of his training aren't explained better as a difference in what properties he perceives than as a difference in the psychological immediacy and phenomenal integration of the same kind of transition in thought.

inputs often are perceptual. Consider not only *Cat* and *Tiananmen Square* but also seeing a hit-and-run, seeing one person push another off a footbridge, observing a street child living in abject poverty, hearing thuds and cries from a dark alley, or (recall Audi's example) hearing someone scream abuse at a conference speaker.[41] Many of us would represent what we are perceiving in such cases as wrong or bad, and our non-moral perceptual inputs would normally be integrated with certain phenomenological responses. But we seem to exercise the same kind of ability when the street child's circumstances as documented in a photograph strike us as bad and evoke affective empathy. And we seem disposed to respond in the same kind of way also when we imagine rather than perceive the thuds and the cries or the hit-and-run, or come to know about the abusive vulgarity or the footbridge incident by description rather than acquaintance.[42]

These different types of moral experience can of course differ in degree, such as with respect to how reliably they trigger the associated phenomenology and how vivid they are.[43] But they seem all to involve the same *kind* of representational ability. At the very least we have yet to see good reasons to individuate the relevant abilities to represent moral properties so narrowly that we have one ability for 'vivid' perceptual cases like *Cat* and another for others. Explaining these moral experiences seems to require only a general capacity to represent moral properties which is responsive to inputs from perception, imagination, supposition, and belief, and which can be psychologically immediate at least when those inputs are reliably and closely bound up with certain emotional and affective dispositions.[44]

The point isn't just that positing a distinctively perceptual representational ability adds no explanatory power with respect to cases like *Cat* and *Tiananmen*

41 Note that this model doesn't presuppose that F cannot figure in the contents of one's perceptual experience if one cannot discriminate Fs from not-Fs (e.g. stabbings that are bad from behaviorally identical stabbings that aren't). The relevance of counterfactual conditions to perceptual representation and knowledge is discussed in Cowan, op. cit. note 1, 183–7.

42 Inputs from imagination can be diverse in further ways. They may involve either mental imagery (the kind of quasi-perceptual process we engage in when, for instance, we close our eyes and imagine seeing something 'in the mind's eye') or propositional imagining.

43 Many philosophers are no doubt desensitized to descriptions of people being pushed off bridges, owing to the ubiquity of such cases in the trolley problem literature.

44 Robert Cowan suggested to me that this picture implies a disanalogy with aesthetic experience. It isn't clear that we can exercise aesthetic discriminative abilities in response to imagination or testimony in the way we can in response to perception. That may be right, but it remains unclear how this disanalogy would bear on (MP) until its source is clarified.

Square. (That was the lesson I took from *Fine Wine.*) It is also that explaining moral experiences that are based on non-perceptual inputs would require positing some distinct representational abilities. The rival to (MP) which I have put on the table explains a wider range of cases by appeal to one general representational mechanism. In this respect it offers a simpler and more unified account of moral experience than alternative packages built around (MP).

This claim is further strengthened when we note that (MP) involves an additional commitment that the parallel claims about many other higher level properties don't. If we perceive what the hoodlums are doing as wrong in *Cat*, we do so in part by perceiving the cat's distress. If we perceive the shooting of the citizens in *Tiananmen Square* to be unjust, we do so in part by perceiving the citizens as peaceably criticizing their government. These non-moral properties, though low*er* level relative to moral properties, are still high level. In typical arguments for high level content, perceptual content gets cognitively penetrated once over: background cognitive states influence perceptual experience by adding high level properties to a framework provided by low level properties. But at least in some cases (MP) would imply that perceptual content is cognitively penetrated twice over: first in the typical way, then by a further addition of moral properties to a framework provided by non-moral high and low level properties.[45] The parallel that many discussions of moral perception draw between moral properties and other higher level properties therefore looks exaggerated.[46] (MP) requires a more complex picture than such cases of high level perceptual content as causal relations and kind properties.[47]

---

45 Robert Cowan suggested to me that one might think instead that if one is presented with a non-moral perception of lower level properties involved in (say) *Tiananmen Square*, CP will involve the activation of a cluster of beliefs and dispositions associated with peaceful criticism, injustice, and so on, which then cognitively penetrate the perceptual experience together, only once over. But I see no principled reason to think this is the only way moral perception could occur if (MP) were true.

46 I have in mind the works cited in note 1 by Greco, Watkins and Jolley, Chappell, Cullison, and McBrayer. For a yet further complication, consider the emotional and affective dispositions manifested in the sorts of overall mental states that generate the phenomenal contrast argument for (MP) are shaped by background moral beliefs. Are these emotional and affective experiences themselves cognitively penetrable? If yes, this may require a yet further layer of CP. The cognitive penetrability of emotional experience is discussed e.g. in Vance, op. cit. note 8.

47 This point bears also on the relationship between perceptual experience and perception in the moral case. Even if (MP) were true, the double-layered CP structure might mean that the relevant perceptual experiences fail to be sufficiently directly related to the moral properties

In sum: There is a principled model of the role of perception in moral experience on which moral properties don't figure in the contents of perceptual experience. It treats the examples that are meant to make (MP) attractive as special cases of a more general and unified account of moral experience. As a simpler explanation of those cases, the model is preferable to (MP).

5. Two Objections

I'll now consider two objections to my argument against (MP) so far. The first objection is that the rival explanatory schema I have sketched carries some further commitments that make it no more unified or simple than (MP). The second is that some features of moral experience cannot be adequately explained without positing distinctively perceptual representations of moral properties.

One way to press the first objection is to claim that my rival explanation carries a commitment to 'cognitive phenomenology' which (MP) avoids. To acknowledge cognitive phenomenology is to acknowledge that propositional attitudes (believing, desiring, hoping, etc.) constitutively involve phenomenal character or 'what-it-is-likeness'.[48] Interpreting this worry isn't straightforward. My account of the phenomenological difference between Norma and Alex in *Cat* requires no particular belief or desire which Norma has but Alex doesn't.

Perhaps the worry is that what it is like for Alex to represent what the hoodlums are doing as bad differs from what that is like for Norma: Alex's experience involves a conscious inference but Norma's doesn't, and this is a difference in cognitive phenomenology. But this worry also fails to bite. On the explanation that I propose, both Norma and Alex come to represent what the hoodlums are doing as bad by an inference or some other kind of transition in thought. In Alex's case the phenomenological manifestation of the transition is 'figuring things out'.[49] In Norma's case the transition is phenomenologically reflected by certain affective states. In each case the transition needs to be involved in the relevant phenomenology only causally, not constitutively.

---

that they represent to count as moral *perception*. (Compare the discussion in Cowan, op. cit. note 1, 182–3.)

48 See e.g. Terence Horgan and John Tienson, 'The Intentionality of Phenomenology and the Phenomenology of Intentionality', in *Philosophy of Mind: Classical and Contemporary Readings*, ed. David J. Chalmers (Oxford: Oxford University Press, 2002), 520–32.

49 Nor do I need to suppose that what Alex's experience is like to him is a constitutive feature of making an inference rather than a phenomenal state that may be causally related to making an inference. Here consider e.g. the broad notion of inference in Sturgeon, op. cit. note 36.

The rival explanation to (MP) that I propose for these cases also doesn't require any phenomenology that we don't need to recognize anyway. The phenomenal contrast argument for (MP) outlined in section 3 is equally committed to whatever phenomenal character are possessed by the emotional and affective states invoked in Norma's case and the experience of 'figuring things out' invoked in Alex's case. The version of (MP) that strikes me as the most promising says that at least in some subjects, perceptual input triggers (directly or indirectly) an emotional or affective response owing to which some relevant moral background beliefs which have shaped those responses cognitively penetrate perceptual experience, adding some moral property to its content. The rival explanation that I propose treats those responses as reflecting an implicit habituated transition in thought from a perceptual input to a moral representation, owing to the way that the relevant emotional and affective dispositions have been shaped by some relevant background moral beliefs which connect non-moral inputs with moral classifications. Each view thus agrees that background moral beliefs can cognitively penetrate some or other aspect of the overall mental state of which a perceptual experience is a part. Each view also treats the capacity to represent moral properties as an acquired capacity that requires suitable training and background moral beliefs. My argument has been that treating this capacity as a perceptual capacity adds no explanatory power.

Characterizations of CP tend to be too weak to discriminate between these rival explanations. To say that a perceptual experience E is cognitively penetrated by a background belief, desire, or other cognitive state C is to say that the phenomenal character of E depends non-trivially on C. The relevant relation of dependence looks to be some suitably internal and mental relation of causal dependence. Now suppose we understand this relation counterfactually: if C didn't occur (antecedent to E), then E wouldn't occur.[50] This characterization is too weak to distinguish, at least in some cases, between the claim that C penetrates the content of the perceptual experience E and the claim that C somehow or other penetrates the overall mental state of which E is a part.[51] On the rival explanatory schema, the relevant background moral beliefs, the emotional and affective dispositions they shape, and non-moral perceptual inputs can stand in suitably internal and mental (indeed, contentual) relations that satisfy the counterfactual. Norma wouldn't have the emotional and affective

50 See e.g. Stokes, op. cit. note 8, whose survey of CP I'm following here.

51 To my knowledge, this problem hasn't been noted in the literature on CP.

responses she has to cases like *Cat* if her non-moral perceptual inputs didn't engage her background moral beliefs.[52] The first objection to my argument fails.

The second objection is that not everything about moral experience that needs explanation can be explained without positing distinctively perceptual representations of moral properties. In particular, one might claim that only (MP) can explain cases where something continues to strike the subject as bad or wrong in spite of stable background beliefs to the effect that it is neither.

One notable feature of perceptual representations is that they are highly insensitive to conflicting background beliefs. For instance, in the Müller-Lyer illusion one line is experienced as being longer than the other. Whatever other adjustments the background knowledge that the lines are really of the same length might effect in my cognitive system, it won't change my perceptual phenomenology. The same is true of certain perceptual illusions, such as the horse illusion (Figure 1), where no sense modality carries all of the information that gets represented in the experience.

Figure 1 The horse illusion

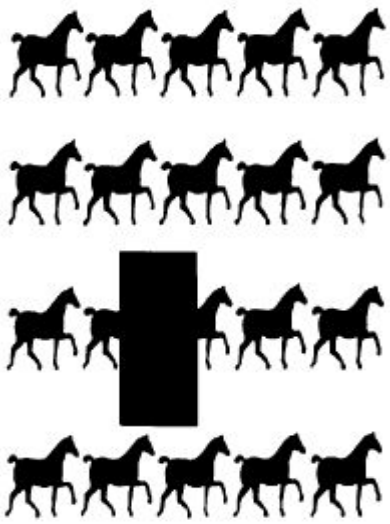

I cannot help seeing the black box as occluding the midsection of one very long horse, although I firmly believe, on the basis of all the other horse contours, that I should complete the occluded part of the picture with the front half of the horse on the left and the rear half of the horse on the right. We cannot help perceptually completing the occluded parts of objects in the simplest way possible, even if we have some conflicting firm beliefs about how we should

52 Again one simple model of how perceptual input can 'engage' background moral beliefs in cognitive processing is that it matches the antecedent of a background moral principle.

complete them.[53] If we completed this shape by forming a non-perceptual belief that the occluded shape is such and such, our completion shouldn't be insensitive to our other beliefs in the way it is.[54]

To defend (MP) on the grounds that moral experiences in the putative cases of moral perception are relevantly analogous to these experiences, it won't suffice to say that the subjects of these moral experiences cannot help having them. For instance, a utilitarian with empirical background beliefs that form a good utilitarian justification for what the hoodlums are doing in *Cat* might well not be able to help responding to it with affective empathy and revulsion or help representing it as wrong.[55] But my rival explanatory schema can explain this much. I can agree that in typical cases the simplest 'moral completion' is the moral representation that we tend to have. In *Cat*, for instance, the likeliest option given the subject's background beliefs and non-moral perceptual evidence is that what the hoodlums are doing is wrong, and this is so irrespective of whether its wrongness is represented perceptually or otherwise. My view is also consistent with the idea that we often cannot help experiencing things the way we do. The kind of transition in thought to which my explanation appeals can occur reliably as a psychologically immediate response to the relevant non-moral inputs, and the representational state it delivers needn't be a full-blown belief. The associated emotional or affective responses that make moral experiences feel so compelling also tend to be at best indirectly sensitive to other beliefs.[56] (Emotional retraining, though not impossible, is often very hard.) So the view that representations of moral properties aren't distinctively perceptual is consistent with the view that moral representations exhibit the kind of insensitivity to the other parts of the subject's cognitive system, such as conflicting background beliefs, which perceptual appearances tend to exhibit.[57]

53 See Bence Nanay, 'Perception and Imagination: Amodal Perception as Mental Imagery', *Philosophical Studies* 150 (2010), 239–54, at 243.

54 Various accounts of such 'amodal perception' are discussed in Nanay, op. cit. note 53. Cowan suggests that the phenomenon can be recruited in support of (MP); see Cowan, op. cit. note 1, 173–5. That would require that the background cognitive states which influence what additional material we represent on top of lower level perceptual content cognitively penetrate specifically the contents of perceptual experience rather than some other aspect of the overall experience. But that is the issue at stake.

55 Nor would (or should) she wish to help responding in this way – but that is another story.

56 Insofar as the proponents of (MP) take affective states to facilitate CP, they incur the same commitments regarding how insensitive those states are to the other parts of the subject's cognitive system, to what extent such insensitivity makes those states irrational, and so on.

57 I'm not (yet) convinced that the relevant moral experiences really are as insensitive to

So what more does this second objection require? Sometimes when we cannot help experiencing an action as bad or wrong, we nonetheless pause to think whether anything supports a morally charitable account of the action. On my account this would be to double-check a psychologically immediate transition in thought. (Compare snap judgments, which get double-checked fairly often even when they are fairly reliable.) According to (MP), by contrast, we would be checking whether a perceptual appearance is illusory (much as one might check the perceptual conditions when doubtful of one's perceptual appearances regarding color, distance, or the like). One might then wonder whether the latter hypothesis is more compelling.

One might pursue the analogy with perceptual illusions by appealing to examples of people who cannot help experiencing something as bad or wrong in spite of having settled that it isn't bad or wrong and having tried cognitive therapy and emotional retraining. Consider, for instance, a person who continues to be disgusted by homosexual affection long after having settled that there is nothing disgusting or immoral about such behavior.[58] However, it is standard to think that such cases of disgust are irrational insofar as they are recalcitrant, and neither of the competing explanations has grounds for denying this. And since each appeals to emotional and affective dispositions, each is committed to a disconnect between the subjects' background moral beliefs and their moral experiences. (MP) implies that if a property like being immoral figures in the contents of a perceptual experience of homosexual affection, it must have got there in some other way than cognitive penetration by background moral beliefs facilitated by suitable emotional or affective dispositions. And my rival explanatory schema implies that the predominant causal influence on which responses are psychologically immediate and phenomenally integrated with non-moral perceptual input is the recalcitrant disgust rather than the subject's background moral beliefs and the transitions in thought they support. So the example of recalcitrant morally laden disgust by itself privileges neither explanation.

---

conflicting background beliefs as perceptual experiences tend to be. This is an empirical issue that I cannot discuss here properly. But there is anecdotal evidence that moral appearances often respond to background beliefs. For instance, many people who used to see eating meat as perfectly permissible but no longer do come to represent instances of eating meat differently as their beliefs change. Thanks to Jacob Sparks for discussion here.

58 Thanks to Preston Werner for suggesting this example to me.

In fact, there may be reason to prefer my hypothesis about what is going on when we 'pause to think'.[59] When we check to see whether a perceptual appearance is illusory, we do so by making other observations, such as looking more closely to see whether what looks like a barn is just a facade or measuring whether the lines really are different lengths. But neither usually does the trick in the moral case. When I cannot help experiencing something as wrong and pause to think about it, I'll usually try to imagine the situation from a different perspective or find (or remind myself of) some other morally relevant feature that I may have overlooked. At least the latter doesn't look normally to happen with perception. But if checking to see whether a moral appearance is illusory were a matter of checking a transition in thought, one would expect it precisely to involve scrutinizing whether the transition holds in view of a fuller range of potentially relevant features grasped from the most apt perspective one can take.

6. Conclusion

I have argued that the phenomenology of the sort of moral experiences we have in response to perceptual input can be well explained without thinking that moral properties figure in the contents of perceptual experience. Treating experiential representations of moral properties as resulting from trained implicit transitions in thought yields a simpler and more unified explanation. But recall that a thesis like (MP) might be motivated not only on phenomenological but also on epistemological grounds.[60] If anything gives us knowledge about the external world, perception does. So might the existence of moral perception help us with the difficult project of developing a plausible moral epistemology?

If I'm right that phenomenal contrast arguments that rely on the cognitive penetration of perceptual content by background moral beliefs fail to support

59 Many thanks to Jacob Sparks for suggesting the line of argument in this paragraph.

60 A third possible type of argument for (MP) is based on the metaphysics of moral properties. Chappell, op. cit. note 1, argues that at least some (moral) patterns are properties and at least some (moral) pattern recognition is perception, so there is moral perception. But the contrast he draws between perception and inference ignores various distinctions drawn in the course of my argument. Heather Logue suggests that if instantiations of some aesthetic properties are token-identical with their non-aesthetic base properties and we can visually experience the former, then we can thereby visually experience the latter; see Logue, 'Can We Experience Aesthetic Properties?', this volume. Whether this generalizes to moral properties depends partly on the issue raised in note 22: can moral properties be perceived if not all of the typically complex non-moral grounds of their instantiations are perceivable? Otherwise nothing in this paper rules out the possibility of defending (MP) on the basis of some controversial metaphysics of moral properties which requires independent argument.

(MP), then delivering on the epistemological promises of (MP) requires a different kind of argument for (MP) itself. Such arguments are yet to be developed. But even setting that aside we should be cautious to ascribe (MP) any significant epistemological benefits. Perceptual experiences of some property are in themselves an antidote to skepticism only if they provide an *epistemically independent* source of justification. This is to say, roughly, that perceptually representing O to be F must be able to confer epistemic justification for believing that O is F independently of whether one has (non-perceptual) justification for some belief B that must be justified for the belief that O is F to be justified. For instance, memory will be an epistemically dependent source of justification if it isn't a positive source of epistemic justification in its own right, but can only transmit justification from non-memorial sources of justification.

Insofar as moral perception would involve cognitive penetration of perceptual content by background moral beliefs or emotions, moral perception seems not to provide an epistemically independent source of justification.[61] For it seems that insofar as one's perceptual moral experiences would be sensitive to one's background moral beliefs or emotions, any justification that moral perception could confer on moral beliefs would be mediated by the justification for the relevant background states. And CP seems to abide by the 'garbage in, garbage out' principle: if the penetrating cognitive states are themselves epistemically unjustified, then so would be the beliefs formed on the basis of the perceptual experiences which those cognitive states penetrate. So perceptual moral experience wouldn't confer justification for moral beliefs independently of whether the relevant background states are justified. But this is just to say that perceptual moral experience would be only an epistemically dependent source of justification.[62] So it seems reasonable to suppose that moral perceptions would

61 Here I follow Cowan's argument that (MP) doesn't provide an epistemically independent source of justification; see Cowan, op. cit. note 1, 187–91. After this paper was completed, a fuller argument has been developed in David Faraci, 'A Hard Look at Moral Perception', *Philosophical Studies* 172 (2015), 2055–72. For a reply, see Preston Werner, 'Moral Perception without (Prior) Moral Knowledge', *Journal of Moral Philosophy* (forthcoming). Note that I'm not supposing that CP *as such* is epistemologically problematic; compare Siegel, op. cit. note 8, and Vance, op. cit. note 8.

62 For all I say, moral perception might not be epistemically dependent on versions of (MP) that don't appeal to CP (see note 19). Versions of (MP) which do appeal to CP require not just that the moral *concepts* one has cognitively penetrate perceptual content, but also that some moral *beliefs* do. Mere conceptual CP isn't enough to explain why different things tend to strike people as wrong as background moral beliefs vary, at least in subjects whose emotional and affective dispositions aren't significantly out of line with their background moral beliefs.

all be epistemically dependent. No doubt this claim needs a fuller defense. But if it is right, then moral perception would make no independent contribution to a plausible moral epistemology even if it did occur.

Talk of 'seeing' things to be right or wrong can still come naturally and won't normally mislead. What we have yet to see are good reasons to think that moral perception occurs in the literal sense of (MP). Recent arguments in the philosophy of perception for thinking that various kinds of higher level properties figure in the contents of perceptual experience aren't plausible when applied to moral properties. Those phenomenological features of moral experience which seem to provide the most promising kind of phenomenal contrast argument for (MP) can be adequately explained without appeal to (MP). Nor would (MP) seem to help us with the project of developing a plausible moral epistemology. Moral experience is a complex and rich phenomenon with various features that I haven't considered here. So I cannot rule out the possibility that other arguments for (MP) might fare better. But on the present showing doubts about the existence of moral perception look robust.[63]

This strengthens the case that moral perception won't be epistemically independent insofar as (MP) appeals to CP.

63 Thanks to audiences at University of Leeds, the Evaluative Perception conference at University of Glasgow, and the Evaluative Perception workshop at Cardiff University for helpful discussions of this paper. Many thanks also to Anna Bergqvist, Robert Cowan, David Faraci, David Killoren, Heather Logue, Jacob Sparks, Preston Werner, and an anonymous referee for valuable written comments.