THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Turning on Default Reasons

Citation for published version:

McKeever, S & Ridge, M 2007, 'Turning on Default Reasons', *Journal of Moral Philosophy*, vol. 4, pp. 55-76(21). <<http://booksandjournals.brillonline.com/content/10.1177/1740468106072786>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Journal of Moral Philosophy

Publisher Rights Statement:

© McKeever, Sean, Ridge, Michael. (2007). Turning on Default Reasons. *Journal of Moral Philosophy*, 4(1), 55-76. 10.1177/1740468106072786

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Turning on Default Reasons*

SEAN MCKEEVER

Dept of Philosophy
Davidson College
Davidson, NC, USA
semckeever@davidson.edu

MICHAEL RIDGE

Dept of Philosophy
Edinburgh University
Edinburgh, Scotland
mrIDGE@staffmail.ed.ac.uk

Particularism takes an extremely ecumenical view of what considerations might count as reasons and thereby threatens to 'flatten the moral landscape' by making it seem that there is no deep difference between, for example, pain, and shoelace color. After all, particularists have claimed, either could provide a reason provided a suitable moral context. To avoid this result, some particularists draw a distinction between default and non-default reasons. The present paper argues that all but the most deflationary ways of drawing this distinction are either implausible or else insufficient to help the particularist avoid flattening the moral landscape. The difficulty can be avoided, however, if we reject particularism's extremely ecumenical view of reasons.

Keywords: default reason; defeasible generalizations; Jonathan Dancy; particularism; primary reason; secondary reason

One leading version of moral particularism holds that moral judgement does not presuppose the possibility of providing a 'suitable supply of moral principles'.¹ Particularists argue for their view primarily on the strength of a doctrine they call 'holism in the theory of reasons'. Holism in the theory of reasons holds that what is a reason in one context can be no reason at all,

* Many thanks to the participants of the conference on particularism hosted by the University of Kent in December 2004. Special thanks to Jonathan Dancy and Alan Thomas for helpful comments. Special thanks also to Christian Piller for some very helpful comments about an earlier version of our discussion of enablers and defeaters.

Some of the material presented here draws heavily on Chapter 3 of our book *Principled Ethics: Generalism as a Regulative Ideal* (Oxford: Oxford University Press, 2006).

1. J. Dancy, *Ethics Without Principles* (Oxford: Oxford University Press, 2004), p. 7.

or a reason with the opposite valence, in another context. We have argued elsewhere that this argument is unsound or question-begging, but we shall not repeat the main lines of that argument here.² Instead, we want to explore a concern which arises out of the radical sort of holism to which particularists are committed. For the combination of holism and particularism, in turn, suggests that an enormous range of considerations can be reasons for action. Here it is worth pausing to remember some of the more memorable examples particularists have offered as considerations which can have moral significance in the right context. Here is Margaret Little:

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 the change from Tuesday to Wednesday makes all the difference.³

Of course, there is a sense in which a hardy generalist can agree with Little's comments. For example, a hedonistic act utilitarian can admit that shoelace colour or the day of the week can assume moral significance if the context is right, since shoelace colour or day of the week can sometimes influence the consequences of various actions. Insofar as shoelace colour or day of the week can figure in the content of a promise, a Kantian can also admit that such features can assume moral significance. In order to distinguish particularism from these hardy generalist accounts of how such eccentric features can matter, Little presumably wants to insist that such seemingly eccentric features as shoelace colour can in some sense have *direct* moral significance if the context is right. In our view, the most plausible gloss on this idea is that such features can figure in the content of moral reasons for action. Jonathan Dancy goes so far as to insist that, 'we can give no sense to the idea that we might now have finished the list of moral principles or of properties that can make a difference sometimes...there is no limit to the number of properties which can on occasion be important'.⁴ Once again, it must be taken as read that Dancy does not intend this point in the anodyne sense in which a utilitarian or Kantian could accept it. Nor should it be surprising that particularists make such strong claims about what sorts of considerations can be moral reasons for action. For if the number of possible reasons were finite and manageable then the idea that we might codify the moral landscape would look much more plausible.

One danger that has not been lost on particularists is that their extremely ecumenical view of moral reasons for action threatens to flatten the normative landscape in an implausible way. After all, even if we think that in the right context shoelace colour can provide a reason for action, there surely is

2. See S. McKeever and M. Ridge, 'What Does Holism Have to Do With Particularism?', *Ratio* (2005), pp. 93-103.

3. M. Little, 'Moral Generalities Revisited', in B. Hooker and M. Little, *Moral Particularism* (Oxford: Oxford University Press, 2000), p. 291.

4. J. Dancy, *Moral Reasons* (Oxford: Blackwell, 1993), p. 67.

an important difference between considerations of shoelace colour and considerations of pain, pleasure, promising and the like. Mark Lance and Margaret Little put the worry very clearly:

In unqualified form, moral holism of the sort just outlined seems to imply that lying, killing and the infliction of pain have no more intimate connection to wrongness than do truth-telling, healing and the giving of pleasure. After all, each, in the right context, can have a positive, negative or neutral moral import. But the morally wise person, one might have thought, is someone who understands that there is a deep difference in moral status between infliction of pain and shoelace colour, even if both can, against the right narrative, be bad-making.⁵

Particularists have tried to show how they can avoid any unfortunate flattening of the normative landscape by drawing a distinction between default reasons and non-default reasons. Roughly, the idea is to privilege certain considerations as default reasons, where the status of such considerations as a reason with a given valence does not stand in need of explanation. By contrast, non-default reasons are sometimes reasons but their status as reasons cannot be assumed and always stands in need of explanation. Drawing some such distinction does seem essential to making particularism plausible, and in this article we explore different ways of understanding the distinction between default reasons and non-default reasons. The distinction can be given metaphysical, epistemological and pragmatic readings. Read metaphysically, we argue that the distinction is untenable (section I). Read epistemologically, we argue that the distinction cannot do the work particularists need it to do (section II). Read pragmatically, we argue that the distinction is tenable and *may* be adequate for the particularist's purposes though this is certainly not the reading of the distinction particularists themselves have favoured (section III). Whether this reading is nonetheless sufficient for their purposes depends on whether we should think the normative landscape is necessarily not flattened or whether we should be confident only that our local normative landscape can reasonably be understood as not flattened. Our own view is that the best way to avoid the flattening of the normative landscape is to reject the particularist's highly ecumenical conception of reasons for action and more generally of what kinds of features can assume direct moral significance.

I. The Metaphysical Reading

To convey an intuitive feel for the basic idea behind the distinction between default reasons and non-default reasons, we can do no better than to quote Dancy's own admittedly metaphorical way of putting it. Some reasons, Dancy claims, 'arrive switched on, though they may be switched off if the circumstances so conspire, while others arrive switched off but are switched on by

5. M. Lance and M. Little, 'Particularism and Anti-Theory', in D. Copp (ed.), *The Oxford Handbook of Ethical Theory* (Oxford: Oxford University Press, 2006), pp. 20-21.

appropriate contexts'.⁶ Making the same distinction in the theory of value, Dancy remarks that, 'some features come switched on already, as it were, though they can be switched off by other features; others do not come switched on, but they can be switched on by a suitable context'.⁷ A 'default reason', then, is a consideration whose reason-giving force has a default setting. Pain, for example, might be in this sense a default reason. It 'arrives' as a reason against performing an action that would promote it, but (in the spirit of holism) can be switched off by context. Though intuitive, the language of arriving and switching suggests that some considerations have the force of reasons independently of context. This suggestion is at odds with at least some particularists who emphasize that considerations can have reason-giving force only in a surrounding context.⁸ Fortunately, the familiar particularist tools of defeaters and enablers can help unpack the metaphor. A default reason is a consideration whose status as a reason does not stand in need of explanation by some further feature of the situation—it does not require an enabler. Non-default reasons do stand in need of such explanation and require an enabler.⁹ One interesting corollary here seems to be that one and the same consideration could be a default reason with respect to one valence and a non-default reason with respect to the opposite valence. The example of pain is a good illustration if we agree with those particularists who have argued that the fact that an action would cause pain is sometimes a reason in its favour. When the fact that an action causes pain is a reason against the action it is plausibly a default reason, for the idea that pain counts against an action stands in no obvious need of explanation. By contrast, when an action causing pain is a reason in favour of the action, its having this valence will require some enabler, e.g. that the pain is constitutive of athletic achievement.¹⁰

Dancy understands the distinction between default and non-default reasons as a metaphysical one. The idea is not simply that some considerations usually do have a particular reason-giving force or that we should expect or presume them to have a given status as reasons. Instead, the distinction concerns the considerations themselves. Pain is a reason against (when it is) just because of the nature of pain. To see things this way, then, is not just to know the

6. Dancy, *Ethics Without Principles*, p. 113.

7. J. Dancy, 'Are There Organic Unities?' *Ethics* 113 (2003), pp. 629-50 (638).

8. Little ('Moral Generalities Revisited') emphasizes this point and understands it as one of the central lessons of holism. To think that a consideration could have a reason-giving force antecedent to its position in a surrounding context is, she claims, to subscribe to a dubious 'ideal model' of moral reasons.

9. G. Cullity, 'Particularism and Presumptive Reasons', *Proceedings of the Aristotelian Society*, Supp. edn (2002), pp. 169-90, marks the distinction in a similar way. What he calls 'presumptive reasons' are 'pro tanto reasons unless undermined'. Though Cullity does not explicitly define non-presumptive reasons, one might think that they are considerations that must first be 'enabled'.

10. The example is from Lance and Little, 'Particularism and Anti-Theory'.

contingent fact that the fact that an action would cause pain is a reason against the action, but rather to know something about pain. By contrast, when the fact that an action would cause pain is a reason in favour of the action there must be an enabler which explains why this is so. This way of drawing the distinction is extremely natural for a particularist, as it builds heavily on the machinery of enablers and defeaters emphasized in their discussions of holism. Moreover, this approach seems to avoid flattening the normative landscape in a fairly robust way, as the distinction between default reasons and non-default reasons on this interpretation is meant to be metaphysically deep and not contingent on local circumstances.

However, the metaphysical reading of the distinction between default reasons and non-default reasons is untenable. For on any plausible version of holism *all* reasons require enablers, in which case *all* reasons are non-default reasons in this sense, in which case the normative landscape remains flattened. Why do we say all reasons require enablers? Here we invoke the plausible idea that the presence of certain considerations are always necessary for any consideration to count as a reason but which are not themselves part of the reasons which they make possible. The best example of this is the fact that no fact can be a reason for an agent to perform an action if the agent cannot perform the action. Nonetheless, as particularists have plausibly insisted, the fact that an agent could do something is not part of the reason to do it. The reason to meet me for lunch is that you promised, not that you promised and can. If the fact that an agent could perform an action is not a reason or a part of a reason, though, then what is it? The tempting answer is to say it is an enabler for all possible reasons, and indeed we think that this answer is correct. So all reasons need at least one enabler, in which case there can be no default reasons in Dancy's canonical sense and the distinction lacks interest.

The preceding objection demonstrates the untenability of Dancy's canonical formulation of the distinction between default reasons and non-default reasons. However, a very close cousin of Dancy's formulation can survive this first objection. For we can and should distinguish what might be called *global enablers* from *local enablers* and redraw the metaphysical version of the distinction between default reasons and non-default reasons in terms of local enablers rather than global enablers.¹¹ A global enabler is a consideration which must obtain before there can be any reason for an agent to perform an action. That the agent can perform the action is a clear example of a global enabler. More controversially, perhaps there can be a reason for someone to do something only if the agent can refrain from doing it as well. The idea would be that reasons function to guide our actions but guidance is possible only when an agent has some choice between alternatives. If an agent cannot help but perform an action then talk of guidance and reasons is simply out of

11. Dancy prefers the terminology of 'generic' and 'specific'.

place.¹² Though just which considerations are global enablers will be a controversial question (and one we will not try to settle), the idea of a global enabler is tied to the general idea of a reason for action. By contrast, some enablers (most of the ones particularists discuss) are local, having to do with specific kinds of reasons rather than with the very idea of a reason for acting as such. For example, suppose that while the fact that an action will involve pain normally is a reason against it, if the pain is constitutive of athletic accomplishment then that very same consideration can be a reason in favour of the action.¹³ In that case, the fact that a given pain would be constitutive of athletic accomplishment can function as a local enabler.

This distinction between global and local enablers shows the way to a more subtle distinction between default reasons and non-default reasons which is very much in the spirit of Dancy's canonical version of the distinction but which avoids the objection from global enablers. We might give the idea of a default reason the following intuitive gloss: a default reason just is a fact which does not depend for its status as a reason on any further facts beyond those required by the very concept of a reason for acting. The existence of 'global enablers' then poses no threat to the proffered distinction. For a default reason is on this account not a fact which stands in need of no enablers whatsoever to function as a reason, but rather is a fact which requires no local enablers to function as a reason.

However, even on this more refined understanding, the metaphysical reading of the distinction between default reasons and non-default reasons ultimately collapses. The tenability of the distinction depends on the tenability of the distinction between the presence of an enabler and the absence of a defeater but this distinction itself is untenable. Suppose that we agree that the fact that an action will promote pleasure is sometimes a reason but also maintain that its status as a reason depends (in part) on whether the pleasure would be sadistic. We can then describe our view in one of two ways. On the one hand, we could maintain that the fact that the pleasure promoted would be sadistic (when this is a fact) is a defeater, preventing the fact that the action would promote pleasure from functioning as a reason. Alternatively, we could maintain that the fact that the action is sadistic merely serves to indicate the absence of an enabling condition, understanding the fact that the pleasure promoted is non-sadistic as an enabler which must be present for the fact that an action would promote pleasure to function as a reason. If we characterize the presence of sadism as defeater then we are still free to characterize the fact that an action would promote pleasure as a default reason which needs no enablers. If instead we characterize the presence of sadism as the absence of an enabler then we no longer can say that fact about pleasure is

12. Kant makes a similar point in the *Groundwork* when he claims that imperatives apply only to 'imperfect wills', wills capable of being determined by rational norms, but also subject to competing motivations (K 413-414).

13. The example is borrowed from Lance and Little, 'Particularism and Anti-Theory'.

a default reason in the intended sense. This naturally leads to a series of questions. How do we determine which of these descriptions (needs an enabler vs. needs the absence of a defeater) is correct? This epistemological question is difficult enough, but even more difficult questions concern the logical independence of the notion of an enabler and the notion of a defeater. Is there any reason in our example to think that both descriptions of our view (the one couched in terms of non-sadism as an enabler and the one couched in terms of sadism as a defeater) could not *both* be correct? Indeed, is there any reason to suppose that claims about enablers and claims about defeaters are not really just notational variants on one another?

In our view, the answer to this last question is 'no' and this spells the end for the metaphysical reading. If a consideration needs an enabler in order to function as a reason then trivially that consideration can be defeated; the absence of the enabler is a defeater. Similarly, if a consideration's status as a reason can be defeated then the consideration needs an enabler; the conjunction of the negation of all of its possible defeaters is an enabler. Lest this seem too swift, note that enablers and defeaters are typically defined in counterfactual terms. Here is Dancy:

...in the absence of (2), (1) would not have favoured the action. In this sense, the presence of (2) *enables* (1) to favour (5).¹⁴

Certainly if we understand enablers and defeaters in counterfactual terms then the fact that a consideration needs an enabler entails that it can be defeated and vice versa. Moreover, Dancy himself goes on explicitly to endorse this necessary connection:

...trivially, the absence of an enabler will disable what would otherwise be a reason.¹⁵

Though Dancy here endorses only the inference from 'needs an enabler' to 'can be defeated' there is no textual or rational reason (given his counterfactual definition) to suppose that he would deny the inference in the other direction.

Perhaps a purely counterfactual conception of defeaters and enablers is too crude. We might instead hold that for a consideration E to be an enabler of a fact F as a reason to x in circumstances C is for E to be a possible feature of C such that E's presence would *explain* why F is a reason to x. This refinement of the conception of enablers and defeaters does not help, though. For being an enabler means that E is a member of a set of possible features of C (E might or might not be the only member of the set) such that the absence of *all* of the features in that set would *defeat* F's status as a reason to x in C. Moreover, the absence of all of a candidate reason's enablers also would *explain* why the candidate is not really a reason. Someone puzzled about why a given

14. Dancy, *Ethics Without Principles*, p. 39.

15. Dancy, *Ethics Without Principles*, p. 41.

consideration was not a reason in a given case could find it illuminating that none of the various further facts which would explain its functioning as a reason is present. Of course, these further facts may be too numerous to list, but we could still perhaps give our interlocutor an idea of the kinds of considerations which would have to be present for the fact to function as a reason by listing representative examples and then explaining that no such feature is present. On this way of thinking, to say something is an enabler in a given context entails that it is a member of a set of possible features, such that the absence of all of the features in that set is a defeater. So once again the fact that a consideration needs an enabler entails that it can be defeated. Similarly, to say a consideration D is a defeater of a fact F as a reason to perform a given action is to say that its presence can explain why the feature is not a reason. This seems to entail that D is a member of a set of possible features such that the absence of all of those features *enables* F to be a reason to perform the action. For if a consideration is a reason in favour of a given sort of action unless defeated then the absence of all defeaters will explain why that consideration functions as reason. So the fact that a consideration stands in need of an enabler entails that it can be defeated and vice versa.

The preceding line of argument is fairly abstract, but the basic point is actually quite simple. An example should help get the idea across. Suppose for the sake of argument that the status of the consideration 'it would cause pain' as a reason has exactly two defeaters: (a) the pain is constitutive of some excellence and (b) the pain is constitutive of a deserved punishment. The fact that considerations of pain have these defeaters also entails that such considerations stand in need of an enabler. The enabler is the conjunctive fact the pain is not constitutive of excellence and is not deserved. For the fact that a given pain is neither constitutive of an excellence nor deserved can explain why the fact that an action causes that pain should function as a reason not to perform the action. Someone who was perplexed at the idea that such a fact might function as a reason could gain understanding if told that the pain is neither constitutive of excellence nor deserved, since the person might have implicitly been assuming one of those conditions was not satisfied. Although a full theory of explanation goes well beyond the present scope, we take it that on any plausible account to figure in an explanation is to track dependence relations in ways that can remove puzzlement about that which is explained thereby answering 'why?' questions about that which is explained. On the toy theory under consideration, the fact that the pain is neither constitutive of excellence nor deserved does track a dependence relation and citing those facts can thereby remove puzzlement and answer the relevant 'why?' questions. This does not mean that the fact that the pain is not constitutive of an excellence nor deserved *fully* explains why the fact that an action promotes pain functions as a reason when it does; they may provide only a partial explanation or be at best part of an explanation (these may not be the same). Since the presence of the pain in question is neither constitutive of excellence nor deserved can at least partially explain why the fact that an

action would produce such pain is a reason not to perform the action, it follows that the fact that the pain is neither constitutive of excellence nor deserved is an enabler. Of course, a particularist may insist that there will be far more possible defeaters for a candidate reason than two, and our example is in this sense just a toy example; the example was given only to illustrate the basic point in a simplified context. The fact (if it is a fact) that any plausible example will involve far more possible defeaters does not gainsay our basic point. That would instead just mean that the enabler in question will be the fact corresponding to a much longer list of possible defeaters.

So far we have used our example to show how the fact that a consideration's status as a reason can be defeated entails that the consideration needs an enabler in order to function as a reason. The example also illustrates how the fact that a consideration needs an enabler entails that it can be defeated. In our example the fact corresponding to the negation of (a)—that the pain is constitutive of an excellence—can explain why the fact that the action produces pain is *not* a reason not to perform the action. Since a defeater just is a fact that can explain why a fact that would otherwise be a reason with a given valence does not function as such a reason, this means that the fact corresponding to the negation of (a) is a defeater. The same point applies, *mutatis mutandis*, to the fact corresponding to the negation of (b). So the claim that a given fact functions as an enabler in the case at hand will entail that certain other facts would function as defeaters in the case at hand. The defeater(s) corresponding to the enabler present in a given case will be the facts corresponding to the negation of each of the conjuncts of the conjunction corresponding to the conjunctive enabler. In the limiting case in which the enabler is not best understood as corresponding to a conjunction we should say that the fact corresponding to the negation of the enabler is a defeater. The main point here is that when a fact's functioning as a reason with a given valence depends on the presence of certain facts, the absence of any one of those facts functions as a defeater and the fact corresponding to the conjunction of those facts functions as an enabler.

So any reason which can be defeated also trivially needs an enabler and vice versa. In that case, though, the distinction between default reasons and other reasons collapses into the distinction between invariable reasons and context-sensitive reasons. For the idea was supposed to be that default reasons are considerations whose status as reasons does not stand in need of explanation *in the sense that they need no enablers*. Given the preceding line of argument, the only reasons which truly need no (local) enablers whatsoever will be reasons which cannot be defeated, and those just are the invariable reasons. In that case, the distinction between default reasons and non-default reasons collapses into the distinction between variable reasons and invariable reasons. So we do not need a further distinction here; we have already marked the distinction between variable and invariable reasons more clearly in terms of variability. More to the point, the collapse of the distinction means that the idea of default reasons (so understood) will not do the kind of work that

particularists needed that distinction to do. For the distinction was supposed to mark an important divide *within* the category of context-sensitive reasons, and not to distinguish the variable from the invariable ones. Hence, drawn in this way the distinction cannot forestall the worry that particularism implausibly flattens the normative landscape. However, the distinction between default reasons and non-default reasons can be drawn in other ways.

II. The Epistemic Reading

If Dancy's metaphysical gloss on the distinction between default reasons and non-default reasons is untenable then it is tempting to try to draw the distinction in epistemological terms. Mark Lance and Margaret Little have developed just such a version of the distinction.¹⁶ Lance and Little endorse the particularist theses that all reasons depend for their status as reasons on context and that virtually any consideration could have moral import provided a suitable context. However, they also hold that only some considerations paradigmatically have a particular moral force. Even considerations, which are paradigmatically reasons, however, are still only defeasibly reasons; there are circumstances in which such considerations are no reason at all or a reason with an opposite valence. To focus on one of their familiar examples, perhaps the fact that an action would cause pain is paradigmatically a reason against the action. Though pain can sometimes be a reason in favour (say, in cases of athletic challenge or justified punishment) of an action it is never in Lance and Little's sense 'defeasibly' a reason in favour of an action. Lance and Little typically refer to such reasons as 'defeasible reasons' to emphasize their defeasibility. However, what is distinctive about these considerations is precisely *not* their defeasibility, for so-called non-defeasible moral reasons (like shoelace colour, for example) are only defeasibly reasons too—if anything, their status as reasons is 'more defeasible'. So 'paradigmatic reasons' seems to capture the idea they are after better than 'defeasible reasons', and we shall here use 'paradigmatic reasons' to refer to the sorts of reasons they have in mind. Little herself at one point uses the 'paradigm' terminology to characterize the idea, suggesting that deviant cases can be understood only 'by reference to a paradigm that carries the privileged valence'.¹⁷

How exactly should we understand what it is for a consideration to be a paradigmatic reason? To say that a given consideration, such as pain, paradigmatically has a certain normative force is to say that it has such a force in certain 'privileged conditions', and they suggest that privileged conditions come in different varieties.¹⁸ Sometimes the privileged conditions are those

16. See Lance and Little, 'Particularism and Anti-Theory', *op. cit.*, and 'Defeasibility and the Normative Grasp of Context', *Erkenntnis* 61 (2004), pp. 435-55.

17. M. Little, 'On Knowing the "Why": Particularism and Moral Theory', *Hastings Center Report* (2001), pp. 32-40 (38).

18. Lance and Little claim that the cases which follow are species of a single genus. It is not clear to us that this is the case.

in which things are morally aright. Thus killing has a negative moral valence unless things are morally amiss. Alternatively, Lance and Little suggest that being a paradigmatic reason can be matter of 'explanatory asymmetry' where explanation is understood as an epistemic idea. As Lance and Little put it, 'We can understand a situation in which lying is wrong-making without resort to any context in which it has the opposite valence; but to understand the moral status of lying in Diplomacy, one must understand the players as having agreed to play a game with these rules in a context in which lying has its typical valence'.¹⁹ On this view, then, even when lying has an abnormal valence, we can only understand this by reference to the normal or paradigmatic case. Put another way, when lying is not a reason against, this must be seen as an exception to a general (but defeasible) norm against lying.²⁰ In at least one context, Little suggests that this epistemological asymmetry is itself explained by a conceptual asymmetry. As she puts it, '...mastery of moral concepts is mastery of defeasible generalizations...one cannot be said to understand moral concepts without appreciating the privileging moves that lie at their heart'.²¹ So whereas one presumably can be competent with moral concepts without appreciating how shoelace colour can be morally relevant, one cannot on this account be competent with moral concepts without understanding that considerations of pain, deception and other familiar moral categories are paradigmatic moral reasons. This conceptual thesis would indeed explain the epistemological asymmetries Lance and Little discuss. However, the epistemological asymmetries presumably could be explained in other ways, so that the interest of what we are here characterizing as the 'epistemological reading' does not depend on the conceptual gloss Little at one point gives it. Here we explore the epistemological conception of default reasons insofar as it can be divorced from this conceptual gloss.

Lance and Little have tapped into some extremely interesting and heretofore neglected territory. In particular, the idea that some considerations are such that we can understand how they can function as a reason with one valence only if we understand how they can function as a reason with the opposite valence is a fascinating one which deserves further investigation. For example, suppose for the sake of argument that a retributive conception of punishment is correct. We do not here mean to endorse retributivism; the point is purely illustrative. We might then think that the fact that an action would cause pain can, in some contexts, be a reason in favour of the action,

19. Lance and Little, 'Particularism and Anti-Theory', p. 30.

20. It is unclear to us whether Lance and Little would claim that every consideration that can be a reason has some practical import defeasibly (even if it defeasibly has the import of being no reason at all either way). Lance and Little claim at once that, 'For something to count as a reason for some given moral status, then whatever its moral import in this case pro or con, we are also attributing to it a defeasible moral import' ('Particularism and Anti-Theory', p. 29). Elsewhere, though, they claim that, 'For while shoelace colour can have various moral imports in various contexts, it has none of them defeasibly' (p. 30).

21. Little, 'On Knowing the "Why"', p. 39.

as when its causing pain of the right sort is part of how the action provides a legitimately imposed punishment for a crime committed. Plausibly we can understand how pain can function as a retributive reason in favour of punishment only if we understand how pain more paradigmatically functions as a reason against actions. For presumably part of the retributive idea is that it is precisely because pain is paradigmatically 'to be avoided' that it is appropriate to impose on someone for certain kinds of crimes. By contrast, one does not need to understand how pain can function as a retributive reason in favour of punishment in certain contexts in order to understand how pain can function as a reason against actions more generally. Indeed, many people who explicitly reject retributivism seem to understand perfectly well how pain functions as a reason against actions more generally. This is a very interesting sort of asymmetry, and it is certainly one which is not true of considerations which can function as reasons in general. So assuming that holism about reasons is true, we can usefully sort candidate reason-giving considerations into two classes—(a) those which are rightly associated with the sort of epistemic asymmetry illustrated by the case of pain/retribution and (b) those which are not. Call those considerations which fall into class (a) asymmetric reasons; call candidate reason-giving considerations which fall into class (b) symmetric reasons.

Our agenda here is not to cast doubt on the general importance of this way of distinguishing reasons, but rather to argue that it cannot perform one of the main jobs for which particularists like Lance and Little invoke it—to avoid the 'flattening of the normative landscape'. Recall that Lance and Little are rightly keen to avoid the conclusion that particularists are committed to treating eccentric considerations like ones of shoelace colour as on a par with considerations of pain and promise-keeping. To equate shoelace colour (for example) with pain is to flatten the normative landscape in a way that is extremely implausible. Intuitively, if we accept the particularist/holist idea that just about any consideration can be a moral reason in the right context we need to draw a distinction between core and periphery such that pain, promise-keeping and pleasure end up in the core while shoelace colour ends up in the periphery. Our main point is that there is no reason whatsoever to think that this important core/periphery distinction will map neatly onto the independently interesting distinction between asymmetric reasons and symmetric reasons.

Before we explain why we think the core/periphery distinction will not map neatly onto the asymmetric/symmetric distinction, we should explain one important respect in which we disagree with Lance and Little about what we are calling asymmetric reasons. Some of Lance and Little's examples suggest that on their view the relevant sort of asymmetry can emerge in cases in which a consideration is sometimes a reason with one valence and sometimes no reason at all. Whereas we maintain that the relevant sort of asymmetry emerges only in cases in which a consideration is sometimes a reason with one valence and sometimes a reason against. So on our reading, only considerations which

can 'go both ways' can fall into the category of asymmetric reasons but on the Lance/Little account (it seems) a consideration can fall into the category of asymmetric reasons even if it can only be a reason with one valence or no reason at all. This emerges most clearly in the example of playing Diplomacy, where lying is no longer a reason against an action. Actually, some particularists seem to suggest that in games like Diplomacy the fact that an action would be a lie actually becomes a reason in favour of the action but in our view this is an extremely implausible gloss on these cases.²² The reason to make a certain move in Diplomacy is that it will increase the odds that you will win or some such, and not that the move is an instance of lying. Poker is a perhaps more familiar example where it should be entirely clear that the reason to bluff is not that it is an instance of deception but rather that it will increase your expected winnings. To their credit, Lance and Little do not fall into this trap, but instead hold that when playing Diplomacy the fact that an action would be a lie is no longer a reason against it. However, this is why their gloss on the case as an instance of what we are calling asymmetric reasons is problematic. For the suggestion is that we can understand the exceptional cases (Diplomacy, poker, etc.) only insofar as we understand how lying functions as a reason against an action in the context in which the players agreed to play this sort of deception-tolerating game. This is fine as far as it goes, but the idea here seems to generalize too easily. For it will be true of any consideration whose status as a reason can sometimes be defeated that we can adequately understand why it is not a reason here only if we understand how it can be a reason elsewhere. For to understand why something is not F here we must in general have some idea of how it can be F elsewhere if it can. If this is enough for a consideration to qualify as an instance of asymmetric reasons then any consideration whose status as a reason can ever be defeated will qualify as an instance of asymmetric reasons, and that makes the distinction far less interesting than it first appeared. Moreover, it certainly would render the distinction incapable of helping the particularist avoid the 'flattening of the normative landscape'.

Incidentally, we think that the preceding point is the kernel of truth in Dancy's remark that, 'one only needs to know the rule in order to understand the *exceptionality* of the case where cruelty is a reason against—but to say this is to say too little to be interesting'.²³ Dancy develops this point as one horn of a dilemma, which is supposed to undermine the general interest of

22. Dancy takes this line about a game called Contraband, on the grounds that, 'if one doesn't do plenty of lying, it spoils the game' (Dancy, *Moral Reasons*, p. 61). Presumably, though, some lies increase one's chances of winning the game while other lies do not. The fact that one must do plenty of lying to avoid spoiling the game suggests that one should lie whenever it would increase one's chances of winning (or at least not decrease one's chances). The reason to tell an individual lie, though, is that it will increase one's chances of winning or that it will make the game more fun; in our view the fact that you would be lying is never the reason in favour of a move in the game.

23. Dancy, *Ethics Without Principles*, p. 115.

the distinction between asymmetric and symmetric reasons. The other horn of Dancy's dilemma insists that, 'It isn't strictly true that in order to understand the case where cruelty is not a reason against, one needs to know the rule. If we demand that, we demand too much'.²⁴ In our view, this dilemma is too quick. For Lance and Little can navigate between the two horns of Dancy's dilemma by insisting that what we must understand is not the exceptionality of a consideration's not being a reason at all; that is trivial. Rather, what is interesting is the idea that we must understand how a consideration can have one valence to understand how it could have the opposite valence but not vice versa, and this does not seem trivial. Nor does this seem to impale the account on the other horn of Dancy's dilemma. For we take it that the example of pain functioning as a retributive reason which can be understood only by reference to how pain normally functions as a reason with the opposite valence suggests that this is not always to 'demand too much'. So long as we allow that pain can sometimes function as a retributive reason, the demand seems to us to be entirely reasonable. Even if we reject the possibility of retributive reasons, the example serves to illustrate that the basic idea is intelligible and interesting; we might only need more imagination to come up with a better example. This does, however, mean that we must depart from what seems to be the Lance/Little conception of asymmetric reasons. On our proposed reading, only what we might call 'bivalent' considerations—considerations which can both function as reasons for and against—can fall into the category of asymmetric reasons. Considerations which are univalent but context-sensitive—that is, can sometimes be reasons with one valence and sometimes no reason at all but never a reason with the opposite valence—cannot fall into the category of asymmetric reasons on pain of trivializing the distinction.

We are now in a position to see why the distinction between asymmetric reasons and symmetric reasons does not map neatly onto the core/periphery distinction and therefore does not help the particularist avoid the flattening of the normative landscape. The Lance/Little idea seems to be that what we have called core reasons (pain, pleasure, promising and the like) will line up with what we have called asymmetric reasons while periphery reasons (shoelace colour, and the like) will line up with symmetric reasons. In our view, neither of these alignments holds. First, there can be core reasons which are not asymmetric reasons. Indeed, any consideration which can never flip valences even though it can sometimes be no reason at all will not be an asymmetric reason, and it would be rather surprising if no core reasons were not in this sense univalent albeit context-sensitive. Perhaps the fact that an action would be deceptive is a good example, as it seems to us to sometimes be a reason against (the normal case) and sometimes is no reason at all (poker, Diplomacy, etc.) but is never a reason in favour of an action. Moreover, even

24. Dancy, *Ethics Without Principles*, p. 115.

some core reasons which plausibly can flip valences do not look like instances of asymmetric reasons. For example, suppose we agree with the standard particularist idea that sadism can turn considerations of pleasure into a reason against an action. We take it that pleasure is plausibly a core reason in favour of an action. However, we do not think that someone must understand how pleasure can provide a reason in favour in order to understand how sadistic pleasure can provide a reason against. To understand the former one might need only to understand 'what pleasure is like' in some suitable sense, while to understand the latter one needs only to understand that to take pleasure in something is a kind of endorsement and that pain ought not be endorsed in any way. Someone with a puritanical bent might not even believe that pleasure can ever provide a reason in favour of an action but still understand perfectly well how in the context of sadism pleasure can provide a reason against. So core and asymmetric reasons can come apart in at least two ways. First, there can be core reasons which are not even bivalent and hence not candidates for being asymmetric reasons. Second, there can be core reasons which are bivalent but which nonetheless are not asymmetric reasons as, in the case of sadism, as understood by particularists. So core reasons and asymmetric reasons can come apart and are not the same.

Furthermore, there is no obvious reason to suppose that there could not be instances of asymmetric reasons which are not core reasons. We are hampered in arguing here in that we do not think that so-called periphery reasons like shoelace colour really can ever be reasons. Moreover, we find it rather hard to 'get inside the particularist's head' on this point and see what follows from a charitable gloss of their view on these matters. For although particularists often seem to suggest that considerations like shoelace colour can sometimes be reasons for action, they rather frustratingly never provide examples to illustrate this suggestion. We suspect that this is because any proffered example will be highly contestable—shoelace colour might influence the consequences of an action but then the generalist can plausibly respond that it is the consequences that really provide the reason for action and not the shoelace colour. The particularist must mean more than this, as we have seen that even a hedonistic act utilitarian or dyed-in-the-wool Kantian can allow that shoelace colour might sometimes matter morally in various indirect ways (by figuring in a promise or influencing the consequences, for example). The more radical suggestion here seems to be that shoelace colour itself can provide a reason, but no credible example has been given to illustrate this rather odd idea. Perhaps there has been a bit too much loose talk about shoelace colour. However, we see no reason *ex ante* to assume that if (contrary to fact) shoelace colour really could sometimes provide a reason that it might not also be an asymmetric reason. From our perhaps naïvely generalist perspective, once we assume that shoelace colour can sometimes be a reason we see no *a priori* grounds for supposing that one could only understand how it can function as a reason with one valence if one understands how it functions with the opposite valence but not vice versa. The challenge for the particularist who wants

to employ this distinction is to provide some reason for thinking that eccentric non-core reasons like shoelace colour could never be asymmetric.

The distinction between asymmetric reasons and symmetric reasons is only one member of a family of distinctions drawn by Lance and Little. We have devoted so much attention to this distinction only because it seemed to us to provide a more *prima facie* promising way of avoiding the flattening of the normative landscape than the other distinctions they draw. Moreover, we have related worries about their other distinctions. For example, we see no reason to suppose that the core/periphery distinction will map neatly onto the distinction between considerations which function as a reason with a given valence only when things are morally aright and those of which this is not true. We should, however, return to the questions raised by the more general characterization of their distinction in terms of privileging conditions instead of going through the various and sundry species of that genus.

The framework for the distinctions drawn by Lance and Little is cast in terms of asymmetry in privileged conditions. An obvious alternative would be to explicate the relevant asymmetry in terms of our local conditions. In our world, instances of pain might have a high probability of being such that there is reason not to promote them. By contrast, shoelace colour has at best a very low probability of being a reason of any valence. Most often, it is no reason at all. Whatever the merits of this idea it is emphatically not what Lance and Little have in mind. To be paradigmatically a reason is not merely to usually be a reason.²⁵ In characterizing pain as paradigmatically providing a reason against performing an action that promotes it, Lance and Little insist we must mean more than that pain is likely to be a reason against in our local conditions. Once we put the conceptual gloss on their distinction to one side, what more could we mean?

To say that pain is paradigmatically a reason against, Lance and Little would argue, is at least in part to claim propriety for a certain inference—the inference from

(P) *X is and instance of pain.*

To:

(R_{Ag}) *There is reason not to promote X.*

The issue is what kind of propriety the inference has. It is not local statistical reliability, for two reasons. Even in a world in which most instances of pain were not reasons against, pain would be defeasibly a reason against. Imagine, for example, a variation on an example of H.L.A. Hart.²⁶ A world is populated by rational creatures whose crab-like bodies feature tough exoskeletons. It is very difficult to inflict pain on such a creature from the outside,

25. See Lance and Little, 'Particularism and Anti-Theory', pp. 27-28. For an example, they suggest that fish eggs normally turn into fish, even if the vast majority of fish eggs never do so.

26. The original example is from H.L.A. Hart, *The Concept of Law* (Oxford: Clarendon Press, 1961), p. 194.

and so there are very few instances of such pain. Nevertheless such creatures are prone to pain when they exert themselves in athletic competition, and most pain is like this. In such a world, the inference from (P) to (R) is unreliable, but even in such a world we might say the pain has the same 'default' status that it has in our world. Second, even if a given inference is reliable in local conditions, this does not establish the kind of norm Lance and Little have in mind. In our world of crab-like creatures, the inference from

(P) *X is an instance of pain.*

To:

(RFor) *There is a reason to promote X.*

is a reliable one. However, even in this world pain is not (in Lance and Little's sense) a paradigmatic reason in favour of actions that promote it. So local reliability is neither necessary nor sufficient for the inferential propriety Lance and Little have in mind.

Instead of thinking of the inferential propriety in question as reliability in our local circumstances, Lance and Little argue that we should think of it as propriety in 'a set of possible worlds whose contours are set by the goals of the given discipline whose inferential proprieties are at issue'.²⁷ In a similar vein, Marc Lange has argued that scientific laws can admit of exceptions and still qualify as laws; laws are associated with reliable inferences, not exceptionless regularities. Moreover, generalizations need not be stable (or the associated inference reliable) under all possible counterfactual circumstances. The counterfactual conditions under which a generalization must be preserved to qualify as a law depends upon the larger set of generalizations employed by a given science and the questions that such a science is trying to answer.²⁸ As Little puts it, 'In saying that explanations situate instances within broader patterns, we are saying, roughly speaking, that when we offer one phenomenon as explanatory of another we are asserting our confidence that the latter always follows from the first within a suitable constellation of possible worlds'.²⁹

This, however, is not enough to distinguish paradigmatic and non-paradigmatic reasons. For every reason will be associated with a reliable inference in a suitable constellation of worlds (or circumstances). The inference will be reliable in all those worlds (or circumstances) in which any relevant enabling conditions are present and no defeating conditions are present.³⁰ It is worth noting that this objection need not apply to Lange's proposal for understanding scientific laws. Here two points are relevant. First, the laws of nature are plausibly thought of as contingent; there are at least other logically possible worlds where the laws of nature that obtain in our world do not and other

27. Lance and Little, 'Particularism and Anti-Theory', p. 28.

28. See M. Lange, 'Who's Afraid of Ceteris Paribus Laws? Or: How I Learned to Stop Worrying and Love Them', *Erkenntnis* 57 (2002), pp. 407-23.

29. Little, 'Moral Generalities Revisited', p. 300.

30. Here we simply set aside possible 'meta-defeaters'.

laws obtain instead. Second, laws of nature, like moral principles, can be hedged in various ways. When a purported natural law of the form *All Fs are G* fails to hold, we might then draw three lessons. First, and most obviously, we might treat the case as a counterexample; the purported law is no law at all. Second, we might understand the law *All Fs are G* to be hedged in various ways and decide that the case is an 'exception' of the kind built into the law. Third, we might decide that the case falls outside the range of circumstances for which the law is to hold. As a possible example of this third case, consider a suggestion made by Marc Lange. According to the 'area law' of biogeography, the biodiversity of an island increases exponentially with the island's area. Such a law would doubtless need to be qualified by a *ceteris paribus* clause. In some cases, the biodiversity of an island might diverge from the predictions of the area law because of the presence of some 'disturbing' factor. In other cases, the divergence might be in a case for which the area law was never designed. As Lange puts it,

Biogeographers are interested in how species would have been distributed had (say) Gondwanaland not broken up, and in how Montserrat's biodiversity would have been affected had the island been (say) half as large. On the other hand, biogeography is not responsible for determining how species would have been distributed had Earth failed to have had the Moon knocked out of it by cataclysm early in its history. Biogeographers do not need to be geophysicists.³¹

Thus even were it successfully argued that if the Moon had not been knocked out of Earth, the biodiversity of islands would not correspond to the predictions of the area law, this would not falsify the area law or show it not to be a law. Lange's claim will be contested, and it is not our aim to defend it. We do wish to argue, however, that Lange's suggestion is not one that can be especially helpful in drawing the distinction between default and non-default reasons.

As the example of the area law makes clear, the exceptions tolerated by the area law are real possibilities. Presumably there is some fact about how much biodiversity there would have been if the Moon had never been. This is a fact about biodiversity and about geography, but, if Lange is right, it falls outside the scope of the concerns of biogeography. Now consider the analogous suggestion for moral theory. There are moral facts, but these fall outside the scope of the concerns of moral theory. Perhaps there are cases in which shoelace colour is a reason, but these are cases that occur in worlds (or circumstances) that fall outside of moral theories' concerns. There are two problems with this suggestion. The first is that we expect moral principles (if any there be) to be necessary truths. Kant gave powerful expression to this idea when he claimed that if there are moral laws at all they must hold of all rational beings with absolute necessity.³² Nevertheless, the scope of the necessity that a valid moral

31. Lange, 'Who's Afraid of Ceteris Paribus Laws?', p. 418.

32. I. Kant, *Foundations of the Metaphysics of Morals* (New York: Macmillan, 1990 [orig. 1784]): K 389.

principle must have is a question that goes to the heart of the nature of moral theorizing and that has divided the Aristotelian and Kantian traditions in normative theory.

In the present context, however, there is no need to beg the question against alternative views of moral theory and moral necessity. For the present suggestion encounters a second difficulty. The distinction between default and non-default reasons was meant to mark a distinction *within* the concerns of moral reflection (whether this takes a particularist or a generalist form). If we apply Lange's suggestion to the moral case, though, then we seem to be marking a distinction between the moral facts that are of concern to normative theory (and so are relevant to what generalities or principles are true) and the moral facts that occur in worlds (or circumstances) that moral theory should ignore. This is not the distinction we were after. We started with the insight that of the many things that can be morally salient (and are rightly seen as such by virtuous agent and moral theorist alike) some are more salient. We cannot serve that insight by introducing moral facts that fall outside the 'constellation of worlds' that concern us.

There is a further worry about the foregoing approach, at least for those particularists who have robustly anti-theoretical aspirations. If the 'exceptions' to moral principles occur, so to speak, offstage³³ in esoteric possible worlds we can reasonably just ignore then particularism poses little challenge to traditional normative moral philosophy. Indeed, it would seem that even an advocate of hedonistic act utilitarianism could embrace particularism by arguing that the only cases in which the right act and the act that maximizes pleasure are not one and the same are cases that lie outside the concerns of moral philosophy and so do not count against his theory. Many particularists have taken precisely such simplistic attempts to codify morality as their target.³⁴

III. The Pragmatic Reading

We have seen that both the metaphysical reading and the epistemological reading of the distinction between default reasons and non-default reasons are problematic. However, there is one final reading of the distinction which is entirely uncontroversial—the pragmatic reading. Indeed, one might plausibly suggest that the pragmatic reading is the default interpretation of default reasons. The pragmatic reading begins with the observation that which features it makes sense to mention when offering an explanation depends on the

33. We owe this apt phrase to Lange.

34. Although we are confident that they would not accept hedonistic act utilitarianism, the basic point we are making here may not be an objection to Lance and Little. According to them, at least, particularists have typically misapplied their own insights. Particularism is not a challenge to moral principles or laws. It simply requires that we 'revisit what those laws look like'. Lance and Little, 'Particularism and Anti-Theory', p. 32. On this view, perhaps, particularism stands to normative theory much as debates about what a law of nature is stands to practising science.

context. In particular, facts which are known to be common knowledge need not be made explicit when giving an explanation. For example, suppose her friends ask Sally why Harry is so upset and she simply replies, 'He ran into Helen today'. In the right context this might be a perfectly good reply even though as an explanation it is highly elliptical. The adequacy of this concise reply relies on the fact that it is common knowledge, and known among the parties to the conversation to be common knowledge, that Harry recently went through a bitter divorce with Helen and has not gotten over her yet. Had her interlocutors not known Harry's relation to Helen then Sally would need to make these further facts explicit. The same point applies to moral explanations. If asked why it would be wrong to invite Helen to the party, Sally might simply tell us that her presence would make Harry upset. Sally would typically not need to make explicit that inviting Sally is not the only way to stop a catastrophic terrorist attack.

This already gives us the resources to draw a distinction between default reasons and non-default reasons against the backdrop of holism and particularism. On this account, whether a given consideration is a default reason with a given valence is itself context-dependent, and in this sense in the holistic spirit of particularism itself. For the suggestion is that whether a given consideration is usefully understood as a default reason with a given valence depends on whether one's interlocutors would take it as given that such a consideration would function as a reason with that valence unless some special story was explicitly told as to why it did not so function. That an action would cause pain is in most contexts a default reason against the action. Why? Because the fact that an action would cause pain almost always is a reason against the action and this is common knowledge and moreover it is commonly known to be common knowledge. If someone was puzzled about why a given action was objectionable we could typically dispel this puzzle simply by pointing out that the action would cause some pain. Moreover, this does not presuppose that the person is an atomist and assumes that pain always counts against an action. For example, our interlocutor might agree that in the context of an S & M room, the fact that an action would cause some pain might well speak in favour of the action. The point then would be that such contexts are known (and commonly known to be known) to be very rare, and this is why it can go without saying that this is not the case when it is not. Hence in most local conversational contexts, that an action would cause pain is a default reason against the action but also is sometimes a non-default reason in favour of an action in certain contexts. We do not need any fancy metaphysical or epistemological background theory in order to draw a perfectly sensible contextualist distinction between default reasons and non-default reasons.

Moreover, in our view, the pragmatic reading is considerably more plausible than its rivals and not only because it avoids the objections pressed against those rivals in sections I and II. It also fits better with a reflective understanding of the sources of our own intuitions. For those intuitions themselves seem

highly sensitive to contingent features of our situation. Consider again our previous example of athletic crab-people. It seems very likely that such people would not find the idea that pain provides a default reason against an action very plausible, assuming that they all agree that being constitutive of athletic accomplishment is sufficient (meta-defeaters to one side) to make pain into a reason in favour of the action which causes it. To take another example, suppose that most instances of pleasure were sadistic. If we agree with the standard particularist idea that sadism is a defeater of the reason-giving force of considerations of pleasure then (meta-defeaters to one side) this would entail that most instances of pleasure would not provide a reason to perform actions which would promote the pleasure. In a world in which this was the case and widely known to be the case, we submit that the intuition that pleasure provides a default reason in favour of an action that promotes pleasure would not be widely held.³⁵ Indeed, the opposite intuition might well be dominant. This fits perfectly with the pragmatic reading, for whether a consideration is a default reason on the pragmatic reading does depend on local conditions. This suggests that it would be a mistake to interpret our own intuitions about default reasons in a metaphysically or epistemologically grandiose way.

Is drawing the distinction pragmatically enough to save the particularist from the objection that their account implausibly flattens the landscape? Perhaps. Here we must simply assume with the particularists that all sorts of considerations can sometimes be reasons for action, even though we do not actually accept this thesis. For the objection that the particularist flattens the normative landscape does not presuppose that such considerations can never be reasons for action, only that their status as reasons is in some way second-rate. Certainly *if* shoelace colour sometimes provides a reason for action this is in our world very rare. Whereas pleasure and pain typically do provide reasons for action. So in our world it is quite reasonable in almost all conversational contexts to treat pleasure and pain as default reasons but not treat shoelace colour and other such bizarre considerations as default reasons even if we agree that the latter can in suitably weird circumstances also provide reasons for action. This avoids the charge that the particularist flattens the actual normative landscape, but is that enough? For it now seems that in a world in which shoelace colour typically did provide a reason with a given valence (and was known to do so, and known to be known to do so) it would be appropriate to treat shoelace colour as a default reason as well. This does indeed follow from what we take to be the best way of drawing the distinction

35. One peculiarity of the case is that sadists might be unlikely to realize that their sadistic pleasures do not count in favour of an action. However, we could handle this situation by stipulating that we are now talking about a species of rational creatures most of whom are incapable of pleasure, but a very few of them are capable of pleasure albeit only sadistic ones. It might well be common knowledge amongst the members of such a species that pleasure does not typically provide a reason in favour of actions which promote it.

between default reasons and non-default reasons and perhaps this is embarrassing for the particularist. However, in our view the embarrassment here stems not from the pragmatic reading of default reasons but rather from the particularist commitment to the thesis that weird considerations like considerations of shoelace colour can even sometimes be a reason for action. In our view, such weird considerations can themselves *never* be reasons for action. Hence we agree that there is a deep and non-pragmatic asymmetry between shoelace colour and pain; the former can never be a (primary) reason for action while the latter can. Which is just to say that the best way to accommodate such deep asymmetries is to reject particularism's highly ecumenical account of reasons for action.

Conclusion

Particularism threatens to flatten the normative landscape, marking no normative distinction between shoelace colour and pain. Particularists typically try to avoid this conclusion by drawing a distinction between default reasons and non-default reasons, but the ambitious metaphysical and epistemological glosses they typically give this distinction are either problematic in themselves (the metaphysical reading) or interesting but unpromising as a way to avoid the flattening of the normative landscape (the epistemological reading). The distinction should instead be understood pragmatically. This might be enough for the particularist to avoid the charge that their account flattens the actual normative landscape. However, in certain counterfactual contexts particularism would flatten the normative landscape. In our view, the solution to this problem is not to resurrect some version of the metaphysical or epistemological reading of default reasons. Rather, we should simply reject the particularist's highly promiscuous attitude to what kinds of considerations can be reasons for action. However, a full defence of this diagnosis of the particularist's error must await another day.