

## Durham Research Online

---

**Deposited in DRO:**

04 September 2018

**Version of attached file:**

Accepted Version

**Peer-review status of attached file:**

Peer-reviewed

**Citation for published item:**

Faraci, David (2014) 'A hard look at moral perception.', *Philosophical studies.*, 172 (8). pp. 2055-2072.

**Further information on publisher's website:**

<https://doi.org/10.1007/s11098-014-0397-6>

**Publisher's copyright statement:**

This is a post-peer-review, pre-copyedit version of an article published in *Philosophical studies*. The final authenticated version is available online at: <https://doi.org/10.1007/s11098-014-0397-6>

**Additional information:**

---

**Use policy**

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

## A HARD LOOK AT MORAL PERCEPTION

*David Faraci*

This paper concerns what I take to be the primary *epistemological* motivation for defending moral perception.<sup>1</sup> Offering a plausible account of how we gain moral knowledge is one of the central challenges of metaethics. It seems moral perception might help us meet this challenge. The possibility that we know about the instantiation of moral properties in something like the way we know that there is a bus passing in front of us raises the alluring prospect of subsuming moral epistemology under the (relatively) comfortable umbrella of perceptual or, more broadly, empirical knowledge.<sup>2</sup>

The good news on this front is that various combinations of metaethical positions and theories of perception arguably have the potential to vindicate moral perception (though I won't do much to defend this claim here).<sup>3</sup> The bad news, I'll argue, is that such vindication is only half the battle where this epistemological goal is concerned. And the other half of the battle is unlikely to be won.

### 1. The Other Half of the Battle

Nearly everyone agrees that we can perceive “low-level” properties including spatial properties, color, shape, motion and illumination. It is more controversial whether we can perceive “high-level” properties like *being a bus* or, relevantly for our purposes, *being wrong*.<sup>4</sup> I will assume that we can perceive at least some high-level properties, since if we cannot, the question of moral perception's

---

<sup>1</sup> As opposed to, for instance, phenomenological motivations. For a recent critique of moral perception along these lines, see Väyrynen (m.s.).

<sup>2</sup> I focus on perception of “thin” moral properties—good, bad, right, wrong, etc. It's controversial whether perception of thicker moral properties counts as moral perception. See, e.g., Väyrynen (2013). Even if it does, one could use thick moral perception to develop a purely perceptual moral epistemology only if thin moral knowledge is grounded in thick moral knowledge. I highly doubt this, though I cannot argue against it here.

<sup>3</sup> In addition to other works discussed herein, Väyrynen (m.s.) mentions defenses of moral perception from Audi (2013), Blum (1994), Chappell (2008), Cowan (forthcoming), Cullison (2010), Cuneo (2003), DePaul (1993), Greco (2000), McNaughton (1988), and Watkins and Jolley (2002).

<sup>4</sup> The list of low-level properties is borrowed from Siegel (2011). Siegel's is a list of “non-kind” properties (as opposed to “kind” properties). This is the same distinction—or, at least, non-kind properties and low-level properties are the same, and kind properties are a subset of the high-level properties. The low-/high-level phrasing is borrowed from Werner (forthcoming).

epistemological implications is moot.<sup>5</sup> I further assume that perceptual experiences can be distinguished (for one) by their phenomenal content.<sup>6</sup>

A few more preliminaries: “Perception of X” is ambiguous between (for one) “perception explained by X” and “perception grounded in an experience with distinctively X-like phenomenal content.”<sup>7</sup> Respectively, these are the senses in which Lois Lane does and does not perceive Superman when she sees Clark Kent. Given that my interest is epistemological, I will be using “perception of X” exclusively in the latter, intensional sense. As I’m using it, “perception of X” is also not a success term—one can have a false perception of X.

When I speak of X-like experiences or experiences *as of X*, I mean experiences with distinctively X-like phenomenal content. Where X is in some class C, I will sometimes talk about C perceptions or experiences—e.g., *being wrong* is a moral property, so a perception of wrongness is a moral perception, and an experience *as of wrongness* is a moral experience.

I will not offer a full theory of what it is for phenomenal content to be X-like. Roughly, the idea is that the phenomenal content of an X-like experience bears an important kind of similarity relation to certain recognizable features of X.<sup>8</sup> For example, the phenomenal content of a square-like experience—an experience *as of a*

---

<sup>5</sup> For grammatical simplicity, I will sometimes talk of properties as objects—e.g., say that Norm perceives some behavior, rather than perceiving that something instantiates the property of behaving in some particular way, or that Vera perceives a bus, rather than that she perceives that something instantiates the property of being a bus. This should not be taken as a leap to assuming that we can perceive objects independently of their properties (though I do not deny this, either.)

<sup>6</sup> For defense—as well as arguments that this claim is compatible with views it is typically taken to be in tension with, such as naïve realism—see Siegel (2011, chap. 2).

<sup>7</sup> What is the difference between perception of X and experience *as of X*? First, I take it perceptions of X need to have a certain kind of “cognitive basicness.” Presumably, both optimists and skeptics about mental state perception would grant that there can be experiences *as of anger*. Their disagreement would be over whether such experiences ground or are grounded *in* judgements about anger. Second, for the optimist, a perception of anger is at an intermediate “cognitive level” between an experience *as of anger* and an anger-judgement. This allows for the possibility of having an experience *as of X* without perceiving X—say, if I know I’m in a context where there are likely to be a lot of fake Xs. This bit about “cognitive level” is intentionally left vague. It serves only to make certain points more intuitive (I hope); nothing substantive is meant to hang on it.

<sup>8</sup> Which features? Answering this question is difficult. For instance, it can’t just be those features that explain the experience. For, intuitively, it seems that a quotidian experience of water is water-like, but not H<sub>2</sub>O-like. For the same reason, it can’t be those features that we associate with the object of experience—at least not given that some people know that water is H<sub>2</sub>O. I suspect the answer has something to do with the features that allow the experiences to meet criteria set by our concepts, but I won’t explore this further here.

*square*—bears certain structural similarities to key recognizable features of actual squares. Arguably, an experience *as of a square* has the same phenomenal content as a particular kind of experience *as of a rectangle*—i.e., one where the rectangle seems to be equilateral. I take it that someone who lacks the concept ‘rectangle’ might have an experience *as of a square* without having an experience *as of a rectangle*. Thus, “experience *as of X*” is also intensional.

Begin with a simple case. Norm and Vera are driving through the country. A bus cuts them off, and Vera becomes irate. Norm sees that she is angry, and tries to comfort her.

Some will accept that Norm really does perceive that Vera is angry. Perhaps before Norm knew Vera as well as he does now, the phenomenology of seeing her reaction would have been different from how it is now. The best explanation for this “phenomenal contrast” might be that Norm is now *perceiving* her anger.<sup>9</sup> Others will deny this. They might hold, instead, that Norm perceives only that Vera is scowling, *judging* that she is angry on the basis of that perception.<sup>10</sup>

Suppose a skeptic about mental state perception is trying to account for Norm’s knowledge that Vera is angry. On the skeptical view just described, there is an important epistemological relationship between Norm’s judgement that Vera is angry and his perception of her scowl. Norm infers (or in some other way “shifts”<sup>11</sup>) *from* perception of that scowl *to* a judgement about her mental state. For this inference to have epistemic merit, Norm must possess some background knowledge about connections between Vera’s facial expressions and her mental states.<sup>12</sup>

Crucially, the epistemological story for *optimists* about mental state perception is unlikely to differ much from the skeptic’s. Assuming Norm isn’t telepathic, his perception of Vera’s mental state is clearly grounded in his experience of her behavior: He perceives anger *because* he has an experience *as of a*

---

<sup>9</sup> This argument mimics the phenomenal contrast arguments Siegel (2011) deploys to defend perception of various complex properties. Werner (forthcoming) deploys such an argument to defend moral perception. Werner’s argument comes up again in §2.3.

<sup>10</sup> I’m assuming here that Norm might judge Vera to be angry partly on the basis of her behavior. This is roughly in line with the “theory theory” about judgements concerning others’ mental states. In contrast, according to “simulation theory,” Norm might judge that Vera is angry after running through empathetic processing starting with the same input (the bus cutting them off). I set simulation theory aside because it doesn’t seem compatible with mental state perception, which I’m focusing on for illustrative purposes. (For one thing, Vera’s mental state plays no role in the simulation theoretic explanation of Norm’s judgement.) For an overview of simulation theory as a reaction to the theory theory, see Gordon (2009).

<sup>11</sup> I take no position here on whether sub-personal processing can count as inferential.

<sup>12</sup> Or, at least, beliefs that themselves have some epistemic merit (e.g., are justified).

scowl. For this relation to have epistemic merit, Norm must again possess some background knowledge of a relation between scowls and anger.<sup>13</sup>

In terms used in the current literature on perception, we may say that optimism about epistemically successful mental state perception relies on *cognitive penetration*<sup>14</sup>—Norm’s perception of Vera’s anger must be cognitively penetrated by background knowledge<sup>15</sup> that scowls indicate anger (in Vera).<sup>16</sup> I will say that Norm’s perception of her anger is *mediated* by this knowledge.<sup>17</sup>

---

<sup>13</sup> Since I completed the paper, a number of people have suggested that this knowledge requirement is too strong. Perhaps Norm is justified in believing that Vera is angry simply because perception is a basic source of justification. And perhaps the explanation for the reliability of that perception is just the fact of some relationship between Vera’s facial expressions and her mental states. What I want to insist on, though, is that in order for Norm’s perception to ground *knowledge*, whatever explains the reliability of that perception must in some sense be *available* to him. If the explanans can be relevantly available to Norm without his believing it, then perhaps the knowledge requirement is too strong. But I am confident that shifting from a knowledge requirement to an availability requirement would make little difference to my arguments herein.

Of course, some—e.g., externalists or coherentists—might reject even the availability requirement. But then Norm’s is just the sort of case that is frequently used to challenge externalism: It seems implausible that an agent can know something when the reliability of his route to that knowledge is deeply mysterious (at least to him). See, e.g., Bonjour’s (1980) case of Norman the Clairvoyant. As for coherentism: If Norm is sensible, he will recognize the need to explain the reliability of his perception. It seems unlikely that his perception-based belief that Vera is angry would cohere with his belief that the reliability of his perception of her anger is deeply mysterious. For application of externalism and coherentism to *moral* knowledge see, e.g., Shafer-Landau (2005) and Sayre-McCord (1996), respectively.

<sup>14</sup> I first encountered this phrase in Siegel (2011), though I owe my use of it Werner (forthcoming), who uses it to make points similar to those made here about why moral perception alone can’t vindicate a purely perceptual moral epistemology.

<sup>15</sup> Since I’m only interested in epistemically successful cases, I’ll continue to refer to background knowledge, rather than beliefs, though arguably one can have perceptions that involve cognitive penetration even by false beliefs.

<sup>16</sup> This is a case of what we might call *augmenting penetration*, where background knowledge leads an experience *as of Y* to ground perception of X. This can be contrasted with what we might call *undercutting penetration*, where background knowledge prevents an experience *as of Y* that otherwise *would* ground perception of X from doing so. For instance, if Norm learned Vera was an android, he might cease seeing her as angry. Except where noted, discussion herein is limited to augmenting perception.

<sup>17</sup> I trust my decision not to have my paper riddled with talk of penetration is self-explanatory. Note that it may also be possible for perceptions to be mediated by non-cognitive attitudes. This might be relevant for those who take certain affective states to themselves be perceptions of moral properties. See, e.g., Oddie (2009) and McBrayer (2010a; 2010b). My arguments apply as well to affective states as to cognitive ones. If Norm’s perception of Vera’s anger is mediated by

So, in developing an epistemology of (others') mental states, both optimists and skeptics about mental state perception have to account for certain background knowledge—e.g., explain how we know that certain behaviors implicate certain mental states. If the optimist hopes to develop a *purely perceptual* mental state epistemology, it follows that she must show that the relevant background knowledge is *itself* perceptual.

At least for the sake of argument, I grant the optimist about moral perception that there is perception of moral properties. However, as with perception of others' mental states, I argue that moral perception would have to be mediated by background knowledge of relations between moral and non-moral properties, which I refer to as *moral bridge principles*. The other half of the battle, for champions of a purely perceptual moral epistemology, is thus to show that knowledge of moral bridge principles is *itself* perceptual. I argue that this cannot be the case. It follows that even if there is moral perception, we almost certainly must give up hope of developing a purely perceptual moral epistemology.

## 2. Against Unmediated Moral Perception

Sarah McGrath (2004) argues, much as I have, that we shouldn't be content with showing that moral perception exists in *some* sense. Rather, as she puts it, we should hope to learn whether we have any moral knowledge *by* perception. She thinks we sometimes do have such knowledge. Consider one of her examples:

[S]uppose Alice believes that homosexuality is wrong, and that she believes this because she has learned that the scriptures say that homosexuality is wrong, and believes that the scriptures are authoritative on this matter. But then she gets to know a couple, Bob and Chuck, who live next door. She gradually comes to believe that it is not wrong for them to be in this relationship. It isn't that she comes to believe this because she detects some non-moral features that she believes are sufficient for having a morally permissible relationship—she doesn't change her mind because she learns that these people are monogamous, or that they prioritize each other's needs. According to the moral principles that Alice believes, these sorts of non-moral facts would be insufficient for having a relationship that is morally permissible. Alice simply comes to

---

an affective state, rather than a cognitive one, we would still require a story about how that affective state epistemically vindicates the grounding relation between Norm's experience *as of a scowl* and his perception of anger: Norm would still need to "know," in some sense, that scowls implicate anger.

believe that there is nothing wrong with this relationship, on the basis of her acquaintance with Bob and Chuck. (McGrath 2004, 224–225)

McGrath denies that Alice perceives the permissibility of Bob and Chuck’s relationship “because she detects some non-moral features that she believes are sufficient for having a morally permissible relationship.” It isn’t clear, though, that McGrath can simply stipulate this. Suppose Norm claims to know by perception that Vera is angry, but denies that this is because he detects aspects of her behavior he believes to indicate that she’s angry. Unless he is telepathic, this is impossible. Here’s a more detailed version of the argument from §1:

P1. [**Mediation**] If perceptions of X are grounded in experiences *as of Y*, then perceptions of X produce knowledge only if they are mediated by background knowledge of some relation between X and Y.<sup>18</sup>

P2. In epistemically successful cases, and barring telepathy, perceptions of others’ mental states are grounded in experiences *as of behavioral cues*.

C1. Therefore, in epistemically successful cases, and barring telepathy, perceptions of others’ mental states are mediated by background knowledge of some relation between certain behavioral cues and the relevant mental states.

P3. Norm’s perception of Vera’s anger is a perception of another’s mental state.

C2. Therefore, barring telepathy, Norm’s perception of Vera’s anger is epistemically successful only if it is mediated by relevant background knowledge—presumably, that her behavior suggests she is angry.

If we are to grant McGrath that Alice’s case, as described, is possible, something must block an analogous line of reasoning with respect to moral perception. In §2.1, I argue for the moral analogue of P2: Epistemically successful

---

<sup>18</sup> This is not to say that the perception and the experience must actually be separate. It might be that certain *aspects* of an overall experiential state with one phenomenal character are grounded in other aspects of that same overall experiential state with a *different* phenomenal character. I am also not assuming that there is only one unique phenomenal content that counts as being *as of permissibility*—only that experiences *as of permissibility* and experiences *as of the relevant base properties* have different phenomenal contents. More on this in §2.2.

perceptions of moral properties are always grounded in non-moral experiences. In §2.2, I consider an objection and revise the argument to accommodate it. This illuminates a path of resistance for the optimist about epistemically successful unmediated moral perception (from here I frequently drop the “epistemically successful” qualifier). In §2.3, I offer some reasons to think that path is closed.

### *2.1 An Argument that Moral Perception is Grounded in Non-Moral Experience*

Suppose Norm accepts the above argument, but claims that he *is* telepathic, and thus that his mental state perceptions are not grounded in behavioral experiences. Here’s one way we might test his claim: Construct a counterfactual scenario much like the one discussed in §1, but make the minimal changes necessary to remove the anger: Norm and Vera are driving along, a bus cuts them off, Vera scowls, shouts, etc., and Norm sees this. But it’s all a setup: Vera knew the bus was going to be there, and is behaving just as she would were she actually angry (she’s a very good actor) so as to test Norm’s claim that he is telepathic. If Norm perceives anger just as he did in the original case, this suggests that his perception of anger is explained by his experience of her behavior.

Nick Sturgeon deploys a similar counterfactual test in a famous exchange with Gil Harman, on the topic of moral explanations of moral judgements.<sup>19</sup> Harman offers a case in which someone comes upon a group of children torturing a cat for fun. The onlooker—call him Sam—perceives that the children are doing something wrong. Harman argues that Sam’s moral perception can be fully explained by Sam’s non-moral experience in combination with his background moral beliefs, and thus the wrongness itself does no explanatory work.

Sturgeon rebuts with a counterfactual test: He introduces a counterfactual scenario that differs from the original to the minimal extent necessary to remove the wrongness, then asks whether Sam’s experience responds to that change. If Sam no longer experiences what he sees as wrong, it seems his experience is (perhaps directly) responsive to the presence (absence) of wrongness. Sam passes the test. For were we to alter the situation’s moral qualities—say by having the children pet the cat instead of torturing it—Sam’s experience would undoubtedly change.

It is not hard to see why we might be suspicious of this result. I characterized the test in terms of constructing counterfactual scenarios that differ “to the minimal extent necessary to remove the wrongness.” Unfortunately, moral supervenience entails that “the minimal extent necessary” requires altering the base properties along with the moral ones. And so it is possible that it is the change *in*

---

<sup>19</sup> The exchange begins with Harman (1977). Sturgeon responds in his (1986). This went back-and-forth a bit; next is Harman (1986).



those base properties, rather than the change in the moral properties themselves—i.e., the fact that the cat is now being pet rather than tortured—that explains the change in Sam’s experience.<sup>20</sup> What we may well have discovered, then, is not that Sam passes an important test, but rather that because moral properties supervene, the test can’t function properly in moral cases.<sup>21</sup>

Luckily, we can do better. For our concern isn’t whether moral *properties* explain moral perceptions, but rather whether moral perceptions are grounded in non-moral *experiences*. Thus, our question in this case is not whether wrongness itself explains Sam’s moral perception, but whether Sam’s experience *as of certain non-moral features of cat-torture* grounds that perception.

To see why the counterfactual test is still useful, return briefly to Norm. The fact that Norm continues to see Vera as angry in the counterfactual case strongly suggests that her behaviors explain his perception of anger. This suggests the principle:

**Explanation** If perceptions of X track the presence of Y even in the absence of X, the best explanation is that the presence of Y explains perceptions of X.

**Explanation** doesn’t seem to cut much ice in the moral case when we substitute some base property for Y, because it isn’t possible for Y to be present while X (the moral property) is absent.<sup>22</sup> But there is another way to think about the results of this test in terms of the phenomenal content of Norm’s experience, rather than in terms of what, externally, explains his perception. It seems safe to assume that when Norm sees Vera’s behaviors, he has an experience *as of those behaviors*. The fact that Norm continues to see Vera as angry in the counterfactual case strongly suggests that his experience *as of her behavior* grounds his perception of her mental state. This motivates the more specific principle:

**Grounding** If perceptions of X track experiences *as of Y* even in the absence of X, the best explanation is that perceptions of X are grounded in experiences *as of Y*.

Crucially, while we can’t remove the wrongness from Sam’s case without removing certain non-moral features, we *can* remove the wrongness while retaining something that will produce in Sam an experience *as of those non-moral features*.

---

<sup>20</sup> This assumes that the base properties and moral properties aren’t, in fact, the same properties. I return to this issue in §2.2.

<sup>21</sup> David Slutsky (2001) argues that the test can be improved, by considering cases where the relevant non-moral base properties are causally inefficacious. There are a number of parallels between his argument and my own.

<sup>22</sup> Though again see Slutsky (2001).

That is, we can construct a *convincing fake*: an animatronic cat, holographic fire, etc., etc. Surely, Sam would continue to experience what he sees as wrong.<sup>23</sup> **Grounding** thus suggests that Sam’s moral experience in this case is grounded in his non-moral experience. If the counterfactual test gives the same results in *all* such cases—if Sam’s moral perceptions consistently track certain non-moral experiences even when the relevant moral properties are absent—the best explanation seems to be that his non-moral experiences consistently ground his moral perceptions.<sup>24</sup>

This brings us back to McGrath’s case: Suppose it turns out that Alice’s neighbors are in fact not a gay couple, but rather (unbeknownst to Alice) two heterosexual method actors preparing for an upcoming performance. Surely, Alice’s views about gay marriage would shift just as they did in McGrath’s original case. And surely this is because Alice’s perception of permissibility is grounded in her experiences *as of certain non-moral features of Bob and Chuck’s relationship*.

Crucially, none of this relies on the details of the particular cases under consideration. Something similar seems true about *every* case involving *every* moral property. Schematically, take any epistemically successful case in which person P perceives some moral property M in the presence of some set of non-moral properties N. I submit that in *every case*, were we to remove M, convincingly fake some subset of N, and change nothing else, P would falsely perceive M. If we can indeed expect this in every case, we should conclude that all epistemically successful moral perception is grounded in non-moral experience.<sup>25</sup>

---

<sup>23</sup> One incredible alternative would be that Sam judges the cat-torture to be fake on the basis of failure to perceive wrongness!

<sup>24</sup> I’m simplifying slightly. It is controversial whether the experience caused by the fake cat-torture has the same phenomenal content as that caused by the real. However, given that the fake cat-torture is *convincing*, we may suppose the contents of the experiences caused by the real and fake cat-tortures are in principle indistinguishable *to Sam*. The counterfactual test thus suggests that Sam’s moral perception tracks experiences that are indistinguishable from experiences *as of cat-torture* better than they track the presence of wrongness. The best explanation for this seems to be that his falsidical moral perception is grounded in a non-moral experience that is indistinguishable from an experience *as of cat-torture*. And the best explanation for *this* seems to be that his veridical moral perception is grounded in his experience *as of cat-torture*. For simplicity’s sake, I will continue to speak as though the two have the same phenomenal contents. For those who reject this, “experience *as of cat-torture*” can be taken to refer both to cat-torture-like experiences and those that are indistinguishable from cat-torture-like experiences. This makes no substantive difference to my arguments.

<sup>25</sup> The idea that this concerns every case is crucial. The point of the counterfactual test is *not* that if I can falsely perceive X, then perceptions of X are grounded in experiences *as of some Y*. I accept that there can be false unmediated perceptions, perhaps even of low-level properties.

We now have the makings of an argument that all epistemically successful moral perception is mediated:

P1. [**Mediation**] If perceptions of X are grounded in experiences *as of Y*, then perceptions of X produce knowledge only if they are mediated by background knowledge of some relation between X and Y.

P4. [**Grounding**] If perceptions of X track experiences *as of Y* even in the absence of X, the best explanation is that perceptions of X are grounded in experiences *as of Y*.

P5. In epistemically successful cases, perceptions of moral properties track non-moral experiences even in the absence of the relevant moral properties.

C3. Therefore, in epistemically successful cases, moral perception is mediated by background knowledge of some relation between moral and non-moral properties—i.e., of moral bridge principles.

Assuming the argument is valid, my opponent has three options: First, challenge **Mediation** by arguing that a perception of X grounded in an experience *as of Y* might be epistemically successful even without being mediated by relevant background knowledge. Second, challenge **Grounding** with a competing explanation for the results of the counterfactual test. Third, challenge P5 by denying that the results of the counterfactual test would be as I've suggested.

**Mediation** is highly intuitive: Suppose A's perception of X is grounded in an experience *as of Y*. But suppose A also has no reason to think there's any relation between X and Y. In that case, there seems to be no way A's experience *as of Y* could produce *knowledge* about X. For the grounding relation in question could have no epistemic merit.<sup>26</sup> Skipping to P5: Though this is technically an empirical matter, I have no doubt that the results of the counterfactual test would be as I've suggested. I thus set the possibility of rejecting these results aside. This leaves the option of proposing a competing explanation for the results of the counterfactual test. In §2.2, I consider one such explanation and revise the argument to accommodate it. In §2.3, I offer reasons for thinking the argument is sound, even so altered.

---

<sup>26</sup> Though see note 13, above.

## 2.2 An Objection to and Revision of the Above Argument

Return to Norm and Vera, but this time, focus on Vera's perception of the bus. Vera's perception of the bus is presumably grounded in an experience *as of a bus*. According to **Mediation**, it follows that Vera's perception of the bus (assuming it is epistemically successful) must be mediated by background knowledge of a relation between the bus and itself. We could accept this, but it seems rather odd. First, it might seem that if there are *any* cases of unmediated perception, Vera's will be one. Certainly, her perception seems more "direct" than Norm's, Sam's or Alice's. So perhaps we should revise **Mediation** to exclude cases where X and Y are identical, as follows:

**Mediation\*** If perceptions of X are grounded in experiences *as of Y*, and X and Y are non-identical, then perceptions of X produce knowledge only if they are mediated by background knowledge of some relation between X and Y.

**Mediation\*** would allow some to resist my argument that all moral perception is mediated—reductive naturalists, for instance. But **Mediation\*** is too weak, for there are cases where epistemically successful perception of X *does* require mediation by knowledge of a relation between X and Y, even though X and Y are identical. To see this, consider the case of Cliff who, after stepping in a puddle of water, cries, "Ack! I've just gotten H<sub>2</sub>O all over my new shoes!"

Distinguish two versions of this case: Cliff<sub>1</sub> has the peculiar ability to see things at the molecular level, and thus experiences the stuff he steps in as groups of hydrogen and oxygen atoms. Cliff<sub>2</sub>, on the other hand, experiences that stuff as most of us would—as colorless, wet, etc. Insofar as Cliff<sub>2</sub> perceives H<sub>2</sub>O, it seems clear that his perception is grounded in his experience *as of water* (and thus mediated by background knowledge that water is H<sub>2</sub>O). This is true even assuming *being water* and *being H<sub>2</sub>O* are the same property.

The counterfactual test can help us determine which case we are dealing with. Suppose Cliff claims to be Cliff<sub>1</sub>, able to directly see molecular structure. Dubious, we head over to Twin Earth and get ourselves some XYZ. Given its causal properties, we can safely assume that XYZ will produce in Cliff an experience *as of water*—*unless* he has Cliff<sub>1</sub>'s "molecular sight."<sup>27</sup>

We then recreate the original case, except this time Cliff steps in a puddle of XYZ. If he falsely perceives H<sub>2</sub>O, this gives us reason to believe that he is Cliff<sub>2</sub>—his perception of H<sub>2</sub>O is grounded in his experience *as of water*. Despite his protestations, his perception is thus almost certainly mediated—again, presumably

---

<sup>27</sup> This is, of course, a reference to Putnam (1975).

by background knowledge that water is H<sub>2</sub>O. (One more piece of intuitive evidence: Consider Cliff<sub>3</sub>, who is exactly like Cliff<sub>2</sub> in every way, except Cliff<sub>3</sub> lacks the belief that water is H<sub>2</sub>O, though he possesses both concepts. Surely Cliff<sub>3</sub>—whether confronted with XYZ or with H<sub>2</sub>O—wouldn't perceive H<sub>2</sub>O.)

Cliff's case shows that **Mediation**\* is too weak because even though water and H<sub>2</sub>O are the same property, water-like experiences don't have the same phenomenal content as H<sub>2</sub>O-like experiences. This suggests that we should instead consider cases where two experiences *do* have the same phenomenal content. This brings us back to Vera's perception of the bus. Here are three true claims about Vera: First, Vera perceives a bus. Second, Vera has an experience *as of certain low-level properties—shape, color, etc.* Third, if there were no bus, but Vera had an experience *as of those low-level properties*—perhaps she was cut off by a motorcyclist wearing a very convincing bus-façade—she would falsely perceive a bus. Given this, **Grounding** suggests that Vera's perception of the bus is grounded in her experience *as of those low-level properties*. Since something's instantiating the relevant low-level properties does not guarantee that it instantiates the property of *being a bus*—after all, the bus-façade instantiates them, too!—both **Mediation** and **Mediation**\* suggest that Vera's perception of the bus must be mediated by background knowledge of some relation between those low-level properties and busses.

This sounds wrong. Again, intuitively Vera's perception of the bus is unmediated. I've already hinted at the explanation: It's not clear that Vera's experience *as of those low-level properties* has different phenomenal content from her experience *as of a bus*. It may just be that Vera can think about one content in two different ways—as bus-like or, more abstractly, as of various shapes, colors, etc.<sup>28</sup>

---

<sup>28</sup> I'm not asserting that this is necessarily the case. I think it may turn out that Vera's perception here *is* mediated by knowledge of a relation between certain low-level properties and the presence of busses. For it may be possible to have those low-level experiences but not see them as bus-like. Consider the relation between perception of a duck (or a rabbit) and experiences *as of the low-level features of the duck-rabbit*. Perhaps *all* high-level perceptions are like this, and so are all mediated. If that were the case, it would be even easier to reach my skeptical conclusion, so I make the optimist's position stronger by assuming Vera's perception might be as suggested here. The general point is that in cases where perceptions *are* unmediated—and surely, some are—one of two things must be happening: (1) There is some Y such that perceptions of X track experiences *as of Y* even in the absence of X, but experiences *as of X* and *as of Y* have the same phenomenal content (or, at least, indistinguishable—see note 24, above); or (2) There is no such Y, because illusory experiences of X are impossible. (This might be the case if certain mental states are *luminous*—if we can't be mistaken which we are in. For an argument against this, see Williamson (2002).)

This suggests a way of resisting **Grounding**. It's true that Vera's perceptions of busses track her experiences *as of certain low-level properties*, even in the absence of busses. But the explanation for this seems not to be that the former is grounded in the latter. Rather, the explanation is that her perception is grounded in an experience *as of a bus*, and since that experience has the same phenomenal content as an experience *as of certain low-level properties*, anything that tracks the former is likely to track the latter as well. And, of course, given that illusions are possible, it's no surprise that Vera's perceptions of busses track bus-like experiences more reliably than they track busses. This motivates revision of **Grounding** (with a similar revision to **Mediation**<sup>29</sup>):

**Grounding\*** If perceptions of X track experiences *as of Y* even in the absence of X, **and experiences as of X have different phenomenal content from experiences as of Y**, then the best explanation is that perceptions of X are grounded in experiences *as of Y*.

**Grounding\*** illuminates a way to resist the conclusion that all epistemically successful moral perception is mediated: Argue that some moral and non-moral experiences have the same phenomenal content. If certain experiences *as of suffering* have the same phenomenal content as certain experiences *as of wrongness*, then the fact that Sam's perceptions of wrongness track his experiences *as of suffering* even in the absence of wrongness is unremarkable. It causes no more epistemic trouble than the fact that Vera's perceptions of busses track both her experiences *as of certain low-level properties* and *as of busses* even in the absence of busses.<sup>30</sup>

I highly doubt that any moral and non-moral experiences have the same phenomenal content. In §2.3, I'll offer some reasons to share my doubt. However, I admit what I say there may not be decisive. Thus, the conclusion of §2 is conditional: If moral and non-moral experiences have different phenomenal contents, then the counterfactual test shows that epistemically successful moral perceptions are grounded in non-moral experiences. It would follow that moral

---

<sup>29</sup> We need to change **Mediation** if we are to avoid the conclusion that a perception of X grounded in an experience *as of X* must be mediated by knowledge of a relation between X and itself. It's also worth noting that if experiences with the same phenomenal content are identical, *only Mediation* needs to be amended, for if X-like and Y-like experiences are identical, it is trivial that a perception grounded in one is grounded in the other. Framing things as I do is meant to make it easier to raise certain objections to my arguments.

<sup>30</sup> This is not to say, of course, that there are no puzzles about how experiences *as of X* generate knowledge about X, given the possibility of illusions, only that these puzzles would pose no special problem for moral perception.

perception can only produce moral knowledge insofar as it is mediated by background knowledge of moral bridge principles.

### *2.3 Some Reasons to Think Moral and Non-Moral Experiences Have Different Phenomenal Contents*

I have suggested that the optimist about unmediated moral perception might argue that some moral and non-moral experiences have the same phenomenal content. If so, we should expect and be untroubled by the fact that epistemically successful moral perceptions track certain non-moral experiences. Crucially, though, we should *also* expect that some epistemically successful *non-moral* perceptions will track certain *moral* experiences. For example, suppose all and only pleasure is good, and that experiences *as of goodness* and experiences *as of pleasure* have the same phenomenal content. Suppose unmediated perceptions of goodness are grounded in, and therefore track, experiences *as of goodness*. Barring exceptional cases, we should thus expect that perceptions of goodness will also track experiences *as of pleasure*. But this goes both ways. So we should *also* expect that, barring exceptional cases, perceptions of pleasure will track experiences *as of goodness*. This suggests that people would rarely, if ever, perceive pleasure without having an experience *as of goodness*.<sup>31</sup>

I have no knock-down argument against this. But I will say a few things about why I doubt it. First, note that neither of the moral cases discussed so far provides any help to my opponent here. Start with Sam. Somewhat ironically, there is an argument *for* moral perception that relies explicitly on the idea that an experience *as of* cat-torture is phenomenally distinct from an experience of its wrongness. Preston Werner (forthcoming) asks us to consider two characters: One is like Sam. The other, Pathos, is an “EEDI”: Someone who has “a fully-functioning ‘theory of mind’ . . . but who nonetheless lack[s] affective empathy in the sense that they fail to have ‘an emotional response to another individual that is congruent with the other’s emotional reaction’.”<sup>32</sup> Werner points out that there would almost certainly be a phenomenal difference between Sam’s experience of cat-torture and the EEDI’s. Werner further “contend[s] that the best explanation is a difference in the perception of moral properties—[Sam]’s experience represents the cat’s burning *as [wrong]*, whereas Pathos’s does not.” If Werner is right, then it seems perfectly

---

<sup>31</sup> Barring this sort of monism, the situation will be more complicated. For example, suppose causing suffering is wrong, but many other things are, too. In that case, there will be, at best, a partial overlap between experiences *as of suffering* and experiences *as of wrongness*. This won’t change the basic issues, though, just make them more complicated.

<sup>32</sup> The inner quotation concerning what it is to lack “affective empathy” comes from Blair (2007, 4).

possible for someone to perceive cat-torture—as both Sam and Pathos do—without having an experience *as of wrongness*—as only Sam does.<sup>33</sup>

Of course, my opponent could argue that this is an exceptional case. As noted in §1, even if two experiences have the same phenomenal content, it might be possible to have one without the other, say if one lacks the ability to apply a certain concept to that content. Perhaps this is the case here: Perhaps Sam and Pathos actually have identical experiences—or, at least, any difference is a reaction to, rather than a part of, the experience grounding their non-moral perception—but only Sam is able to conceive of it as an experience *as of wrongness*. This simply doesn't have the ring of truth, as they say—especially if there are cases where Pathos *can* appropriately apply moral concepts (perhaps ones not requiring empathy). It seems much more natural to think that Sam's experience *as of wrongness* is something over and above the experience *as of cat-torture* he and Pathos share.

Alice's case is even less helpful to my opponent. For Alice's case is entirely intrapersonal. Clearly, it is possible for Alice to experience the non-moral features of Bob and Chuck's relationship without having an experience *as of permissibility*, for this is precisely what she does at the beginning of the story! It is only after constant interaction with them that she comes to see their relationship as permissible. It's hard to see how this could be an exceptional case.<sup>34</sup>

If I'm right, the unhelpfulness of these cases is not accidental. Rather, it is a symptom of a consistent phenomenal gap between moral and non-moral experiences. I am inclined to accept a sort of phenomenal parallel of some classic arguments about moral semantics:<sup>35</sup> For any way I might experience the world *non-*

---

<sup>33</sup> The claim that EEDIs have a theory of mind serves to block the alternative that only Sam perceives certain relevant non-moral features of the cat-torture, such as the cat's suffering.

<sup>34</sup> Another way to push back here stems from the idea of *undercutting penetration* (*i.e., mediation*). See note 16, above. The idea would be that experiences *as of suffering* and experiences *as of wrongness* have the same phenomenal content, but certain background beliefs can penetrate the experience to undercut perception of one while retaining perception of the other. This seems possible, but I doubt it can be what's going on in these cases. Think back to Vera. Suppose she learns that there are a lot of motorcyclists wearing bus-façades on her current route. When she gets cut off, she might fail to perceive a bus. But surely her experience would still be relevantly bus-like. If we asked her if she saw a bus, she'd presumably reply that she *seemed* to, but knows it very likely wasn't a bus. In contrast, it seems that Pathos and early Alice don't even have the relevant moral *seemings*.

<sup>35</sup> E.g., Moore's (1903) open question argument, Hare's (e.g., 1952) missionaries and cannibals, and Horgan and Timmons' (e.g., 1991) Moral Twin Earth.



*morally*, it is possible for me to fail to experience it *morally* in any particular way.<sup>36</sup> So moral and non-moral experiences must have different phenomenal contents.

Of course, not everyone shares these intuitions. Analytic naturalists, for instance, maintain that there are some conceptually necessary connections between moral and non-moral properties. It is not hard to believe that they would further (indeed perhaps *because* of this) maintain that some moral and non-moral experiences have the same phenomenal content. They may be able to defend unmediated moral perception. I have done what I can to cast doubt on this. I will thus move on to consider the implications for a purely perceptual moral epistemology, assuming I'm right that epistemically successful moral perception would have to be mediated.

### 3. Against Perceptual Knowledge of Moral Bridge Principles

In this section, I argue that if epistemically successful moral perception is mediated by knowledge of moral bridge principles, there can be no purely perceptual moral epistemology. Those seeking to offer a purely perceptual moral epistemology would have to offer a purely perceptual epistemology for those bridge principles. If moral perception is mediated, they cannot.

Suppose Sam knows that cat-torture is wrong. Here are three ways—arguably, the *only* three ways—he might know this:

- (1) Sam knows the principle a priori.
- (2) Sam knows the principle through perception.
- (3) Sam infers the principle from background knowledge.

I will start by setting aside views on which moral bridge principles are analytic. This is not because I think such views are false but because, as discussed in §2.3, they may be able to vindicate unmediated moral perception, and thus needn't

---

<sup>36</sup> I'm inclined to think this just *is* a Moorean intuition. Can one accept both Moorean intuitions *and* that moral and non-moral experiences can have the same content? Sort of. Suppose Carla gained her concept 'wrong' by ostension. At a young age, when confronted with certain things, she was told: That is wrong. She might thus claim that when confronted with suffering today, she has an experience *as of wrongness*. Nevertheless, she might still accept the Moorean intuition that, conceptually, being suffering isn't actually sufficient for being wrong. I'm not sure how psychologically plausible this is. But even if it is, it provides a rather hollow victory for the optimist about unmediated moral perception. For in order to develop a purely perceptual moral epistemology, we need to account for the knowledge of whoever pointed those wrong things out to her. Suppose we trace this back to the origin of the concept. Moorean intuitions seem to block the idea that this origin could itself be ostensive. This is importantly different from certain non-moral cases. For instance, not only might Vera have gained the concept 'bus' via ostension; it seems the origin of that concept might have been ostensive, too.

worry about my arguments in this section. On the other hand, the principle's being an instance of *synthetic* a priori knowledge is obviously incompatible with an attempt to develop a purely perceptual epistemology for it. So (1) provides no additional hope for those who would defend a purely perceptual moral epistemology.

Moving to (2): If perception of principles is even possible, it surely must be mediated. After all, principles are not the sorts of things that can directly explain our perceptions of them. What sort of background knowledge might mediate such perception? So far as I can see, the only candidate would be knowledge of the principle's *relata*. If this is to be purely perceptual, knowledge of the relata would have to be perceptual. In a moral bridge principle, though, one of the relata is a moral property. Since the possibility of purely perceptual knowledge of moral properties is precisely what is at issue, (2) serves only to move the bump in the rug.

This brings us to (3). In order for (3) to help the defender of a purely perceptual moral epistemology, the inference in question would have to be somehow grounded in perception. To see how this might go, return to Cliff<sub>2</sub>'s perception of H<sub>2</sub>O. Cliff<sub>2</sub> has an experience *as of water*. He also perceives H<sub>2</sub>O—a perception mediated by his knowledge that water is H<sub>2</sub>O. How does he know that water is H<sub>2</sub>O? A while back, he collected some samples of water and looked at them under a high-powered microscope. He noticed that, in all cases, the samples had the molecular structure of two hydrogen atoms and one oxygen atom. Given this, plus further (we may presume perceptual) knowledge about molecular behavior, Cliff<sub>2</sub> concluded that the best explanation for this consistent correlation is that water is H<sub>2</sub>O.<sup>37</sup> We might thus conclude that, despite being mediated, Cliff<sub>2</sub>'s knowledge is purely perceptual.<sup>38</sup>

Suppose Sam claims that his knowledge that cat-torture is wrong works much the same way: At some point in the past, he had multiple correlated perceptions of wrongness and of cat-torture (yikes!), and inferred that the best explanation for this is the truth of the principle that cat-torture is wrong. I accept, for the sake of argument, that this is a live possibility. Importantly, it still won't permit Sam to conclude that his knowledge here is purely perceptual.

To see why, recall that we are assuming that Sam's perceptions of wrongness are mediated by knowledge of moral bridge principles. Clearly, if he

---

<sup>37</sup> Drawing this conclusion about property identity might require further, non-empirical philosophical knowledge about the nature of property identity, but Cliff could make do with a weaker relation, such as consistent correlation.

<sup>38</sup> One might wish to resist this conclusion. Perhaps Cliff<sub>2</sub>'s abductive inference relies on something a priori, for instance. Since this would only serve to weaken the case for purely perceptual knowledge of moral bridge principles, I set it aside.

claims to know that cat-torture is wrong via abduction from correlated perceptions of wrongness and cat-torture, the relevant background knowledge can't be that cat-torture is wrong. His perceptions of wrongness must have been mediated by other knowledge—say, the bridge principle *causing suffering is wrong* as well as his knowledge that cat-torture causing suffering.

We are now left to ask how Sam knows that causing suffering is wrong. Suppose that, again, he claims to have arrived at this principle via abduction from multiple correlated perceptions of wrongness and of suffering (it's been a bad year). If so, there must be some *further* bridge principle that mediated his perceptions of wrongness in *those* cases. And we must ask how Sam knows *that* principle. If we are to avoid an unhelpful regress, at *some* point Sam must gain the relevant knowledge in some other way.<sup>39</sup>

Assuming these are indeed the only options, we may conclude that while there may be perceptual knowledge of *some* moral bridge principles, others—specifically, the *fundamental* ones—must be known non-perceptually. It follows that there can be no purely perceptual epistemology for moral bridge principles. If all moral perception is indeed mediated by knowledge of such principles, it further follows that there can be no purely perceptual moral epistemology.

#### 4. Conclusion

The potential epistemological advantages of defending moral perception are highly seductive. Given the myriad explanatory challenges facing any metaethical view, it would be an incredible boon to be able to say that we know about the instantiation of moral properties in much the same way as we know that there is a bus passing in front of us. Unfortunately, even if we sometimes perceive moral properties, this alone isn't sufficient to grant that benefit. For moral perceptions may be mediated by background knowledge which is not itself perceptual. I have argued that unless some moral and non-moral experiences have the same phenomenal content, moral perception would have to be mediated. I further argued that if moral perception is mediated, knowledge of the mediating moral bridge principles could not itself be purely perceptual.

Before concluding, I'll note two potential further implications of my arguments in this paper. The first is rather obvious: My arguments might be taken to support moral *rationalism*. Looking back to §3, it seems that of the three options presented, all but (1) depend on background moral knowledge. This might indicate that knowledge of the *fundamental* moral principles must be a priori.

---

<sup>39</sup> Compare Zangwill (2006).

Second, my argument that epistemically successful moral perception must be mediated by knowledge of moral bridge *principles* might be taken to have problematic implications for moral particularism. Whether this is the case will depend, for one, on whether the kinds of principles involved in mediation of moral perception are the same as those particularists reject.

Both issues are too complex to be addressed here. I mention them by way of sowing seeds to be cultivated elsewhere, as it were. Here, my conclusion is that, insofar as moral and non-moral experiences have different phenomenal content, moral knowledge cannot be purely perceptual.

### Acknowledgements

My thanks to Christian Coons for turning me on to this use of the counterfactual test, to Kelly Trogdon for introducing me to the perception literature that helped me frame it, to Tristram McPherson for raising numerous objections to it, and to all three for subsequent invaluable discussions. Thanks also to Molly Gardner, Graham Oddie, Geoff Sayre-McCord, Walter Sinnott-Armstrong, Jacob Sparks, Christina Van Dyke, Mark van Roojen, Pekka Väyrynen, Preston Werner and participants at the 2013 Rocky Mountain Ethics Congress and the 2014 PPE Workshop at UNC Chapel Hill.

### References

- Audi, R. 2013. *Moral Perception*. Princeton University Press.
- Blair, R.J.R. 2007. "Empathic Dysfunction in Psychopathic Individuals." In *Empathy in Mental Illness*, 3–16. Cambridge University Press.
- Blum, L.A. 1994. *Moral Perception and Particularity*. Cambridge University Press.
- Bonjour, L. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5 (1): 53–74.
- Chappell, T. 2008. "Moral Perception." *Philosophy* 83 (4): 421–437.
- Cowan, R. forthcoming. "Perceptual Intuitionism." *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12023>.
- Cullison, A. 2010. "Moral Perception." *European Journal of Philosophy* 18 (2): 159–75. <https://doi.org/10.1111/j.1468-0378.2009.00343.x>.
- Cuneo, T. 2003. "Reidian Moral Perception." *Canadian Journal of Philosophy* 33 (2): 229–258.
- DePaul, M.R. 1993. *Balance and Refinement: Beyond Coherence Methods of Moral Inquiry*. Taylor & Francis.
- Gordon, R.M. 2009. "Folk Psychology as Mental Simulation." In *The Stanford Encyclopedia of Philosophy*, edited by E.N. Zalta, Fall 2009. <http://plato.stanford.edu/archives/fall2009/entries/folkpsych-simulation/>.

- Greco, J. 2000. *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge University Press.
- Hare, R.M. 1952. *The Language of Morals*. Oxford Clarendon.
- Harman, G. 1977. *The Nature of Morality: An Introduction to Ethics*. Oxford University Press, USA.
- . 1986. “Moral Explanations of Natural Facts—Can Moral Claims Be Tested Against Moral Reality?” *The Southern Journal of Philosophy* 24 (S1): 57–68.
- Horgan, T., and M. Timmons. 1991. “New Wave Moral Realism Meets Moral Twin Earth.” *Journal of Philosophical Research* 16: 447–465.
- McBrayer, J.P. 2010a. “A Limited Defense of Moral Perception.” *Philosophical Studies* 149 (3): 305–320.
- . 2010b. “Moral Perception and the Causal Objection.” *Ratio* 23 (3): 291–307.
- McGrath, S. 2004. “Moral Knowledge by Perception.” *Philosophical Perspectives* 18 (1): 209–228.
- McNaughton, D. 1988. *Moral Vision: An Introduction to Ethics*. Oxford, UK; New York, NY: Wiley-Blackwell.
- Moore, G.E. 1903. *Principia Ethica*. At the University Press.
- Oddie, G. 2009. *Value, Reality, and Desire*. Oxford University Press.
- Putnam, H. 1975. “The Meaning of ‘Meaning.’” *Minnesota Studies in the Philosophy of Science* 7: 131–93.
- Sayre-McCord, G. 1996. “Coherentist Epistemology and Moral Theory.” In *Moral Knowledge? New Readings in Moral Epistemology*, edited by W. Sinnott-Armstrong and M. Timmons. Oxford University Press.
- Shafer-Landau, R. 2003. *Moral Realism: A Defence*. Oxford University Press.
- Siegel, S. 2011. *The Contents of Visual Experience*. Oxford University Press.
- Slutsky, D. 2001. “Causally Inefficacious Moral Properties.” *Southern Journal of Philosophy* 39 (4): 595–610.
- Sturgeon, N. 1986. “Harman on Moral Explanations of Natural Facts.” *Southern Journal of Philosophy* 24 (S1): 69–78.
- Väyrynen, P. m.s. “Doubts about Moral Perception.”  
<http://www.personal.leeds.ac.uk/~phlpv/papers/moralperception.pdf>.
- . 2013. *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics*. Oxford University Press.
- Watkins, M., and K. D. Jolley. 2002. “Pollyanna Realism: Moral Perception and Moral Properties.” *Australasian Journal of Philosophy* 80 (1): 75–85.
- Werner, P. forthcoming. “Moral Perception and the Contents of Experience.” *Journal of Moral Philosophy*.
- Williamson, T. 2002. *Knowledge and Its Limits*. Oxford University Press.
- Zangwill, N. 2006. “Moral Epistemology and the Because Constraint.” In *Contemporary Debates in Moral Theory*, edited by J. Dreier, 263–281. Wiley-Blackwell.