



BRILL
NIJHOFF

JOURNAL OF MORAL PHILOSOPHY (2014) 1-24

JOURNAL OF
MORAL
PHILOSOPHY



brill.com/jmp

Particularism Doesn't Flatten

Amelia Hicks

Visiting Assistant Professor, Department of Philosophy,
Kansas State University
ahicks2@k-state.edu; amelia.hicks@gmail.com

Abstract

Sean McKeever and Michael Ridge object that moral particularism 'flattens the moral landscape', that is, that particularism treats reasons of *different* kinds as if they were reasons of the *same* kind. This objection is misguided in two respects. First, particularists need not say that every feature can be a moral reason. Second, even if particularists *were* committed to saying that every feature can be a moral reason, they would still not be committed to the view that every feature can have *direct* moral relevance. The failure of this objection shows that the objection exploits side-constraints that need not be placed on moral particularism.

Keywords

Sean McKeever – particularism – reasons – Michael Ridge

Introduction

Sean McKeever and Michael Ridge write:

One danger that has not been lost on particularists is that their extremely ecumenical view of moral reasons for action threatens implausibly to 'flatten the normative landscape'. After all, even if we think that in the

Amelia Hicks is a Visiting Assistant Professor of Philosophy at Kansas State University. She received her Ph.D. in Philosophy from the University of Notre Dame. Her research interests include moral particularism, moral deliberation, subjective obligations, and animal ethics.

right context... shoelace color can provide a reason for action (which, to be clear, we do not), there is an important difference between considerations of shoelace color and considerations of pain, pleasure, promising, and the like.¹

McKeever and Ridge have regularly² attacked moral particularism on the grounds that it treats moral reasons of obviously *different* kinds as if they were reasons of the *same* kind (or, as McKeever and Ridge put it, that particularism ‘flattens the normative landscape’). The particularist response to these attacks has varied widely; some particularists have attempted to amend their theory so as to better accommodate McKeever and Ridge’s intuitions about types of reasons, while others have claimed that this allegedly objectionable feature of particularism is simply one of the core tenets of the view, and that McKeever and Ridge are begging the question against particularists.

In this paper, I will set aside the worry about question-begging and argue that moral particularists can accommodate McKeever and Ridge’s intuitions without adding further commitments to their theory. I will show that (a) McKeever and Ridge’s objection applies only to versions of particularism that include a certain set of side-constraints and that (b) particularists can, in fact, do *without* those side-constraints.

In the first section, I will provide some background on the debate about the flattening objection. (Those readers who are already familiar with the debate can skip section 1.1.) I will then provide a diagnosis of the debate, and suggest what I take to be a better resolution of it. In the second section, I will describe a ‘minimal’ version of moral particularism—a version of moral particularism that includes the core tenets of the view (and thus issues a challenge to many generalist approaches to moral theory), but includes very few further non-entailed commitments. In section three, I present what I take to be the most charitable reading of McKeever and Ridge’s flattening objection. In sections four and five, I show how the minimal particularism described in section two avoids McKeever and Ridge’s objection. In section six, I address several worries about my approach. Finally, I conclude with a discussion of what my defense of minimal particularism can tell us about the particularism debate more generally.

1 Sean McKeever and Michael Ridge, *Principled Ethics: Generalism as a Regulative Ideal* (New York: Oxford Univ. Press, 2006), p. 47.

2 See also Sean McKeever and Michael Ridge, ‘Turning on Default Reasons’, *Journal of Moral Philosophy* 4.1 (2007), pp. 55–76.

1 Some Background

1.1 *Initial Exchanges*

After the publication of Jonathan Dancy's *Ethics Without Principles*,³ a conference on moral particularism was held at the University of Kent. It was at this conference that McKeever and Ridge first presented their worry that moral particularism 'flattens the moral landscape'. Their worry—which I will describe in greater detail shortly—was this: given that particularists are open to many different kinds of contextual features potentially acting as moral reasons, particularists must admit that those features, whether they be *pain* or *shoelace color*, are not in fact very different from each other. According to McKeever and Ridge, this result is at odds with our common sense intuitions about moral reasons, and thus counts against particularism.

Different particularists have had different sorts of responses to this worry. Jonathan Dancy, in response, has suggested that his account of default reasons (from *Ethics Without Principles*) can deal with the flattening objection. According to this account of default reasons, certain features have 'default' moral valences,⁴ and contextual features can override those default valences. A feature such as *pain* arguably has a *negative* default moral valence, while a feature such as *shoelace color* arguably has a *neutral* default moral valence; thus, although both features are potential moral reasons (and even potential moral reasons of the same valence), we can account for the intuition that they're different by examining their *default* valences.⁵ In a similar vein, Margaret Little and Mark Lance have argued that particularists can account for the apparent difference between *pain* and *shoelace color* by appealing to a theory of defeasible generalizations (generalizations that have explanatory power but admit of exceptions).⁶

But while Dancy, Lance, and Little have been concerned with accommodating McKeever and Ridge's intuition concerning the difference between features such as *pain* and *shoelace color*, other particularists have bitten the bullet. For example, Alan Thomas responded 'The normative landscape is flat' and

3 Jonathan Dancy, *Ethics Without Principles* (Oxford: Oxford University Press, 2005).

4 A feature has a negative 'moral valence' relative to some action in a context just in case? it morally counts against that action in that context; a feature has a positive 'moral valence' relative to some action in a context just in case it morally counts in favor of that action.

5 For an account of how particularist default reasons might work, see Dancy, *Ethics Without Principles*, pp. 111–117; and Dancy, 'Defending the Right', *Journal of Moral Philosophy*, 4.1 (2007), pp. 85–98.

6 Mark Lance and Margaret Little, 'Defeasibility and the Normative Grasp of Context', *Erkenntnis* 61 (2004), pp. 435–55.

thus '[i]t is no objection to particularism that it represents the normative landscape as flat...'⁷

Thus, the current status of the debate about flattening is this: McKeever and Ridge have argued that particularists are committed to the flattening of the normative landscape, and that this fact counts against particularism. In response, some particularists have argued that they can supplement their view and thereby avoid flattening the normative landscape, while others have argued that the flattening of the normative landscape is not bad at all—in fact, the normative landscape *is* flat, and the fact that McKeever and Ridge think it is not flat is simply a symptom of generalist bias. I do not intend to argue that this latter group of particularists is wrong; it may be that the moral landscape is flat, and that the conclusion of McKeever and Ridge's *reductio* isn't absurd at all. What I will do instead is put particularism into a rhetorically stronger position by arguing that one can adopt an interesting version of moral particularism even if one *denies* that the moral landscape is flat.

1.2 *Diagnosis*

What I would like to suggest here—and argue for in the remainder of this paper—is that the debate about the flattening objection has taken this course because of its focus on just a few canonical versions of particularism. These canonical versions—especially those versions held by Dancy and Little—are certainly worthy of discussion. However, these canonical particularist views include additional commitments that (a) are not essential to particularism and (b) make the views susceptible to McKeever and Ridge's flattening objection. The commitments I have in mind are:

Relevance: every property is potentially morally relevant; and

Direct Relevance: every property is potentially directly morally relevant.

To understand what Relevance means, begin by imagining a scenario in which some feature of the context in that scenario is, intuitively, morally relevant. For instance, imagine that Pat assaults a stranger on the street. Intuitively, the fact that the assault causes pain to the assault victim is morally relevant, because the fact that the assault causes pain *makes a moral difference*—had the assault caused no pain whatsoever, the moral facts would have been different. The same goes for Pat's false but justified belief that she was acting in

⁷ Alan Thomas, 'Practical Reasoning and Normative Relevance: A Reply to McKeever and Ridge', *Journal of Moral Philosophy*, 4.1 (2007), p. 77 and p. 84, respectively.

self-defense—that belief of hers makes a moral difference. Moreover, we might even say that the assault victim's visible neo-Nazi tattoo was also morally relevant, because it was partially responsible for Pat's false but justified belief that she was acting in self-defense. Thus, although the idea of 'moral relevance' is difficult to pin down, we can think of a morally relevant feature in a context as a feature that makes a moral difference in that context.⁸ Relevance, then, says that *any given feature* can, in some possible context, make a moral difference.

Direct Relevance is stronger than Relevance—it says not only that every feature is potentially morally relevant in some context, but that every feature is potentially *directly* morally relevant in some context. To understand what this is supposed to mean, think back to the previous scenario. Why do we think that Pat's causing pain to the stranger makes a moral difference? Many of us would say that it *just does*—in this case, pain makes some sort of moral difference, but not because it makes any *other* sort of difference. Pain, then, is distinct from the neo-Nazi tattoo—presumably the tattoo makes a moral difference only because it makes other sorts of differences (it expresses morally reprehensible views, causes pain and fear, etc.). We might say, then, that in this context *the fact that Pat caused pain* is directly morally relevant, while *the fact that the assault victim had a neo-Nazi tattoo* was indirectly morally relevant—the latter feature of the context makes a moral difference only by making a difference to *other* features that are directly morally relevant.⁹ Thus, Direct Relevance says that *any given feature* could, in some context, make a moral difference—and not simply by making some *other* sort of difference.

Relevance and Direct Relevance connect to moral particularism by way of the idea of a *moral reason*. Particularists often present their view as concerning moral reasons, where a moral reason is a feature that accounts for *why* an action is morally wrong (right, permissible) or a feature that *morally counts* against (in favor of) an action.¹⁰ Moral reasons, then, are morally relevant features. Thus, as we'll see later, particularists end up committing themselves to

8 This understanding of moral relevance makes sense of McKeever and Ridge's flattening objection, and so I think this rough characterization is sufficient for our purposes.

9 The reader might disagree with me about which features we should count as directly and indirectly morally relevant in the case described, but this isn't a question I need to settle. (After all, utilitarians and deontologists are going to disagree about this point.) What's important for now is simply that the reader understands the general idea behind the distinction.

10 These two types of moral reasons parallel general and contributory moral principles, which I discuss in the following section.

theses such as Relevance and Direct Relevance by making certain kinds of claims about moral reasons.

In the following section, I'll describe a version of particularism that *avoids* commitments to Relevance and Direct Relevance, and thus show that one does not need to accept either of the above commitments in order to be a particularist.

2 Minimal Particularism

In this section, I will describe the main tenets of a minimal version of moral particularism. I will then explain why the view I describe can properly be called 'moral particularism' and will briefly explain why it entails neither Relevance nor Direct Relevance.

2.1 *The Tenets of Minimal Particularism*

Minimal particularism is a view about the existence of true moral principles—it says that there can be no true moral principles (of a specific sort) because any such principle would have to be infinitely long. Thus, the view can be made clear only if we have a reasonably clear idea of what a moral principle is. I will be working with the following two sorts of moral principles:

General moral principle: a finite generalization stating that some descriptive feature F (or finite set of descriptive features) is sufficient for an action being wrong (or right). E.g., 'If an action is an intentional utterance of a falsehood, then that action is wrong'.

Contributory moral principle: a finite generalization stating that some descriptive feature F always counts in favor of (or against) an action—a generalization stating that F has a consistent moral valence. E.g., 'Being the intentional utterance of a falsehood always counts against an action'.¹¹

Minimal Particularism has five commitments, some of which are entailed by others. The five commitments are:

Valence Incompactness: for any descriptive feature F, any descriptive sufficient condition for F's having a particular moral valence in a context must specify infinitely many facts.

¹¹ See Dancy, *Ethics Without Principles*, p. 76, for a description of what I've called a contributory moral principle.

Verdict Incompactness: for any moral property M, any descriptive sufficient condition for the instantiation of M will include infinitely many descriptive facts.¹²

Denial of Contributory Moral Principles: there exist no true contributory moral principles.

Denial of General Moral Principles: there exist no true general moral principles.

Holism about Moral Reasons: a feature that is a moral reason in one case may be no reason at all, or an opposite moral reason, in another.¹³

The relationship between the five tenets is this: Valence Incompactness entails the Denial of Contributory Moral Principles, Verdict Incompactness entails the Denial of General Moral Principles, and the Denial of Contributory Moral Principles entails Holism about Moral Reasons. Together, these commitments entail that there are no true moral principles of an important kind: those that 'link' descriptive and moral properties. In the remainder of this section, I will clarify these commitments and their relationships. Note that I will not be arguing in favor of these commitments—I simply intend to present them as clearly as possible so as to focus our attention on a specific version of particularism.

A feature has a moral valence relative to some action in a context just in case in that context it counts morally for or against that action; if it counts against the action it has a *negative* moral valence, and if it counts in favor of the action it has a *positive* moral valence. Thus, Valence Incompactness says that in order to guarantee that a descriptive feature has a specific moral valence in a context, one will need to specify infinitely many contextual features.

Valence Incompactness entails the Denial of Contributory Moral Principles, because a contributory moral principle specifies a moral valence had by a descriptive feature regardless of context. If there are some contributory moral principles, then Valence Incompactness is false. That is, if there is a descriptive feature that has the same moral valence regardless of context, then it is

12 For the sake of brevity, I put these two views in terms of 'incompactness' because the term expresses the idea that though there is an infinitely complex descriptive sufficient condition for the instantiation of a moral property (or for a feature's having a particular moral valence), there is no finitely complex descriptive sufficient condition.

13 See Dancy, *Ethics Without Principles*, p. 73, for his version of the holism thesis.

possible to provide a finite descriptive sufficient condition for that feature's having a specific moral valence. Thus, Valence Incompactness entails the Denial of Contributory Moral Principles.

Similarly, Verdict Incompactness entails the Denial of General Moral Principles, because if there *are* true general moral principles, then Verdict Incompactness is false. A general moral principle would provide exactly the sort of finite descriptive sufficient condition for the instantiation of a moral property that Verdict Incompactness says is impossible. Thus, Verdict Incompactness entails the Denial of General Moral Principles.

Finally, the Denial of Contributory Moral Principles entails Holism about Moral Reasons. Consider what happens when one rejects holism about moral reasons. If one rejects holism, then one is immediately committed to a weak version of *atomism* about moral reasons, a version of atomism according to which at least one descriptive feature retains the same moral valence across contexts. But to affirm that a descriptive feature retains the same moral valence across contexts *just is* to affirm the truth of at least one contributory moral principle. For example, let's say that you think the feature of *being the intentional utterance of a falsehood* always has the same (presumably negative) moral valence. In that case, you are committed to the truth of a contributory moral principle: '*Being the intentional utterance of a falsehood* always counts against an action'. Thus, the Denial of Contributory Moral Principles entails Holism about Moral Reasons.

What we have here, then, is a network of related commitments concerning moral generalizations. To summarize the view expressed by the network, we might say that there are no true descriptive-to-moral bridge principles (principles 'linking' descriptive properties to moral properties), because (a) principles must be finite and (b) any strictly true generalization about the relationship between descriptive and moral properties must be infinitely complex. But why call this network 'moral particularism'?

2.2 *Is 'Minimal Particularism' Particularism?*

Many readers will worry that the minimal particularism I have described does not get at the heart of what moral particularism is supposed to be about. To understand this worry, consider two issues on which minimal particularism remains silent. First, minimal particularism, in and of itself, says nothing about the nature of proper moral deliberation. For all that minimal particularism says, it could be that we should use moral principles—albeit false ones—in the course of moral deliberation. Second, minimal particularism says nothing about the nature of moral competence. For instance, it is compatible with minimal particularism that true moral principles (of either the general or

contributory sort) are necessary for moral competence. But, of course, many particularists have quite a bit to say about the nature of proper moral deliberation and the nature of moral competence. So perhaps minimal particularism has missed the mark by focusing solely on moral metaphysics.

However, there are three main reasons why we should think of minimal particularism as a version of moral particularism. First, it is a view (or closely resembles a view) that is held only by paradigmatic particularists. Second, it is the *kind* of view to which McKeever and Ridge's flattening objection was intended to apply. And third, it issues the sort of challenge to 'generalism' that particularism is supposed to issue—it entails that normative theory, insofar as normative theorizing is driven by certain generalist assumptions, is mistaken. Let's consider each of these points in more detail.

First, to see that minimal particularism is held by prominent particularists, consider the particularist views of Margaret Little and Jonathan Dancy. Little describes her particularism as a form of pessimism about the existence of moral principles; she says that, although the existence of true moral principles is not impossible, it would be 'philosophically serendipitous'.¹⁴ Particularists, according to Little, think there are most likely no true moral principles—no 'generalizations that are both accurate and contentful enough to be action-guiding'.¹⁵ Moreover, Little does not simply think that a true moral principle would likely be *merely* very complicated; instead, she thinks that a true moral principle would likely be irreducibly complex because of the very nature of the moral domain.¹⁶ Thus, Little denies the existence of descriptive-to-moral bridge principles, since she denies the existence of finite, action-guiding principles.

Dancy also appears to endorse minimal particularism, or else a very similar view in moral metaphysics. In his most recent descriptions of the view, he writes that in addition to believing that moral principles need not play a role in good moral deliberation, particularists also believe that there are most likely no true moral principles:

Particularists take their holism to be a reason to reject any invariance of reasons of either sort—whether at the overall or at the contributory level.... What the particularist says... is that the possibility of morality in no way depends upon a suitable provision of invariant reasons of the

14 Margaret Little, 'Moral Generalities Revisited', in B. Hooker and M. Little (eds.), *Moral Particularism* (Oxford: Oxford University Press, 2000), p. 277.

15 Little, 'Moral Generalities Revisited', p. 278.

16 Little, 'Moral Generalities Revisited', p. 279.

sorts that principles are attempting to specify. Principle-based accounts of morality... are left looking rather peculiar.¹⁷

Moreover, there are many passages in *Ethics Without Principles* in which Dancy describes particularism in language resembling minimal particularism. For example, in Dancy's discussion of the right-making relation, he expresses that he is concerned not only with whether moral principles are necessary for moral judgment, but also with whether there exist generalizations that capture *how moral properties work*.¹⁸ Thus, although Dancy clearly endorses a broader version of particularism than the version I just presented, it is clear that at least one element of his particularism is a thesis in moral metaphysics closely resembling minimal particularism.

Second, minimal particularism is the sort of particularist view—a view in moral metaphysics about how moral properties work—McKeever and Ridge intend to attack with their flattening objection. Notice that the flattening objection is a non-starter when it comes to particularist views about deliberation and moral competence. This is because (as we will see shortly) McKeever and Ridge's worry is, in and of itself, unrelated to claims about what is necessary for moral competence or for good moral deliberation.

Finally, minimal particularism has the capacity to issue the sorts of challenges to generalism that are characteristic of particularism. The reason moral particularism is theoretically significant is that, if it is true, much of traditional moral theorizing will turn out to have been misguided. Minimal particularism has the same consequence, and thus is of the same significance. If minimal particularism is true, then many versions of deontology and consequentialism will turn out to be false (namely, all of those versions that endorse the existence of finite descriptive-to-moral bridge principles, of either the contributory or general sort). Moreover, if it is impossible to formulate true descriptive-to-moral bridge principles, then *any* project that assumes there are finite, purely descriptive sufficient conditions for moral properties is doomed to failure. Finally, although minimal particularism is, in and of itself, silent on the nature of moral competence and proper moral deliberation, minimal particularism will have consequences for competence and deliberation when combined with other assumptions. For instance, if a minimal particularist

17 Jonathan Dancy, 'Moral Particularism', *Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), < <http://plato.stanford.edu/archives/fall2013/entries/moral-particularism/> >. Dancy's 'overall' principles are the same as my 'general' principles.

18 Dancy, *Ethics Without Principles*, p. 80.

assumes that moral competence is possible, then they must think that the existence of true moral bridge principles are not necessary for moral competence; if a minimal particularist thinks good moral deliberation is possible, then they must think that good moral deliberation need not involve the use of true moral principles.

2.3 *Minimal Particularism and the Two Additional Commitments*

I would like to flag the fact that minimal particularism is *not* committed to the two side-constraints described in 1.2, Relevance and Direct Relevance. As I will show in sections 4 and 5, this means that minimal particularism is not susceptible to McKeever and Ridge's flattening objection.

There are two reasons one might think that minimal particularists *are* committed to Relevance. First, one might think that a commitment to holism is sufficient for a commitment to Relevance; after all, it looks as if, according to holism, contextual features are always changing their moral valences. Second, one might think that a commitment to holism *alongside* a commitment to the non-existence of moral principles is sufficient for a commitment to Relevance; even if holism does not by itself entail Relevance, perhaps it does when combined with the view that the way in which valences change cannot be captured by any strictly true generalizations.

But holism does not entail Relevance. Holism simply says that any feature *that has a valence in one context* can have a different sort of valence in another context—but that claim is compatible with some features never having any moral valence at all. Holism essentially tells us to look at the set of features that have a moral valence in some context, and then tells us that those features can change their valences from context to context—it *does not* say that every feature is in that set. Since a feature must have some moral valence in order to be morally relevant, and since holism does not entail that every feature can have a moral valence, holism does not entail Relevance.

Moreover, holism conjoined with the thesis that there are no true moral principles does not entail Relevance. Recall that, according to minimal particularism, there are no general moral principles because of Verdict Incompactness, and there are no contributory moral principles because of Valence Incompactness. One might think that holism combined with Verdict Incompactness and Valence Incompactness entails Relevance, because they entail that every time a feature has a moral valence (or we have a moral verdict about an action) there are infinitely many morally relevant features of the context that are 'fixing' that valence or verdict; and one might think that if there are *infinitely many* morally relevant features in those contexts, then every feature is potentially morally relevant. But that is not so; even if one must specify infinitely

many contextual features in order to fix a moral verdict (or a feature's moral valence, as Valence Incompactness would have it), it still may be that *not all* contextual features will appear in that specification—the fact that a list is infinitely long does not entail that everything appears on that list.

Thus, minimal particularism does not entail Relevance. And this fact is sufficient to show that minimal particularism does not entail Direct Relevance (since Direct Relevance entails Relevance). In the following section, I will describe McKeever and Ridge's flattening objection in more detail; then in sections 4 and 5, I will show exactly how minimal particularism—because it avoids both Relevance and Direct Relevance—is immune to the objection.

3 The 'Flattening' Objection

McKeever and Ridge (MR), in *Principled Ethics*, argue that moral particularism 'flattens the moral landscape', that is, that particularism treats moral reasons of different kinds as if they were reasons of the same kind. MR develop their worry in the following way. Particularists must accept holism, as defined above. However, 'the combination of holism and particularism... suggests that an enormous range of considerations can be reasons for action'.¹⁹ MR's worry is that if one accepts holism together with the even stronger claim that there are no general or contributory moral principles,²⁰ then one is committed to thinking that all sorts of features can act as moral reasons in certain contexts.

MR's case is strengthened by the fact that both Little and Dancy have made assertions to the effect that, given a holism-motivated form of particularism, all sorts of surprising features are potential moral reasons. For instance, Little writes,

Depending on which case the comparison is made to, any feature may assume moral significance, from shoelace colour to the day of the week; after all, against a rich enough story, there are cases in which a change from Tuesday to Wednesday makes all the difference.²¹

19 McKeever and Ridge, *Principled Ethics*, p. 46.

20 Note that the version of particularism McKeever and Ridge may intend to focus on is anti-transcendental particularism, according to which moral thought doesn't rely on a suitable supply of moral principles. However, insofar as this sort of worry applies to that sort of particularism, it applies to the kind of particularism I'm concerned with in this paper.

21 Little, 'Moral Generalities Revisited', p. 291; quoted in McKeever and Ridge, *Principled Ethics*, p. 46.

And Dancy writes,

[W]e can give no sense to the idea that we might now have finished the list of moral principles or properties that make a difference sometimes... there is no limit to the number of properties which can on occasion be important.²²

So, MR take these particularist claims at face value and assume that particularists are committed to the view that every (or almost every) feature is possibly *morally relevant*, in the sense that any feature could—in the right context—make some sort of moral difference.²³

MR then note that many people who accept the existence of true moral principles—such as utilitarians and deontologists—can *also* accept the view that every feature is possibly morally relevant.

Of course, there is a sense in which a hardy generalist can agree with Little's comments [about the possible moral significance of any feature]. For example, a hedonistic act-utilitarian can admit that shoelace color or the day of the week can assume moral significance if the context is right, since shoelace color or day of the week can sometimes influence the consequences of various actions.²⁴

In this way, a utilitarian could admit that a feature such as shoelace color could make a moral difference, provided that one is in a strange context in which shoelace color will significantly affect the distribution of pleasures and pains. Since features such as shoelace color can figure into promises, even a deontologist could admit that a feature as trivial-seeming as shoelace color can make a moral difference.

MR infer, then, that particularists mean something different; after all, particularism is not very interesting if non-particularists can sign on to its commitments.

In order to distinguish particularism from these hardy generalist accounts of how such eccentric features can matter, Little presumably wants to

22 Jonathan Dancy, *Moral Reasons* (Oxford: Blackwell, 1993), p. 67; quoted in McKeever and Ridge, *Principled Ethics*, pp. 46–7.

23 McKeever and Ridge, *Principled Ethics*, p. 46.

24 McKeever and Ridge, *Principled Ethics*, p. 46.

insist that such seemingly eccentric features as shoelace color can in some sense have *direct* moral significance if the context is right.²⁵

MR point out that, in the case of the utilitarians and deontologists mentioned above, features such as shoelace color can make a moral difference only *indirectly*—for instance, in the case of the utilitarian, shoelace color can only make a moral difference by making a difference to the distribution of pleasures and pains (and only pleasures and pains are *directly* morally relevant to the rightness of actions). So, MR infer that the particularist must believe not merely that features such as shoelace color can make some moral difference (since non-particularists can agree with this claim), but that such features can make a *direct* moral difference; that is, MR infer that the particularist must think features such as shoelace color can affect the rightness of an action in the same way that a utilitarian thinks features such as pain and pleasure can affect the rightness of an action.

But this, MR argue, leads to the problem of ‘the flattening of the moral landscape’. Apparently the particularist thinks that, at least in some possible context, shoelace color can be a wrong-making feature in the very same way that pain is a wrong-making feature in many contexts. But this is obviously absurd, and many particularists *agree* that there is some important difference between these kinds of reasons.²⁶ Here, then, is my interpretation of MR’s argument:

1. Particularists accept holism.
2. Particularists accept the claim that there are no true general or contributory moral principles.
3. Therefore, particularists are committed to the claim that every feature is potentially morally relevant (is a potential moral reason). (Call this claim *Relevance*.) (1, 2)²⁷
4. Some generalists accept *Relevance*.
5. If particularists *and* some generalists accept *Relevance*, then particularists must believe something *stronger* than *Relevance*.

²⁵ McKeever and Ridge, *Principled Ethics*, p. 46.

²⁶ McKeever and Ridge, *Principled Ethics*, p. 47.

²⁷ McKeever and Ridge do not explicitly say that the combination of holism and particularism entails *Relevance* (they use the term ‘suggests’ rather than ‘entails’). However, I think that the foregoing illustrates that *Relevance* is often treated as a view that one naturally adopts when one adopts particularism. Everything I say in the next section is applicable to those who see *Relevance* as merely part of the ‘particularism package’.

6. Therefore, particularists must believe that every feature is potentially *directly* morally relevant (is a direct moral reason). (Call this claim *Direct Relevance*.) (3, 4, 5)
7. If particularists believe *Direct Relevance*, then particularism is unacceptable—it 'flattens the moral landscape', i.e., it treats moral reasons of different kinds as if they were reasons of the same kind.
8. Therefore, particularism is unacceptable. (6, 7)

After MR present this argument, they spend the remainder of their discussion surveying the different ways in which particularists have attempted to deal with this problem, and show that each strategy has failed in some respect. I will not try to defend the ways in which particularists have thus far dealt with the 'flattening' objection. Instead, I will show that minimal particularism can avoid the inference from (1) and (2) to (3), and that particularists should deny premise (5). Thus, minimal particularism isn't susceptible to the flattening objection.

4 Minimal Particularism Avoids the First Inference

So, first: minimal particularism can avoid the inference from (1) and (2) to (3).

MR—as well as Dancy and Little—sometimes write as if they think a commitment to holism entails Relevance. However, holism alone does not entail Relevance. There are many different ways of understanding holism in the theory of reasons; for now, though, let's work with the definition of holism supplied earlier in this paper, since that is the definition assumed by Dancy, Little, and MR. According to this definition, holism is true just in case every feature that is a moral reason in one context may have a different moral valence in other contexts. (Notice that to be a moral reason *just is* to be a feature with a positive or negative moral valence.) But no matter how we interpret the 'may' in the definition of holism,²⁸ it doesn't follow from holism that absolutely any feature can be a reason. Holism says that for any feature with a moral valence in a context, there is another possible context in which that feature has a different moral valence (or no moral valence at all); it *does not* say that for any feature, there's a possible context in which that feature has a moral valence (i.e., is a moral

²⁸ Crisp points out that there are strong and weak interpretations of this statement of holism. According to the weak version, a moral reason need not maintain the same moral valence; according to the strong reading, every moral reason has a variant moral valence (i.e., the possibility of invariance is ruled out). See Roger Crisp, 'Ethics Without Reasons', *Journal of Moral Philosophy* Vol. 4.1 (2007), pp. 40–49.

reason). For all that the definition of holism says, it could be that the set of features that can be moral reasons *excludes some features*. Thus, nothing about holism entails that every, or even almost every, feature can be a moral reason.

Moreover, holism does not entail Relevance *even when holism is combined with the claim that there are no true moral principles*. This is something minimal particularism can help us see. Assume that the reason there are no true moral principles is that spelling out any strictly true moral principle would require specifying an infinite number of morally relevant contextual features, as minimal particularism holds.²⁹ In this case, it could still be that features such as *shoelace color* are never moral reasons: the fact that a list of features is infinitely long does *not* entail that every feature appears on that list. Thus, even if particularists are committed to there being infinitely many morally relevant features in a given context, they are not immediately committed to the claim that every feature is potentially morally relevant.

However, this is not to say that particularists have *no* good reason for endorsing Relevance—all I've pointed out so far is that there is at least one interesting version of moral particularism that isn't *immediately committed* to Relevance. One reason particularists—even those who start off as minimal particularists—might want to adopt Relevance is to avoid a commitment to negative moral principles. If Relevance is false, that means that there are some features that are never morally relevant; but then a statement such as 'Shoelace color never makes a moral difference'³⁰ would be *true*. Perhaps even these descriptive-to-moral generalizations reek of generalist bias.³¹

Another reason particularists might want to adopt Relevance is that it is plausible. Relevance merely says that any feature *could* make a moral difference if it were placed in the right sort of context. And it turns out to be difficult to think of a feature that could not make a moral difference, given that one is allowed to take into account very unusual situations. Having brown hair, standing on one foot, Pluto's being a Kuiper belt object—we can all think of situations in which these features could make a moral difference. Thus, perhaps particularists could legitimately accept Relevance because of their inability to locate a counterexample to Relevance.

To summarize so far, holism does not entail Relevance, and minimal particularism illustrates that holism *combined with the view that there are no*

29 See McKeever and Ridge, *Principled Ethics*, pp. 41–43 for their discussion of the role that 'uncodifiability' plays in particularist views. I disagree with their assessment since I think it can be shown that arguments from 'radical holism' are not, in fact, question-begging.

30 The reader can replace *shoelace color* with whatever feature is a candidate for never being morally relevant.

31 I'd like to thank David Faraci for pointing this out to me.

moral principles does not entail Relevance. Thus, not all versions of particularism are susceptible to MR's flattening objection. Nevertheless, there may be other reasons why minimal particularists would want to take up Relevance, and so we might wonder: is minimal particularism supplemented by Relevance susceptible to MR's objection?

5 Minimal Particularists Are Not Committed to Direct Relevance

So far, I have argued that a commitment to holism, even when combined with the view that there are no true moral principles, does not entail a commitment to Relevance, the claim that every feature can be a moral reason (where to be a moral reason in a context is to have a moral valence in that context). But let's now assume that the minimal particularist *is*—for whatever reason—also committed to Relevance. MR note that Relevance is a view with which many non-particularists (utilitarians, deontologists) can agree. So, in premises 5 and 6 of the flattening objection, MR saddle the particularist with the even stronger view that every feature can be *directly* morally relevant (Direct Relevance), so as to assign a distinctive view to the particularist and thus enable the particularist to robustly disagree with the non-particularist. However, this move of MR's is misguided, because the particularist and the non-particularist need not disagree about *everything* in order to have a robust disagreement about *something*.

To clarify this idea—that two parties need not disagree about everything in order to disagree about something—let's work with the following dialectical principle:

Permissible Agreement: given two opposing views, *p* and *not-p*, if *q* is compatible both with *p* and with *not-p* then, all else being equal, people who accept *p* and people who accept *not-p* may agree about *q*.

If one accepts Permissible Agreement—and I think one should—then it turns out that particularists and non-particularists may agree that *some* but *not all* features can be directly morally relevant (that is, they may agree on a view that is inconsistent with Direct Relevance).

Let's call the view that some but not all features can be directly morally relevant *Limited Direct Relevance* (LDR). LDR is obviously compatible with non-particularist views; for instance, we saw earlier that utilitarians think that only features such as pain and pleasure can be directly morally relevant, and that all other features can be only *indirectly* morally relevant (in the sense that

all other features can make a moral difference only by affecting the distribution of pains and pleasures). To see that LDR is also compatible with particularism, notice that a minimal particularist could consistently adopt the view that only a certain set of features can be directly morally relevant, and that all other features (perhaps including *shoelace color*) are morally relevant only insofar as they affect the distribution of the former features. Thus, there is no good reason for MR to saddle the particularist with the very strong view expressed by Direct Relevance.

But can we be sure that particularists can consistently adopt LDR? One might have the following worry: if the particularist believes that some features can be directly morally relevant, does that not entail the existence of at least *some* contributory moral principles? For instance, if the feature *being an instance of dishonesty* is one of those distinctive features that can be directly morally relevant, does that not entail that the feature of *being an instance of dishonesty* counts against any action that has it?

The answer to these questions is 'no'; the fact that a feature is sometimes directly morally relevant to the rightness (or wrongness) of an action does not entail that that feature has the same moral valence across contexts, and thus does not entail the existence of a contributory moral principle. For instance, it is consistent with LDR that the feature *being an utterance of a falsehood* is often directly morally relevant and yet is not in *every* context directly morally relevant. Moreover, it is consistent with LDR that when *being an utterance of a falsehood* is directly morally relevant, it sometimes counts *in favor* of an action and at other times counts *against* an action. To summarize, the particularist can say that a feature is of direct moral relevance *in a context* without being committed to the stronger claim that that feature retains the same valence across contexts; and at the same time, the particularist can say that not all features are possibly directly morally relevant. So, the claim that only some features can be directly morally relevant is compatible with minimal particularism, and so LDR is compatible with minimal particularism.

Thus, because LDR is compatible with both particularism and the negation of particularism, particularists and non-particularists need not disagree about LDR in order to have a robust disagreement. And this shows that MR's inference that the particularist must accept Direct Relevance is a bad inference.

6 Objections

Even if everything I've said so far is correct, one might have several reasons for worrying that the conclusion I've reached is problematic or uninteresting.

Is a non-flattening version of particularism sufficiently particularistic? Will the minimal particularism I've defended be susceptible to other kinds of flattening? And even if particularism is not in and of itself susceptible to flattening, could it be that the standard *motivations* for particularism commit particularists to flattening?³²

6.1 *Is this Particularism?*

Earlier in this paper, I argued that we should consider minimal particularism to be a version of particularism because it (a) is endorsed only by paradigmatic particularists, (b) is the sort of view to which MR's objection is intended to apply, and (c) offers the kind of theoretical challenges to generalism that particularism is supposed to offer. But one might nonetheless worry that a minimal particularism not committed to Direct Relevance is not very particularist.

Notice that it is not enough for the minimal particularist to deny Direct Relevance in order to avoid flattening; even if Direct Relevance is false, it could still be true that *some* properties such as *shoelace color* are sometimes directly morally relevant, and thus 'on a par' with properties such as *causing pain*—in other words, the minimal particularist who rejects Direct Relevance cannot immediately rule out the kinds of unacceptable parities that motivate the flattening objection. So, in order for a minimal particularist to avoid flattening, she must say that there are restrictions on which properties can be directly morally relevant. Perhaps the imposition of such restrictions is not particularist—after all, it seems as if '*Shoelace color* is never directly morally relevant' is a contentful moral generalization.

In response to this worry, I'd first like to flag that minimal particularism paired with such restrictions can remain particularistic in other respects. For instance, minimal particularism is compatible with the view that properties such as *causing pleasure* and *causing pain*—properties that are sometimes of direct moral relevance, but that usually have opposite moral valences—are nevertheless on a par. (I elaborate on this point in 6.2.)

Second, we should notice that minimal particularism, even with these restrictions in place, meets the three conditions mentioned at the beginning of the paper: it (a) is endorsed only by paradigmatic particularists, (b) is the sort of view to which MR's objection is supposed to apply, and (c) offers the sorts of challenges to generalism particularists claim to offer. Thus, if one thinks that meeting these three conditions is sufficient to make a view 'particularist', then minimal particularism is particularist.

³² I'd like to thank two anonymous referees for bringing these three objections to my attention.

But, most importantly, I'd like to emphasize that the goal of this paper is to evaluate (in at least one important respect) a *weak* particularist view that is also theoretically *interesting*. The particularism I'm describing is not a logically strong version of particularism; but so long as it's theoretically interesting (in the way that particularist views are supposed to be), then we should study it. In fact, looking at the logically weakest interesting version of a view allows one to see *exactly which view* is generating interesting theoretical consequences, and also allows one to see who can and cannot consistently adopt the view.

6.2 *Another Sort of Flattening?*

One might worry that while minimal particularism avoids one type of flattening, it nevertheless faces another, and thus doesn't obviate the need for a theory of default reasons.

Consider two properties that are often directly morally relevant, but that typically have different moral valences, such as *causing pain* and *causing pleasure*. Most of us think that there is an important difference between these two sorts of properties. But according to minimal particularism, there's no *deep* difference at all—it's simply true of each of them that their moral valences depend on context. Thus, although a denial of Direct Relevance avoids one sort of flattening—namely, the erasure of a distinction between (a) those properties that intuitively seem potentially directly morally relevant and (b) those properties that intuitively seem like they could never be directly morally relevant—it doesn't avoid erasing an important distinction between directly morally relevant properties that typically have different moral valences.

However, this worry does not pose a problem for minimal particularism, but rather captures one of the theoretically interesting consequences of the view. It's a basic commitment of particularism that generalizations such as '*causing pleasure* counts in favor of an action' and '*causing pain* counts against an action' are strictly speaking false, and thus that both properties can take on a variety of valences. It's not a compelling criticism of minimal particularism to claim that a basic particularist commitment amounts to a flattening of the moral landscape—that would be to criticize particularism by restating particularism.

Moreover, I don't think that this objection shows that minimal particularism is *so* revisionary as to be unacceptable in light of our strongly held moral intuitions. It would be surprising if we had to say that there is little difference between *causing pleasure* and *causing pain*—but the minimal particularist doesn't have to say that. The minimal particularist merely has to say that there's no deep, *metaphysical difference in kind* between the two properties; but this is

consistent with saying that these two properties are significantly different because of the different roles that they play in the course of actual human lives. Maggie Little expresses a similar point (about the proper understanding of moral presumptions):

The judgement that a given principle such as 'lying is wrong' will help rather than mislead a moral novice reflects a judgement about the sorts of contexts she is likely to encounter, just as our agreement that it is better training to tell beginner drivers 'Never slam on the brakes' instead of 'Stomp on the gas whenever you see another car' reflects a judgement that the student will most likely be facing our world of crowded highways and not the post-apocalyptic world of Mad Max movies.³³

Of course, Little thinks that particularists do need a theory of defeasible generalizations to respond to the flattening objection. However, I think we can use her point to illustrate that particularists—including minimal particularists—need not adopt a theory of default reasons (or defeasible generalizations) to account for the datum that properties such as *causing pain* and *causing pleasure* seem different. One need not identify a deep, metaphysical distinction (in terms of moral valences) in order to say that two properties are different in important ways; the minimal particularist can say that, as a matter of fact, they play different sorts of roles in our lives, and are importantly different for that reason.

6.3 *Flattening and Motivations for Particularism*

Earlier, I stated that I wouldn't be evaluating motivations for particularism (or minimal particularism, more specifically). But one might think this is problematic. Although minimal particularism may not in and of itself be committed to Relevance or Direct Relevance, it could be that the traditional *motivations* for adopting particularism (and thus candidate motivations for minimal particularism) commit one to those views. If it turns out that the good reasons for believing particularism involve commitments to Relevance and Direct Relevance, then it looks as if those views are, in fact, part of the particularism 'package' after all.

One might worry specifically about the 'non-monotonicity' motivation for particularism. Particularists often support their view by appeal to the

33 Little, 'Moral Generalities Revisited', p. 295.

non-monotonicity of good moral reasoning.³⁴ An inference is monotonic if no addition to its premises can make the inference less good,³⁵ and we can say that a domain of inferences is monotonic if all of the good inferences made in that domain are monotonic. Particularists often point out that the domain of moral inferences is not monotonic—one could properly issue a moral verdict *M* given the pieces of information *p*, *q*, and *r*, and then properly issue the moral verdict not-*M* given the additional piece of information *s* (even while still accepting *p*, *q*, and *r*). The non-monotonicity of good moral reasoning is often used as a point in favor of particularism, because it suggests that moral conclusions are *defeasible* in a way that makes strictly true, finite, descriptive-to-moral principles impossible to formulate. One might think that the non-monotonicity of good moral reasoning indicates that just about anything could, in the right context, be morally relevant.

To begin my response to this concern, notice first what the non-monotonicity of good moral reasoning really is. For the domain of moral reasoning to be non-monotonic is for that domain to *not* be monotonic. Given the definition of monotonicity above, the non-monotonicity of good moral reasoning just means that in the domain of moral reasoning, there is at least one inference and at least one premise such that the inference can be made worse by the addition of that premise. This is not a very strong claim, and it doesn't come close to supporting the defeasibility of *most* moral conclusions that would in turn support particularism. Thus, the non-monotonicity of good moral reasoning is not in and of itself an adequate source of support for particularism in the first place.

Nevertheless, let's say we have good reason to think that non-monotonicity is *widespread* in the domain of moral reasoning; that is, assume that *every* good moral inference is non-monotonic. This would mean that for every good moral inference, there is an additional piece of information that (if added to the premises) would make that inference worse. But even this very strong form of non-monotonicity does not entail Relevance (and thus does not entail Direct Relevance). It could be true that for every good moral inference, there is an additional piece of information that would make that inference worse, *and at the same time* it could be true that there are some pieces of information that don't make any inferences worse. To be a morally relevant property is (as discussed earlier) to potentially make a moral difference; thus, to be morally relevant, a property must either be referred to by a predicate in a

34 See, for instance, Dancy, *Ethics Without Principles*, pp. 8–9, 80. Note that Dancy doesn't think that the fact that good moral reasoning is non-monotonic entails particularism. This is because non-monotonicity is compatible with atomism about moral valences.

35 Dancy, *Ethics Without Principles*, 8.

necessary premise in a good moral inference, or else be referred to by a predicate in an additional piece of information that could make that inference worse. Since it's consistent with the widespread non-monotonicity of the moral domain that some properties play neither of these roles, it's consistent with the widespread non-monotonicity of the moral domain that some properties are never morally relevant. Thus, widespread non-monotonicity does not entail Relevance, and so does not entail Direct Relevance. My response to this worry is, essentially, simply another way of stating a point from earlier in this paper: just because the list of potentially morally relevant properties is infinitely long, it doesn't follow that every property appears on that list.

I conclude, then, that the motivation for particularism most likely to be linked to Relevance and Direct Relevance is, in fact, not so linked. Thus, the typical motivations for particularism do not make Relevance or Direct Relevance part of the particularism 'package'.

7 Conclusion

In summary, I have argued that there is a version of moral particularism—minimal particularism—that can accommodate McKeever and Ridge's worry about 'flattening' *without* committing itself to any *further* theory (such as a theory of default reasons or a theory of defeasible generalizations).³⁶ The reason minimal particularism is able to do this is that it is not immediately committed to Relevance or to Direct Relevance—two commitments that the flattening objection exploits.

What this shows is that the two standard ways of responding to McKeever and Ridge—accepting the objection and developing a further theory, or accusing McKeever and Ridge of begging the question—aren't the only possible particularist responses. Particularists can respond to this objection by committing themselves to less, all while remaining genuine moral particularists. I think that this, in turn, tells us a couple of things:

1. While we should continue our discussion of canonical versions of particularism, we should not let those versions dictate the terms of the

³⁶ One might worry that I have, inadvertently, committed myself to something like a theory of default reasons. But that is not correct. The particularists who endorse a theory of default reasons admit that any feature is potentially directly morally relevant, but account for the intuitive difference between features such as *pain* and *shoelace color* by appealing to a further theory (a theory of default reasons). I've shown that particularists need not say that every feature is potentially directly morally relevant in the first place.

particularism debate; there are many possible, as-of-yet unexplored versions of moral particularism, and some of those are theoretically interesting in their own right.

2. When considering objections to particularism, we should not *simply* ask whether any of the canonical versions of the view can respond to them; we must also consider what the *pros hen* of particularism is, and whether any plausible version of particularism can respond to those objections.

If we take these two lessons into account in discussions of moral particularism, we will be more likely to accurately assess the viability of particularist views, and thus more likely to see what it is moral particularism has to offer when it comes to moral theorizing.³⁷

References

- Roger Crisp, 'Ethics Without Reasons?' *Journal of Moral Philosophy*, 4.1 (2007), pp. 40–49.
- Jonathan Dancy, 'Defending the Right', *Journal of Moral Philosophy*, 4.1 (2007), pp. 85–98.
- , *Ethics Without Principles* (Oxford: Oxford University Press, 2005).
- , 'Moral Particularism', *Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), < <http://plato.stanford.edu/archives/fall2013/entries/moral-particularism/> > .
- , *Moral Reasons* (Oxford: Blackwell, 1993).
- Mark Lance and Margaret Little, 'Defeasibility and the Normative Grasp of Context', *Erkenntnis* 61 (2004), pp. 435–55.
- Margaret Little, 'Moral Generalities Revisited', in B. Hooker and M. Little (eds.), *Moral Particularism* (Oxford: Oxford University Press, 2000), pp. 276–304.
- Sean McKeever and Michael Ridge, *Principled Ethics: Generalism as a Regulative Ideal* (New York: Oxford Univ. Press, 2006).
- , 'Turning on Default Reasons', *Journal of Moral Philosophy*, 4.1 (2007), pp. 55–76.
- Alan Thomas, 'Practical Reasoning and Normative Relevance: A Reply to McKeever and Ridge', *Journal of Moral Philosophy*, 4.1 (2007), pp. 77–84.

37 I'd like to thank Michael DePaul, Terence Cuneo, Amy Lara, Graham Leach-Krouse, and Susan Sterrett for their very helpful comments on earlier drafts of this paper. I'm also grateful to audiences at the 2014 meeting of the Kansas Philosophical Society and the 2014 meeting of the North Carolina Philosophical Society for their feedback.