IS MORAL PERCEPTION ESSENTIALLY RULE-GOVERNED? A CRITICAL ASSESSMENT OF GENERALISM AND A LIMITED DEFENSE OF PARTICULARISM

Peter Shiu-Hwa Tsu
National Chung Cheng University
Department of Philosophy
u4079238@gmail.com

SUMMARY: Moral perception, for the purposes of this article, is taken to be the perception of moral properties, unless contexts dictate otherwise. While both particularists and generalists agree that we can perceive the moral properties of an action or a feature, they disagree, however, over whether rules play any essential role in moral perception. The particularists argue for a 'no' answer, whereas the generalists say 'yes'. In this paper, I provide a limited defense of particularism by rebutting several powerful generalist arguments. It is hoped particularism can thus be made more attractive as a theory of moral perception. Positive arguments for particularism will also be provided along the way.

KEY WORDS: rules of moral salience, connectionism, frame problem, practical wisdom, moral phenomenology

RESUMEN: La percepción moral, para los propósitos de este artículo, se considera la percepción de propiedades morales, a menos que los contextos dicten lo contrario. Si bien tanto los particularistas como los generalistas están de acuerdo en que podemos percibir las propiedades morales de una acción o un rasgo, no están de acuerdo, sin embargo, sobre si las reglas desempeñan un papel esencial en la percepción moral. Los particularistas abogan por una respuesta "no", mientras que los generalistas dicen "sí". En este artículo, ofrezco una defensa limitada del particularismo refutando varios argumentos generalistas poderosos. Se espera que el particularismo pueda hacerse más atractivo como teoría de la percepción moral. A lo largo del camino, también se proporcionarán argumentos positivos para el particularismo.

PALABRAS CLAVE: reglas de relevancia moral, conexionismo, problema de encuadre, sabiduría práctica, fenomenología moral

1. *Introduction: The Problem of Moral Perception*

Suppose that you witnessed a cold-blooded killing. You saw the killer stab a man with a dagger. And you heard the victim cry out and scream, begging the killer to stop. Yet, the killer didn't relent and just stabbed the man again and again, finishing off with a fatal stab at the man's heart. When the killer pulled out the dagger,

you saw it dripping with blood. You feel horrified at the sight of this, hiding in the corner of the room, hoping that the killer would not see you.

Such a scenario is indeed horrifying. And ceteris paribus, you would rightly judge that what you witnessed is morally horrendous and wrong. But how do you come to know its moral wrongness? After all, the property of moral wrongness, unlike the property of killing you witnessed, is not something you can perceive through your five senses. The property of moral wrongness does not have a size or color, nor does it smell, make a noise, or have a weight that you can feel for that matter. To prevent ourselves from falling into the trap of moral perception skepticism, according to which moral properties, even if they exist, are beyond our perceptual reach, it seems imperative for us to provide an account of how we come to be perceptually acquainted with moral properties.[1] The task of providing such an account constitutes what I call 'the problem of moral perception'.[2]

In fact, the problem of moral perception has been philosophers' one million dollar question since Hume, not the least because moral perception is often seen as a potential source of moral knowledge.[3] In addressing this problem, Hume himself takes a sentimentalist view, maintaining that the moral properties are nothing more than the sentiments or emotions we have in our mind, and we come to be acquainted with them via introspection. However, nowadays, not too many find this account convincing. For one thing, the killing still seems to be morally wrong, even if the witness, or indeed anyone, doesn't have any negative sentiments towards it.

Since Hume, various alternative accounts have been proposed to meet the challenge to provide a plausible account of moral perception.[4] Some argue that we can literally perceive moral properties,

---

[1] See Väyrynen (2018b) for a defense of a version of moral perception skepticism. But since moral perception skepticism, if true, is a common problem that both generalists and particularists have to face, I will set it aside for the purpose of this paper.

[2] The problem of moral perception, as I understand it, is to be distinguished from a closely related problem, or what Mower and McGraw (2015, p. 274) call the 'hard problem of normativity', which is a problem of sorting out the ethically relevant features one *should* be morally sensitive to. The solution to this problem of normativity involves, according to Mower and McGraw (2015, p. 276), "*doing* ethics", which I concur with.

[3] McGrath 2004; Cullison 2009; Bagnoli 2011, p. 96.

[4] See for instance Watkins and Jolley 2002; Jacobson 2005; McGregor 2015; McBrayer 2010, and Cullison 2009.

by extending the scope of our sensory modalities beyond the five senses, whereas others maintain that moral properties are reducible to natural properties, and therefore we can perceive them without postulating extra-sensory moral faculty.[5] Still others maintain that we have a rational capacity that enables us to capture moral properties in pretty much the same way as it enables us to capture mathematical properties.[6] Finally, Audi (2013) recently came up with an intriguing proposal: moral properties, though non-reducible to natural properties, can be *indirectly* perceived through those that ground them.[7] For the purposes of this paper, I shall focus on a prominent account proposed by the moral generalists, according to which the structure of moral perception is essentially governed by (exceptionless) moral rules that connect natural properties with moral properties (such as 'killing is wrong/wrong-making').[8,9] It is through the application of these 'natural-moral rules', as we may call them, that we come to perceive the moral properties of things.[10]

[5] For an example of the former view, see Wisnewski 2015, whereas the latter view is generally endorsed by the analytic utilitarians.

[6] For the purpose of this article, I assume it harmless, modulo Kagan 2001 (p. 48), to refer to this rational capacity as a kind of a faculty of moral perception, through which we can (intellectually?) perceive the moral properties of an action (cf. McGrath 2018). This sort of view is generally dubbed as 'rational intuitionism'. See Darwall 1998, chapter 5. Chudnoff (2016) espouses what I regard as a version (or a cousin?) of this view. According to Chudnoff, in at least some cases of moral perception, we intellectually grasp a general moral truth. Wodak (2019) also offers an illuminating defense of a version of intuitionism, according to which the moral properties of hypothetical cases can be known through intuition.

[7] Cf. Dancy 2010, p. 105.

[8] It's to be noted here that moral generalism *per se* need not be committed to the claim here, as one reviewer correctly observed. After all, 'moral generalism' is a term of art that people use to refer to various different doctrines (see Tsu 2018). For the purposes of this article, I focus on the specific version of moral generalism as stated in the text, because, inter alia, it is a very prominent one advocated by several eminent philosophers, as will be detailed in what follows. Moreover, it has a direct bearing on how 'the problem of moral perception' is to be solved, which is the central question investigated in this article. Due to the limited space, the evaluation of other versions of moral generalism will have to wait for another occasion.

[9] I put 'exceptionless' in parentheses for the following reason: while the exceptionless version of generalism is no doubt endorsed by Jackson, Smith, Pettit, and McKeever and Ridge, it is less clear, as I will discuss later in sections 2 and 3.2, whether the exceptionless version is endorsed by Herman or Clark. Yet, all of them can be properly called 'generalists' in that they all think moral rules play a central and indispensable role in moral perception.

[10] My characterization of moral generalism largely coincides with Faraci's (2015) characterization of what he calls 'impurism', according to which moral perception is dependent for its epistemic merit (e.g. providing justifications for our moral

This generalist model has gradually taken shape in the discussion of moral perception, being endorsed by a number of eminent philosophers, including, for instance, Andy Clark (1996), Frank Jackson, Michael Smith, Philip Pettit (2000), Barbara Herman (1985), Sean McKeever and Michael Ridge (2006). These generalists argue from both empirical and conceptual fronts that moral perception must be essentially rule-governed. For instance, Clark argues from the perspective of cognitive sciences that in order for collaborative moral problem solving to be actually possible, rules are indispensable in our moral perception. Advocating a classical model of moral perception, Jackson, Smith and Pettit maintain that our moral perception is not naturally shapeless and can thus be codified by rules. Herman, on the other hand, argues on the conceptual front that the Kantian model of moral judgment, to be plausible, needs to be supplemented with an RMS generalist model of moral perception, according to which moral perception is essentially structured by what Herman calls 'rules of moral salience' (RMS). Finally, McKeever and Ridge contend that for practical wisdom not to be a chimera, moral perception must be rule-governed.

These generalists' arguments are quite powerful, but wrong, or so I will argue. Although there have been quite a few anti-generalist (or particularist) theorizations about moral perception,[11] these powerful generalist arguments have not received any *systematic* examination from the particularist camp, as far as I know.[12] Nor do they receive their deserved attention in the current literature on moral perception on the whole, much to one's dismay.[13] This paper aims to take up

beliefs) on background knowledge of bridge principles that link moral and non-moral (natural) properties. While Faraci further argues that such background knowledge is not purely perceptual (and hence our moral perception is 'impure'), moral generalists are neutral with respect to the debate between 'impurism' and 'purism'. For a defense of 'purism', see Werner 2018. For a defense of (a version of) 'impurism', apart from Faraci 2015, see also Cowan 2014, which argues that our moral perception is 'cognitively penetrated'.

[11] I use 'anti-generalism' interchangeably with 'particularism' in this paper.

[12] 'Particularism' is a term of art which different philosophers use to mean different things. For useful recent surveys, see Tsu 2018 and Väyrynen 2018a. For the purpose of this paper, I take it to be the view that our moral perception is not essentially rule-governed. The particularists who support particularism in this sense include, inter alia, Murdoch 1991; Dancy 1993, 2004; McDowell 2002a; Blum 1991; Nussbaum 1992; Millgram 2005; Garfield 2000; Churchland 2000; McGrath 2018, and on a certain interpretation, Kukla 2002.

[13] The most updated survey article on moral perception (Werner 2020) does not mention any of the generalists' works. One reason for this (undeserved) neglect in the moral perception literature, I suspect, is that Jackson, Smith, and Pettit (2000)

the slack by addressing them and provide a limited defense of par-
ticularism, or the anti-generalist view that moral perception is not
essentially structured by moral rules. According to the particular-
ists, we perceive moral properties of things via a kind of non-rule-
governed moral sensitivity that is essentially embedded in our 'form
of life', to borrow a term from Wittgenstein. Positive arguments for
particularism will also be provided along the way. Hopefully, we will
end up seeing particularist theory of moral perception in a more
favorable light.

Here's the plan of how we will proceed. First of all, I will argue
that the sort of moral rules advocated by Herman not only are
insufficient to account for moral perception, but are not necessary for
it, either. Drawing on the works of John McDowell, I will provide a
particularist explanation of how moral perception is possible without
being governed by the Hermanian rules of moral salience.

Second, as some particularists invoke connectionism in cognitive
sciences to support particularism, some generalists counteract by ar-
guing against connectionism. For instance, Jackson, Smith, and Pet-
tit, qua generalists, bolster classical rule-based model of moral per-
ception against connectionism, maintaining that particularism rests
on a shaky foundation. I will argue, however, that Jackson, Smith
and Pettit's classical model faces insurmountable difficulties account-
ing for typicality. On the other hand, others contend that although
connectionism itself is plausible, the particularists misunderstand its
implications. Connectionism, Clark argues for instance, does not war-
rant the marginalization of moral rules. I agree, but the sorts of moral
rules Clark has in mind (i.e. the summary moral rules, as we shall
see later) are too weak to harm particularism, or so I will argue.

Finally, I will argue against McKeever and Ridge's claim that the
moral perception of the practically wise is essentially governed by a
finite and manageable set of rules. Instead, I will shoulder the task
of explaining how a particularist practically wise person can come to

is overtly concerned with moral concepts, arguing from the fact of our competence
with moral concepts to the view that the structure of moral concepts is governed by
rules (or what they call 'patterns' or 'shapes'). On top of this, Clark 1996 is overtly
concerned with moral cognition. Yet, as we will see in section 3, their arguments,
mutatis mutandis, can be applied equally to moral perception as well. Moreover,
to the best of my knowledge, the common generalist thread (i.e. the claim that
moral perception is essentially governed by natural-moral rules) that runs through
the works of Jackson, Smith and Pettit (2000), Clark (1996), Herman (1985), and
McKeever and Ridge (2006), hasn't been well-identified in the moral perception
literature. Yet, this is only an explanation of the neglect, not a justification.

have moral perception, given the fact that the number of potentially morally relevant features is large or even infinite.

2. *Herman's RMS Model*

According to Herman, an eminent Kantian scholar, a Kantian model of our practice of moral judgment needs to be supplemented with a model of moral perception in order for it to be viable. This model of moral perception has to supply the moral agent with morally salient features of the situations he faces, so that he can go on to formulate his maxim (or the subjective principle for his action) in terms of them, put the maxim to the test of universalizability, and finally make a moral judgment on the basis of the test result.

What is distinctive about this model of moral perception, according to Herman, is that the moral agent perceives the morally salient features through the lens of what Herman calls 'rules of moral salience' (or RMS for short). To be more specific, the content of the RMS is as follows: 'x is morally salient' (x is a variable for natural features). When the moral agent perceives through the lens of RMS, he comes to perceive the moral salience of lying, stealing, helping others, fulfilling one's promise, and the likes. That is, the RMS highlight the moral salience of some natural features. Without the RMS in operation, there is no way for us to perceive the moral salience of the natural features, according to Herman. For Herman, the RMS constitute the structure of our moral perception or sensitivities. If so, a generalist model of moral perception, according to which moral rules play an indispensable role in our moral perception, can be vindicated.

Against this, it can be immediately questioned whether the morally salient features are exhausted by RMS. If not, then this would show that there are at least some morally salient features that are not specified by RMS. That is, we might well perceive the moral salience of some natural features without the help of RMS. And it follows from this that it would be an exaggeration to contend that RMS constitute the essential structure of moral perception.

And I do think there is reason to believe that the morally salient features, being potentially infinite in number, cannot be exhausted by RMS, which are presumably finite. The RMS have to be finite in number, for no human being, being cognitively limited, can be mentally equipped with an infinite number of RMS. On the other hand, the morally salient features are potentially infinite. How so? This is due to the context-dependent nature of moral salience. Any natural

feature, when situated in the right context, can be morally salient. For instance, wearing yellow shoelaces, though normally morally un-salient, can become morally salient if you promise your girlfriend to do so when you go out with her on a date. Or an action performed on Tuesday rather than any other day of the week can also be morally salient, if Tuesday is the day of a funeral for a beloved friend who sacrificed his life in the war to save yours. But in Herman's list of RMS, there does not appear to be any rule that specifies wearing yel-low shoelaces or a Tuesday action as morally salient. For Herman, the content of the RMS is determined by whether a feature is relevant to Kant's second formulation of the categorical imperative, i.e. treating humanity always as an end. It is not unimaginable that on Herman's list of RMS, features such as stealing, lying, intentional ending of a life, helping others, and the likes would be included, whereas features such as wearing yellow shoelaces or doing something on a Tuesday would not.

Two potential responses can be made by Herman. First, it may be argued that *wearing a yellow shoelace*, for instance, *never is a morally salient feature*. Its seeming moral salience is derived from the promise one makes. It is really the promise that has moral salience. Personally, I do not find this a convincing response, for phenomenologically, the feature of wearing yellow shoelaces does strike me as morally salient in this scenario and is thus distinguished from other features that are not, such as the colors of my necktie or suit. But I don't want to stake too much on my personal moral experience. The major defect of this response lies in the fact that it would amount to an implausible view, according to which an agent may perceive his promise as morally salient in the scenario while not perceiving wearing yellow shoelaces as morally salient. This would amount to treating the agent's commitment to RMS as some sort of 'abstract commitment', in the words of Talbot Brewer (2000, p. 68) —the agent can be seriously committed to the RMS about promise-keeping while not seeing the content of his promise as morally salient. But abstract commitment seems far-fetched. One naturally doubts how serious the agent really is about the RMS about promise-keeping if he does not *ever* see the content of his promise as morally salient.

In fact, Herman herself seems to repudiate abstract commitment. So the above response, though available to her, is unlikely to win her approval. So now, let's turn to her second potential response, according to which, Herman takes the opposite tack by admitting the moral salience of the color of the shoelaces in the depicted scenario, but argues that it is entailed by the agent's moral perception through

the lens of the RMS about promise-keeping. That is, Herman might well argue that a true commitment to RMS about promise-keeping would require the moral agent to develop a sensitivity to the content of the promise, such that if the color of one's shoelace figured into one's promise, one would perceive it as morally salient.

But the trouble with this reply is that one might not necessarily perceive the moral salience of the color of one's shoelace even if one possesses the RMS about promise-keeping and is serious in one's commitment to it. This is not unimaginable, as one might fail to notice the moral salience of the color of one's shoelace in this scenario after a long day's work. To quip, even Homer nods. And it would be inaccurate to describe this person as lacking a commitment to the RMS about promise-keeping, just because his sensitivity to the content of his promise is dulled by his physical exhaustion or illness. Suppose that he were reminded about the color of his shoelaces by his girlfriend, and he immediately put on yellow ones, this would suggest that he is actually seriously committed to the RMS about promise-keeping.

So it seems that if one sees the moral salience of the color of the shoelace in this scenario, one does not necessarily see it through the lens of the RMS about promise-keeping; for, even granting a thicker understanding of the RMS (which involves developing a sensitivity to the content of the promise), the possession of the RMS doesn't guarantee that we will see the moral salience. So we have shown that what the moral agent perceives as morally salient may not be exhausted by the RMS.

Now it might well be retorted by Herman that a commitment to RMS is not meant to be sufficient for moral perception, but it is only meant to be necessary.[14] But is it necessary? Can't one perceive the moral salience of a natural feature without help of RMS? Here, I will argue that one can indeed do so without the help of RMS, drawing chiefly on John McDowell's works.

Unlike Herman-Kantian moral sensitivity or perception that is structured by rules of moral salience that are ultimately grounded in

---

[14] This is not to render the preceding discussions of the sufficiency point otiose. For one, Herman did not clearly indicate that RMS is only meant to be necessary. In fact, she appears to endorse the sufficiency point by contending that the possession of RMS (e.g. lying is morally salient) enables one to perceive the moral salience of the natural features (e.g. lying) should they appear. For another, the sufficiency point is of interest per se, insofar as one is concerned to figure out whether the possession of RMS would be enough for the perception of all the morally salient features there can be.

the second formulation of categorical imperative, McDowell (2002a)
contends that one's moral sensitivity is embedded in "forms of life".
Being a competent user of moral concept (in this case, the concept
of 'moral salience') in a specific form of life, one will naturally see
some natural features as morally salient, and others as not, just as
we will see "when an utterance is an assertion, when an appeal,
when an explanation" (p. 202) if we are competent with the uses of
concepts such as 'assertion', 'appeal' and 'explanation' in that form
of life, knowing when they are applicable and when not. And this
way of 'seeing' is not mediated through rules (but rather through a
'form of life'). For what we see as morally salient may not have any
commonality or 'shape' at the level of natural properties, such that
they can be codified into rules of the form 'x is morally salient' ('x'
stands for natural features that unify the morally salient). That is,
according to McDowell, what is morally salient to us might well be
naturally 'shapeless'. McDowell writes:

> [H]owever long a list we give of items to which a supervening term
> applies, described in terms of the level supervened upon, there may be
> no way, expressible at the level supervened upon, of grouping just such
> items together. (p. 202)

To paraphrase for our purposes:

> However long a list we give of features to which a moral concept (such
> as 'moral salience') applies, described in natural terms, there may be
> no way, expressible at the natural level, of grouping just such features
> together.[15]

In other words, the morally salient features may not have any natural
commonality or shape, such that a natural-moral rule of the form 'x
is morally salient' can be formed ('x' refers to a unifying natural
feature). This is not to beg the question against Herman, because
there is some evidence to believe that what is morally salient is indeed
naturally shapeless (or some evidence against the RMS model).

For instance, causing pain is not necessarily morally salient when
I roll up my sleeve to let the nurse to draw blood for a test of
cholesterol. Lying in a game called 'Bluffing' is not so either, given
that doing so is permitted by the rules of the game. Nor is helping,
when a three-year-old helps his kindergarten classmate build a Logo

---

[15] The above two quotes from McDowell are also used in Tsu 2013.

castle. Moreover, neither might the intentional ending of a life be particularly morally salient, when doing so doesn't cause its voluntary victim any pain or suffering, and our technology is advanced enough such that the voluntary victim can be revived in a split second in a rejuvenated form. And with all due respect, groping a woman's buttock, though morally salient on most occasions, is not particularly so when it is done by her two-year-old granddaughter. That is to say, what is morally salient may not have any natural commonality such that it can be codified into a rule of moral salience of the form that 'x is morally salient' ('x' stands for natural features that unify the morally salient). As we have shown, features such as causing pain, lying, intentional ending of a life, helping others, or groping a woman's buttock do not seem to be particularly morally salient in some scenarios. This suggests that the rules of moral salience can have exceptions.

It is important to note here that Herman seems to acknowledge this when she claims that

> [i]f the function of the RMS is to guide the normal agent to the perception [. . .] of the morally relevant features [. . .], *the rules cannot be very complex*. That is, *they will not include the range of exceptions [. . .] that one would require of rules of [overall] judgment*. (p. 420) [emphasis added]

In the above claim, Herman explicitly acknowledges that the RMS can have exceptions, because to serve their guiding function, they cannot be made very complex (and thus be made exceptionless) by including the exceptions into their contents.

Two comments are in order here. First, if the RMS have exceptions, it seems to me that it is an exaggeration to maintain, as Herman (p. 419) does, that "[t]he rules of moral salience constitute the *structure* of moral sensitivity [or perception]" [emphasis added]. For if the RMS have exceptions, this would suggest, contra Herman, that moral perception is *not* essentially structured by them. What the RMS indicate as morally salient may in fact be morally irrelevant (That's why the RMS have exceptions!). That is, the RMS might lead to 'moral hallucination' (rather than 'moral perception'). Very often, the RMS merely summarize what we perceive as morally salient. They are valid up to a point. But they do not *determine* in a top-down fashion (nor do they exhaust for that matter) what we perceive as morally salient. Our moral perception (or perception of what is morally salient) is thus not rule-governed.

Second, although Herman shuns the strategy of making the rules more complex (and thereby exceptionless) by building the exceptions into the contents of the rules, this is a strategy, *mutatis mutandis*, embraced by McKeever and Ridge 2006, as I will show in section 4. And I will argue there that this strategy fails, chiefly for the reason that it risks making the number and the complexity of the rules become cognitively unmanageable for cognitively limited beings like us.

3. *The Canberrans' Classical Model and Clark's Moral Collaboration Model*

In the last section, I argued that our perception of moral salience is essentially shapeless at the natural level and that Herman's RMS model of moral perception, being essentially rule-bound, cannot quite accommodate this without making significant sacrifices on her claim that our moral perception is essentially structured by RMS. Other things being equal, the fact that our moral perception (or perception of moral salience) is essentially naturally shapeless gives particularism an edge over generalism, as particularism, which maintains that our moral perception is not essentially rule-governed, is not incompatible with the shapelessness thesis. Now, if our moral perception is essentially naturally shapeless, then when we perceive the moral salience of a natural property, we do not (have to) perceive it via any rule of moral salience as advocated by Herman (but might well perceive it instead through a vision embedded in a form of life, as suggested by the McDowellian account I favor). But this is mostly a merely negative thesis, and it may rightly be wondered how on earth we come to perceive the moral salience of a natural property without the aid of any rule of moral salience. What is the sort of cognitive or perceptual mechanism underlying our vision embedded in a form of life that enables us to perceive moral salience without the aid of rules of moral salience? This would require some explanation.

To support particularism, Churchland (2000) rests it on the foundation of connectionism, which, in the eye of many cognitive scientists, can offer a compelling story about how our moral perception works.[16] Churchland argues that the natural shapelessness of our moral perception is essentially due to how moral perception is implemented in our brain, which consists of a connectionist network of neurons. The activations of neurons falling under a moral concept

---

[16] The distinction between perception and cognition, though perhaps real and non-arbitrary, is not especially germane here. What Churchland says about moral cognition can naturally translate into arguments about moral perception.

might share some similarities with one another; however, there is no commonality or shape amongst them. The fact that there is no shape or commonality amongst them doesn't mean, however, that they are completely messy. Instead, there are what the cognitive scientists call the 'prototypes', or complexes of central features the typical members of a concept more or less share, at the higher level of our cognitive structure. It is just that not all of the central features are possessed by all members of a concept such that hard and fast rules can thus be formulated. Thus, according to connectionism, we do not perceive the moral salience of natural features via any hard and fast rules; rather we perceive the moral salience of a natural feature, when the natural feature activates the neural network and corresponds to (at least one of) the central features which the members of the morally salient typically have.

The above model can be attacked from two fronts, however. One is an external critique from classical rule-based model of moral perception endorsed by Jackson, Pettit, and Smith (or "the Canberrans", as they once all worked in Canberra in Australia). The other is an internal connectionist critique from Clark. I will argue that neither of them succeeds.

3.1. The Canberrans' Classical Model

Jackson, Pettit, and Smith (2000) argue from the fact that we are competent with the use of moral concepts to the claim that the morally right or wrong, or the morally salient for that matter, must have a natural finite shape or commonality for us to latch onto. This is essentially because without a finite shape or commonality amongst the infinite number of features the morally salient can refer to, we, as cognitively limited beings, can never become competent in the use of moral concepts.

In this case, the Canberrans might well allow that at the implementation level, the activations of neurons might well be shapeless. But at the higher level, they are unified by shape, without which our moral perception becomes impossible. So the model championed by the Canberrans is distinguished from the connectionist one advocated by Churchland, not at the implementation level, but at the higher level.

However, can our competence with moral concepts only be explained via a grasp of a natural finite shape? In fact, we have already seen some hints in the connectionist model that this need not be so. It might well be explained in terms of our grasp of the prototypes,

or the complexes of central features which the members of a concept have. If so, there is no compelling pressure for us to accept the Canberrans' line of reasoning.

Anticipating this sort of objection, the Canberrans argue that the prototypes can actually be reduced to a special sort of shape or commonality —a *disjunctive* commonality. For instance, the connectionists at least would have to admit that x is morally salient, if x shares enough of the prototypes, or the central features, whatever they are, of the morally salient. And here, one can parlay the central features with a raft of disjuncts in the form of (c1 v c2 v c3 v c4 v c5), for instance. That is, if x is c1 v c2 v c3 v c4 v c5, then it is morally salient. This we might well call a Rossian disjunctive commonality, for it is very much in the spirit of Rossian prima facie duties.

So what is wrong with the Canberrans' idea of Rossian disjunctive commonality? I think there are at least two major problems with it. First, the prototypes cannot actually be reduced to disjunctive commonality. The prototypes are a conjunctive rather than disjunctive set of features that is only possessed by the most typical members of a concept. They are not shared by all members of a concept. Second, there is also the gradation of typicality which the idea of disjunctive commonality wouldn't be able to explain. On the Canberran's classical model, there is no explanation for why, for instance, A and B, though sharing the same disjunctive commonality, can nevertheless differ in their degrees of typicality. Appealing to disjunctive commonality won't help, because, by assumption, they have the same disjunctive commonality. By contrast, on the connectionist prototype model, one can explain by how similar A and B respectively is to the most typical member that has all the central prototypical features.

3.2. Clark's Moral Collaboration Model

Clark (1996) argues against particularism about moral perception from two fronts.[17] First, connectionism at the implementation level does not necessarily lead up to particularism about moral perception. The existence of rules is entirely compatible with connectionism. Connectionism only warrants a re-conception of the moral rules rather than a marginalization of them. Second, for our collaborative moral problem solving to be possible, rules are indispensable for they

---

[17] Like Churchland (2000), what Clark (1996) says is about moral cognition. All of what he says, however, can be translated naturally into arguments about moral perception. Again, although the distinction between cognition and perception might be real and non-arbitrary, it is not especially germane here.

provide "the essential starting point of informed moral debate". They constitute the "external scaffolding" that is integral to our moral perception, according to Clark.[18]

More slowly, connectionism at the implementation level is believed by Clark to be compatible with the existence of rules at the higher level of moral perception. To assess whether this claim can hold water, we need to further clarify what is meant by 'rules'. After all, we have shown in the last section that at the higher level of moral perception, it is shapeless with respect to the natural and therefore does not seem to be governed by hard and fast rules.

In fact, Clark (1996, p. 115) agrees with the natural shapelessness of moral perception. For him, qua connectionist, connectionism does have this implication. Yet, for him, there are still some rules integral to our moral perception that are not only compatible with connectionism but are central to our collaborative moral problem solving. These rules, according to Clark (1996, p. 115), are "summary moral rules", rules that are a summary of our past moral judgments.

Why are summary moral rules compatible with connectionism, it might rightly be wondered. To this question, Clark does not supply an answer. Yet, the answer can be gleaned from Churchland (2000), who argues that connectionism, due to its commitment to natural shapelessness of the moral, rules out the existence of overarching moral rules (e.g. the principle of utility, etc.) that purport to unify morality. But since summary moral rules are not moral rules of this sort, they are not the proper targets of connectionism; connectionism may well be reconciled with them. Now, Clark, qua connectionist, agrees (or can agree) with all this. But Churchland goes on to suggest that the status of the rules is merely peripheral. It is here that Clark's view diverges from Churchland's. Clark is of the view that they still play a central role in our moral perception.

The marginalization of such rules, according to Clark, is due to a misconstrual of the central purposes of the rules. Clark urges that the rules are not to be seen as tools that manipulate people's moral perception in the way suggested by Churchland or many fellow connectionists. That is, the rules are not meant to subsume people's perception of natural features under a particular moral category. They are not meant to determine what people morally perceive.

Rather, the rules are to be seen as tools for collaborative moral problem solving, Clark urges. This is because Clark thinks that for

---

[18] On a certain interpretation, Clark (2000, p. 275) regards moral rules as public social artifacts that are external to our skulls and biological brains

collaborative moral problem solving to be possible, rules are indispensable. For instance, in a conflict within a multicultural educational program, where Muslim parents want their daughter to be excused from events involving physical proximity to boys, whereas the head teacher is inclined to let the child decide, Clark (p. 122) maintains that the only solution would lie in each party articulating the moral rules or principles that inform their perspectives. But surely the rules or principles might well conflict, it may be wondered. Clark admits this, but argues that the articulation of moral principles which one is partisan with is the "*essential starting point*" for informed moral debate or collaborative moral problem solving.

But is it? Can't we engage in collaborative moral problem solving without rules or principles? A 'no' answer seems far-fetched. For the reasons offered by the disputants may not take the form of moral rules or principles. Rather, they more often take the form of a narrative or the particular considerations that are unique to one's circumstances. In the case of the Muslim parents, the reasons might well take the narrative form of a fable or a parable in their scriptures. As noted by David McNaughton (1988, p. 205): "What the great moral teachers of the past, such as Buddha and Christ [or Muhammad for that matter], have done is to bring us, by parable, story and paradox, to see the world in new and revealing ways. What we need is not a better set of principles but better moral vision."

It has to be noted here that I agree with Clark that connectionism does not warrant the marginalization of the summary rules. But this is not for the reasons stated by Clark. That is, the summary rules are not indispensable for collaborative problem-solving. Rather, it is because of the mundane fact that the summary rules are useful for individual decision-making in cases of emergency, where time is too pressed for deliberations. However, even granted an important role to play in emergent situations or collaborative moral problem-solving, the summary rules do not seem to harm particularism about moral perception at all. After all, given the summary nature of the moral rules, the particularists might well maintain without incoherence that while they are useful for informed moral debate, moral perception is not essentially determined or governed by them.

4. *Practical Wisdom and Rules*

McKeever and Ridge (henceforth M and R for short) in their thought-provoking book, *Generalism as a Regulative Ideal* (2006), rail against particularism and argues that the moral perception of the

practically wise (or what they call 'the moral experts') has to be rule-governed for their practical wisdom to be possible. The practically wise, according to M and R, are people who can reliably perceive the morally relevant features of their circumstances. However, they are also cognitively limited beings like you and me. So if the number of potentially morally relevant features were infinite or finite but cognitively unmanageable, the practically wise would not perceive so reliably the morally relevant features. So, according to M and R, the number of potentially morally relevant features must be not only finite but also cognitively manageable so that the practically wise could reliably discern them should they arise. And since they are finite and cognitively manageable, they can be specified in terms of rules. Conclusion: the moral perception of the practically wise is governed by a finite and cognitively manageable set of rules.

The sort of rules M and R have in mind is *unhedged* and *Rossian* in nature.[19] For W.D. Ross, there are exactly seven morally relevant features specified by rules of prima facie duty. For M and R, Ross might or might not overlook a few morally relevant features that ought to be included in the list of the rules of prima facie duty, but they think Ross is right to regard the list as finite and manageable. The point where M and R depart from Ross is where Ross holds an atomistic view of moral relevance, according to which the moral relevance of the features specified by rules of prima facie duty remains the same across different contexts whereas M and R are holists who maintain that their moral relevance might change from one context to another.

Given their holism, one question immediately arises. That is, are their rules really unhedged? For instance, although promise-keeping is normally positively valenced, it is not so, when it is coerced. M and R agree with this, but contend that the rules are unhedged in the sense that they are not open-ended. In fact, they think that the exceptional conditions can be fully articulated. For instance, the rule regarding promise-keeping might look like this: promise-keeping is

---

[19] The 'unhedged' 'Rossian' rules may sound oxymoronic. After all, the Rossian rules are generally understood to be hedged in the sense that they are merely 'prima facie' (or better, 'pro tanto'), purporting to specify the moral status of morally relevant features, rather than the overall moral status of actions. As I see it, this is a sensible terminological issue one may reasonably put to M and R. Yet, to be fair to them, when M and R regard these Rossian rules as 'unhedged', they do not take them to be specifying the overall moral status of actions, but rather to be 'not open-ended', as I will explain in the next paragraph. In short, in M and R's usage, 'unhedged rules' mean rules that are not open-ended, not rules that specify the overall moral status of actions.

positively valenced, unless it is coerced. But this rule has exceptions
too, for instance, when I coerce you into promising that you would
not harm the innocent. M and R anticipate this and argue that this
exception can be built into the rule as well. For instance, we may have
the following rule instead: promise-keeping is positively valenced,
unless it is coerced, except when the coercion is to prevent innocent
lives from being harmed. But as it stands, it looks too unwieldy
and overly-specific to be of much use in application. For instance,
promise-keeping is also not positively valenced, when the content of
the promise is to kill your friend's enemy for him. But this is not
covered by the above-mentioned rule. It might well be retorted by M
and R that there is another independent rule to cover this. But if so,
the number of rules can quickly expand out of control. For instance,
promise-keeping is not positively valenced when the content of the
promise is to strangle your friend's neighbor's cat. Of course, to
bring the number of rules under rein, M and R can reply by making
the moral rule more abstract: promise-keeping is always positively
valenced unless its content is immoral. However, there are still cases
that are not covered. For instance, promise-keeping doesn't seem
to be positively valenced, when the promisee has explicitly asked
the promisor not to keep his promise, or when you made a foolish
promise to yourself, etc. The number of rules regarding promise-
keeping alone might well still be finite, but it is unclear whether
the total number of rules is still cognitively manageable, given that
the rules regarding promise-keeping are not the only ones, and other
numerous rules have their own numerous exceptional conditions.

To bring the number of the potentially morally relevant features
under rein, M and R might even go more abstract. For instance,
we may have the following abstract rule regarding promise-keeping:
promise-keeping is always positively valenced unless there is a good
explanation for why it is not so. But this sort of principle is too weak
to harm particularism. It amounts to little more than the hedged
default rules much welcomed by the particularists, according to
which promise-keeping is positively valenced by default unless there
is a good explanation for why it is not so. In fact, many particularists
embrace moral rules of this sort as very useful rules of thumb (Dancy
2004) or presumptive rules, rules that specify certain natural features
as presumptive reasons unless undermined (Cullity 2002). As granted
by M and R themselves (2006, pp. 138–139), the particularists can
still live on in the hedged defaults. So this escape route is not
available to M and R.

Furthermore, even if, against all the odds, the realm of morality might well be codified by a finite and manageable set of rules, it doesn't seem to be the case that the practically wise perceive the moral salience of features through them. The exceptional escape clauses are simply irrelevant in many cases. When I just promised my wife to meet her for dinner after work, it never crossed my mind to check whether my promise is coerced (thank God!) or if the content of my promise is immoral. The fact that the practically wise can detect the morally relevant features reliably doesn't seem to be due to their perceiving through rules. Phenomenologically speaking, I simply see keeping my promise to meet my wife for dinner as positively valenced (or something to be done) without inference through rules of the complicated sort advocated by M and R. The rule that promise-keeping is positively valenced if its content is not immoral, even if it is exceptionless, doesn't seem to have much bearing on the phenomenology of the current case.

Now, it might well be objected, pace Väyrynen (2008) and Wodak (2019), that the inference through rules of the complicated sort might be *implicit* without our conscious awareness of its operation. In reply, I concur largely with Philip Stratton-Lake (2000, p. 127), who perceptively observes that the implicit rule-governed account that we perceive implicitly through rule-inferences is "artificial" because "none of us [perceive] like this, or at best do not think that we do, and we cannot get ourselves to believe that we should". Here, I don't assume that Stratton-Lake is right. I'm open to the suggestion that he is wrong. Yet, the fact that this observation is made by Stratton-Lake, an eminent Kantian whom you would least expect to be a detractor of the (implicit) rule-governed account of moral perception, does seem to suggest that the (implicit) rule-governed account may overreach itself. And I will argue three paragraphs later, this implicit rule-governed account of moral perception, as well as the explicit one, are both vulnerable to the so-called 'frame problem'.[20] For now, let's turn to a more aggressive (rather than defensive) strategy that M and R might employ when edged into the corner.

Instead of negatively defending their position, M and R may actively charge against particularism. They might argue, for instance, that if the number of the potentially morally relevant features were infinite or cognitively unmanageable, practical wisdom would thus become impossible, given the fact that the practically wise are cognitively limited beings. This inference is, however, too hasty. For,

---

[20] See the next footnote for an illustration of the frame problem.

from the particularists' perspective, which in my view is not un-reasonable, the practically wise, qua cognitively limited beings, do most of the time reliably perceive the moral salience, even if the number of the potentially morally relevant features can be large or infinite. It is just that the particularists owe us an account for why the practically wise can perceive reliably the morally salient features they encounter in the moral situations, given the fact that they are cognitively limited. And, contra M and R, I don't think such an account is unavailable for the particularists.

The particularists' story might go as follows: Even if the number of the potentially morally relevant features is infinite or very large, they are not all instantiated all at once on a single occasion. The number of the potentially morally salient features that get instantiated on a particular occasion is generally quite limited, as admitted by M and R themselves. This being so, in order for the moral perception of the practically wise to be reliably accurate, it only needs to be the case that they are sensitive to the morally relevant features that turn up in the particular circumstances. To use an aesthetic analogy to illustrate, although the number of features that might influence our aesthetic perception of an artwork can in principle be infinite too, this doesn't mean that no one can thus reliably discern the aesthetically relevant features that actually get instantiated in an artwork. An experienced art critic or a connoisseur may still reliably notice the aesthetically relevant features. And there is no magic involved.

Yes, given reason holism, there is an infinite number of features that are normally irrelevant but might come to acquire relevance on a particular occasion. And if we have to check through them all to ensure that none is lost on us, then, due to the fact that we are cognitively limited and finite beings, we apparently wouldn't be able to reach a moral conclusion at the end of the day. But the fact that we do suggests that we do not (have to) check, explicitly or implicitly (without our conscious awareness), through them all to be reliable in the detection of the morally relevant features. The so-called 'frame problem' is a well-known problem for AI to distinguish situationally relevant features from irrelevant ones, but as such it is widely ac-cepted that it is not a problem for human beings.[21] The structure of

---

[21] See Dennett (1992, pp. 147–148) for an illuminating account of the frame problem. In Dennett's famous example, the AI, designed to remove a time bomb on the wagon before pulling the wagon out of a room, is operating with a pre-set-up frame (or rules) of relevance, and is still busy checking (explicitly or implicitly) through all the circumstantial features (which are infinite in number) to determine whether they fall under the frame (or rules) of relevance, when the bomb goes off.

our aesthetic perception, and moral perception for that matter, is constituted in a way such that they are generally sensitive to how a feature that is normally irrelevant comes to acquire its aesthetic or moral significance by interacting with other aesthetically or morally relevant features on a particular occasion. To reverse back to our explanation of our perception of the morally relevant in section 2, it might come down to our competence with the concept of 'moral relevance' in a particular form of life, which enables us to perceive a feature that is normally irrelevant (such as the feature of wearing yellow shoelaces) as morally relevant when it figures into the content of a promise.

Furthermore, there is an advantage to the particularist story in terms of moral phenomenology. For the biggest problem with the generalist account, as we have argued, is that even if the morally relevant features can be specified by a finite and cognitively manageable set of rules, we, insofar as we are the practically wise in the sense specified by M and R, do not seem to perceive the moral relevance of natural features through them (or via rule-inferences). Insofar as the natural features impinge on our moral sensitivities, we simply see them as morally relevant. To use an analogy to illustrate, when a joke strikes us as funny, it's not as if we infer its funniness through a rule about what counts as a funny joke. Even if there were such a rule (which is rather doubtful), we probably would not perceive the funniness of the joke through it. Rather, we would simply perceive its funniness if we were equipped with the right sense of humor.[22] Similarly, we would simply perceive the torture of a kitten as cruel or wrong-making, if we were furnished with the right sort of moral sensitivities.

It is to be noted here that while our moral sensitivities might well be informed by moral rules, they are not bound by them.[23] It is

It is important to note that the frame problem is a problem for AI, but not for us, for we can normally quickly identify what is of relevance.

[22] LaFollette and Shanks (1993, p. 330) advocates a similar view. Thanks to one reviewer for bringing this article to my attention. I should also add that, my view here is inspired by Nussbaum (2001, p. 303): "Excellent choice cannot be captured in universal rules, because it is a matter of fitting one's choice to the complex requirements of a concrete situation, taking all of its contextual features into account. A rule, like a manual of humor, would do both too little and too much: too little, because most of what really counts is in the response to the concrete, and this would be omitted; too much, because the rule would imply that it was itself normative for response (as a joke manual would ask you to tailor your wit to the formulae it contains) —and thus would impinge too much on the flexibility of good practice."

[23] This is very much in line with Murdoch's particularism, as interpreted by Driver (2012, p. 305).

not (conformity with) the moral rules that make the features of the actions right-making or wrong-making just as it is not (conformity with) the joke rule that makes the joke funny.[24] Rather, it is the particular features of the joke (e.g. who tells the joke, who is the listener, on what occasion is the joke told, and the content of the joke, etc.) that make it funny.[25] Similarly, it is the particular features of abusing the cat (e.g. setting fire to it, beating it relentlessly) that make it wrong.

Before we conclude, there is one more objection regarding the scope of the particularist moral sensitivities we have to address. According to M and R (2006, pp. 140–158), the particularist moral sensitivities are pretty much shaped by local culture, so might not detect the morally relevant features in a different culture. Hence they are rather insular and limited in scope. However, this objection is really beside the point. If it were really an objection to the particularists, it would equally constitute an objection to the generalists too. For the rules regarding moral relevance might well be shaped by our local culture too. They might not apply in a different culture.[26] The concern about moral relativism is no more of an objection to the particularists than to the generalists. It is a common problem both of them have to face and doesn't prejudge in favor of either.[27]

5. Conclusion

In this article, we have examined four very powerful generalist arguments in favor of rule-governed moral perception; however, as we have argued, none of them succeeded. This of course doesn't

[24] Here, I believe my view departs from Stratton-Lake's 2000 (pp. 126–127), which maintains that the moral rules are the 'transcendental conditions' that make a feature wrong-making or right-making. As I will argue below, it is the particulars that make it so.

[25] My view here is entirely consistent with the view that our sense of humor might well be wrong, as emphasized by one reviewer, citing Smuts (2010), and George and Richardson (2019). That is, I can be partial to the view, without incoherence, that although one may in fact find sexist or racist jokes funny, one *shouldn't*. All I want to emphasize here is that what makes jokes funny is not due to their conformity to the purported rules, whose existence is quite dubious, about what count as funny jokes.

[26] The rule against eating pork, while being a rule of moral relevance in Muslim culture, does not apply in most non-Muslim cultures.

[27] For readers who are interested in the prospects for particularists to develop 'critical moral perception' that transcend above parochialism, see Clarke (2012) and Flanagan (2015) for illuminating accounts.

prove particularism to be right.[28] But at least this helps boost our confidence in particularism. This article is not merely a negative defense of particularism. Along the way, we have in fact produced many positive arguments to support particularism. What emerges is at least a prima facie plausible particularist conception of moral perception. What speaks in favor of it, first of all, is that it is compatible with the natural shapelessness of morality and does not, as generalism does, impose unwarranted constraints in the form of rules of moral salience on our moral perception. Our moral perception is not bound by the rules of moral salience as advocated by Herman. Second, phenomenologically speaking, particularism seems to fit better with our moral phenomenology than generalism. We do not seem to perceive through rule-inferences, not even *implicitly*. For if we were to perceive (correctly) through rule-inferences, then the rule would have to be very specific, excluding all the possible exceptions, in order for our moral perception not to degenerate into moral illusion/hallucination.[29] But this is not how we typically per-

---

[28] For, as the reviewers rightly pointed out, other versions of generalism that haven't been considered in this article might remain alive. And there is even the possibility that arguments stronger than the ones advocated by Herman, the Canberrans, Clark, and Ridge and McKeever remain to be discovered for the specific version of generalism considered in this article. That's why I suggest in the title of my paper that my defense of particularism is 'limited'. This article might well be seen as issuing a challenge to the generalists to either come up with stronger arguments for the specific version of generalism at issue, or articulate a more plausible version of generalism that can bear scrutiny.

[29] One reviewer correctly observed here that moral perception, for the rule-governed generalist account, would thus require the existence of exceptionless rules or RMS. As a result, Herman's RMS, if they are to be understood to be exceptionless, must be meant to be sufficient for moral perception. And this seems to render my earlier discussions of Herman's 'necessary claim' somewhat off-target.

In reply, I think the reviewer has put his/her finger on something very subtle and significant here. And I very much appreciate the reviewer's suggestion. After all, if Herman could not treat the possession of RMS as merely necessary, but must treat it as sufficient for moral perception, then she would be impaled on one horn of the dilemma I posed to her (sorry about the horrid metaphor).

Recall that I argue in section 2 of my article that the sufficiency point cannot go through, because the possession of (exceptionless) RMS cannot guarantee that the moral agent would thus be enabled to perceive the moral salience of the features specified by RMS, whenever these features turn up. At the risk of oversimplification here, the moral agent, due to his physical fatigue or illness, I argue in section 2, might well fail to perceive the moral salience of the RMS features, despite his possession of the RMS and his serious commitment to them.

If Herman were impaled on the horn of the sufficiency point already, this would actually be favorable to my defense of particularism (or my critique of Hermanian generalism). So I actually love the reviewer's perceptive observation. The reason(s)

ceive the moral salience of a natural feature, for, phenomenologically speaking, we do not check all the possible exceptions in normal cases. Nor is it the case that our mind has *implicitly* checked all the possible exceptions for us without our conscious awareness. For, in a rule-governed framework, in order to ensure that none of the possible exceptions has arisen, we would be returned once again to the 'frame problem' faced by AI. And it is important to recognize that this is a hurdle AI has to cross over, but not a problem for a normal-functioning human being.[30] Third, particularism about moral perception has the support of connectionism, which is a respectable doctrine in cognitive sciences. The Canberrans' attack on connectionism can be neutralized. Clark's charge against connectionism's implications for the marginalization of the rules can be addressed. Fourth, particularism can give a plausible story regarding how we can perceive the moral salience of a natural feature, given the fact that the number of potentially morally relevant features is infinite or large. In the light of the pitfalls of generalism and the attractions of particularism, I think particularism remains as a serious contender for the right theory of moral perception.[31]

why I go on to consider the necessary point (or the other horn of the dilemma) is because I want to be as charitable to Herman as possible and see if Herman still has some room to wriggle (again, apologies for the unpalatable metaphor). And Herman does seem to have this wriggle room, at least initially. And this is for the following reason.

In order to hold on to the idea that the RMS are exceptionless, Herman need not be committed to the idea that the possession of RMS is sufficient for moral perception. She could of course be so committed, if she wished. Yet, as I have argued, her sufficiency point would then stand to be refuted. So, a natural escape route for Herman is to argue/clarify that the exceptionless RMS are only meant to be necessary (rather than sufficient) for moral perception (That is, if one were to perceive the moral salience of lying, let's say, one would have to perceive it through the lens of the exceptionless RMS about lying).

As I see it, the necessary point is indeed available to Herman, even if the RMS are exceptionless, although, ultimately, I argue in section 2 that this necessary point cannot go through, for reasons that are related to McDowllian/Wittgensteinian 'forms of life'. In the last analysis, I think Herman would be impaled on the horn of the necessary point as well, if she were to adopt it.

[30] It is an undeniable fact that, unlike AI, in most cases we can immediately grasp what is situationally salient without getting paralyzed by checking explicitly or implicitly through myriads of irrelevancies.

REFERENCES

Audi, Robert, 2013, *Moral Perception*, Princeton University Press, N.J.

Bagnoli, Carla, 2011, "Moral Perception and Knowledge by Principles", *New Intuitionism*, Jill Hernandez (ed.), Continuum, New York, pp. 84–105.

Bergqvist, Anna and Robert Cowan (comps.), 2018, Oxford University Press, Oxford.

Blum, Lawrence, 1991, "Moral Perception and Particularity", *Ethics*, vol. 101, no. 4, pp. 701–725.

Boder, Margaret (comp.), *The Philosophy of Artificial Intelligence*, Oxford University Press, Oxford.

Brewer, Talbot, 2000, *The Bounds of Choice*, Garland Publishing, New York.

Breyer, Thiemo and Christopher Gutland (comps.), *Phenomenology of Thinking*, Routledge, New York.

Broakes, Justin (comp.), 2012, *Iris Murdoch: Philosopher*, Oxford University Press, Oxford.

Chudnoff, Elijah, 2016, "Moral Perception: High-level Perception or Low-level Intuition?", in Thiemo Breyer and Christopher Gutland (eds.), *Phenomenology of Thinking*, New York, Routledge, pp. 207–219.

Churchland, Paul, 2000, "Rules, Know-How, and the Future of Moral Cognition", *Canadian Journal of Philosophy*, Supplementary Volume 26, pp. 291–306.

Clark, Andy, 2000, "Word and Action: Reconciling Rules and Know-how in Moral Cognition", *Canadian Journal of Philosophy*, Supplementary Volume 26, pp. 267–289.

———, 1996, "Connectionism, Moral Cognition, and Collaborative Problem Solving", in Larry May et al. (eds.), *Mind and Morals*, MIT, Cambridge, M.A., pp. 109–127.

Clarke, Bridget, 2012, "Iris Murdoch and the Prospects for Critical Moral Perception", in Justin Broakes (ed.), *Iris Murdoch: Philosopher*, Oxford University Press, Oxford, pp. 196–253.

Cowan, Robert, 2014, "Cognitive Penetrability and Ethical Perception", *Review of Philosophy and Psychology*, vol. 6, no. 4, pp. 665–682.

Cullison, Andrew, 2009, "Moral Perception", *European Journal of Philosophy*, vol. 18, no. 2, pp. 159–175.

Cullity, Garrett, 2002, "Particularism and Moral Theory: Particularism and Presumptive Reasons", *Aristotelian Society Supplementary Volume*, vol. 76, issue 1, pp. 169–190.

Dancy, Jonathan, 2010, "Moral Perception", *Proceedings of the Aristotelian Society*, Supplementary Volume, vol. 84, no. 1, pp. 99–117.

———, 2004, *Ethics Without Principles*, Oxford University Press, Oxford.

———, 1999, "Can a Particularist Learn the Difference Between Right and Wrong?", *The Proceedings of the Twentieth World Congress of Philosophy*, vol. 1, pp. 59–72.

———, 1993, *Moral Reasons*, Blackwell, Oxford.

———, 1983, "Ethical Particularism and Morally Relevant Properties", *Mind*, vol. 92, no. 368, pp. 530–547.

Darwall, Stephen, 1998, *Philosophical Ethics*, Westview, Colorado.

Dennett, Daniel, 1992, "Cognitive Wheels: The Frame Problem of AI", in Margaret Boder (ed.), *The Philosophy of Artificial Intelligence*, Oxford University Press, Oxford, pp. 147–170.

Driver, Julia, 2012, " 'For Every Foot Its Own Shoe': Method and Moral Theory in the Philosophy of Iris Murdoch", in Justin Broakes (ed.), *Iris Murdoch: Philosopher*, Oxford University Press, Oxford, pp. 293–305.

Faraci, David, 2015, "A Hard Look at Moral Perception", *Philosophical Studies*, vol. 172, no. 8, pp. 2055–2072.

Flanagan, Owen, 2015, "Foreword: Three Kinds of Moral Sensitivity", in Deborah Mower et al. (eds.), *Developing Moral Sensitivity*, Routledge, London, pp. 9–14.

Garfield, Jay, 2000, "Particularity and Principle", in Brad Hooker and Margaret Little (eds.), *Moral Particularism*, Oxford University Press, Oxford, pp. 178–204.

George, Shane and C.L. Richardson, 2019, "The Feministecles: A Dialogue on Humor and Sexism", *Israeli Journal for Humor Research*, vol. 8, no. 1.

Herman, Barbara, 1985, "The Practice of Moral Judgment", *The Journal of Philosophy*, vol. 82, no. 8, pp. 414–436.

Hernandez, Jill (comp.), 2011, *New Intuitionism*, Continuum, New York.

Hooker, Brad and Margaret Little (comps.), 2000, *Moral Particularism*, Oxford University Press, Oxford.

Jackson, Frank, Philip Pettit and Michael Smith, 2000, "Ethical Particularism and Patterns", in Brad Hooker and Margaret Little (eds.), *Moral Particularism*, Oxford University Press, Oxford, pp. 79–99.

Jacobson, Daniel, 2005, "Seeing by Feeling: Virtues, Skills, and Moral Perception", *Ethical Theory and Moral Practice*, vol. 8, pp. 387–409.

Kagan, Shelly, 2001, "Thinking about Cases", *Social Philosophy and Policy*, vol. 18, no. 2, pp. 44–63.

Korsgaard, Christine, 1996, "Kant's Formula of Humanity", *Creating the Kingdom of Ends*, Cambridge University Press, Cambridge, pp. 106–132.

Kukla, Rebecca, 2002, "Attention and Blindness: Objectivity and Contingency in Moral Perception", *Canadian Journal of Philosophy*, Supplementary Volume 28, pp. 319–346.

LaFollette, Hugh and Niall Shanks, 1993, "Belief and the Basis of Humor", *American Philosophical Quarterly*, vol. 30, no. 4, pp. 329–339.

May, Larry et al. (comps.), *Mind and Morals*, MIT, Cambridge, M.A.

McBrayer, Justin, 2010, "Moral Perception and the Causal Objection", *Ratio*, vol. 23, pp. 291–307.

McDowell, John, 2002a, "Virtue and Reason", in *Mind, Value and Reality*, Harvard University Press, Massachusetts, pp. 50–73.

———, 2002b, "Non-cognitivism and Rule-following", in *Mind, Value and Reality*, Harvard University Press, Mass., pp. 198–218.

McGrath, Sarah, 2018, "Moral Perception and Its Rivals", in Anna Bergqvist and Robert Cowan (eds.), *Evaluative Perception*, Oxford University Press, Oxford, pp. 161–169.

———, 2004, "Moral Knowledge by Perception", *Philosophical Perspectives*, vol. 18, no. 4, pp. 209–228.

McGregor, Rafe, 2015, "Making Sense of Moral Perception", *Ethical Theory and Moral Practice*, vol. 18, pp. 745–758.

McKeever, Sean and Michael Ridge, 2006, *Principled Ethics: Generalism as a Regulative Ideal*, Oxford University Press, Oxford.

McNaughton, David, 1988, *Moral Vision*, Blackwell, Oxford.

Millgram, Elijah, 2005, "Murdoch, Practical Reasoning, and Particularism", *Ethics Done Right*, Cambridge University Press, Cambridge, pp. 168–197.

Mower, Deborah et al. (comps.), *Developing Moral Sensitivity*, Routledge, London.

Mower, Deborah and David McGraw, 2015, "The Normativity of Moral Sensitivity", in Deborah Mower et al. (eds.), *Developing Moral Sensitivity*, Routledge, London, pp. 264–279.

Murdoch, Iris, 1991, *Sovereignty of Good*, Routledge, New York.

Nussbaum, Martha, 2001, "Non-Scientific Deliberation", *The Fragility of Goodness*, Cambridge University Press, Cambridge, pp. 290–317.

———, 1992, "Finely Aware and Richly Responsible: Literature and the Moral Imagination", *Love's Knowledge*, Oxford University Press, Oxford, pp. 148–167.

Ridge, Michael and Sean McKeever, 2016, "Moral Particularism and Moral Generalism", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/win2016/entries/moral-particularism-generalism/> [date of consultation: 04/11/2020]

Smuts, Aaron, 2010, "The Ethics of Humor: Can Your Sense of Humor Be Wrong?", *Ethical Theory and Moral Practice*, vol. 13, pp. 333–347.

Stratton-Lake, Philip, 2000, *Kant, Duty and Moral Worth*, Routledge, New York.

Tsu, Peter Shiu-Hwa, 2018, 'Particularism in Ethics', in Duncan Pritchard (ed.), *Oxford Bibliographies Online*, Oxford University Press, Oxford, <http://oxfordbibliographies.com/view/document/obo-97801953965 77/obo-9780195396577-0367.xml> [date of consultation: 04/11/2020]

———, 2013, "Shapelessness and Predication Supervenience: A Limited Defense of Shapeless Moral Particularism", *Philosophical Studies*, vol. 166, no. 1, pp. 51–67.

Väyrynen, Pekka, 2018a, "Reasons and Mora Principles", in Daniel Star (ed.), *Oxford Handbook of Reasons and Normativity*, Oxford University Press, Oxford, pp. 839–861.

———, 2018b, "Doubts about Moral Perception", *Evaluative Perception*, in Anna Bergqvist and Robert Cowan (eds.), Oxford University Press, Oxford, pp. 109–128.

———, 2008, "Usable Moral Principles", in Mark Lance et al. (eds.), *Challenging Moral Particularism*, Routledge, London.

Watkins, Michael and Kelly Jolley, 2002, "Pollyanna Realism: Moral Perception and Moral Properties", *Australasian Journal of Philosophy*, vol. 80, no. 1, pp. 75–85.

Werner, Preston, 2020, "Moral Perception", *Philosophy Compass*. (doi:10.1111/phc3.12640)

———, 2018, "Moral Perception without (Prior) Moral Knowledge", *The Journal of Moral Philosophy*, vol. 15, no. 2, pp. 164–181.

Wisnewski, J., 2015, "The Case for Moral Perception", *Phenomenology and the Cognitive Sciences*, vol. 14, pp. 129–148.

Wodak, Daniel, 2019, "Moral Perception, Inference, and Intuition", *Philosophical Studies*, vol. 176, no. 6, pp. 1495–1512.