

PENROSE MEMORIAL LIBRARY
WHITMAN COLLEGE
WALLA WALLA, WASHINGTON 99362

Moral Particularism

Edited by
BRAD HOOKER and
MARGARET OLIVIA LITTLE

CLARENDON PRESS · OXFORD
2000

The Particularist's Progress

Jonathan Dancy

Many people seem willing to call themselves feminists with little idea of what they are committing themselves to. The same, in my experience, is true of particularism in the theory of moral reasons. There is a common suggestion that to be a particularist is, at the outset, only to admit that circumstances can make a difference. But if that were all that particularism amounted to, it would be uncontentious. In this chapter I lay out what I think one commits oneself to if one accepts the general claim that reasons are sensitive to context—a claim sometimes called holism in the theory of reasons, and of which moral particularism is merely one expression.

Of course holism here, as elsewhere, does come in degrees. The strongest form of context-sensitivity would be the claim that every reason is somehow altered with every change of context. The weakest form is the claim that some reasons are on occasions capable of being altered by a change in context. The form of holism that I recommend is pretty weak on *this* scale, so far as the modality goes, but strong on the extent of the domain. I maintain that *all* reasons are *capable* of being altered by changes in context—that there are none whose nature as reasons is necessarily immune to changes elsewhere.

When I talk of altering a reason, I mean to suggest not that the consideration which is a reason is altered, but that its nature as a reason changes. Instead of being a reason in favour of some course of action, it ceases to be a reason for action at all, or even becomes a reason against. One could express this by saying that the practical relevance of the consideration at issue is sensitive to changes in context, and the practical relevance of the

I am grateful to Eve Garrard and David McNaughton for many discussions of the issues discussed in this paper; also to Roger Crisp, Brad Hooker, Derek Parfit and to all those audiences around the world that have let me try to persuade them of the merits of particularism. I owe special thanks to All Souls College, Oxford, where as a Visiting Fellow I wrote an early draft of this paper.

consideration includes its *polarity*. A consideration reverses its polarity when, having been a reason in favour of action, it becomes a reason against, or vice versa. My holism holds that every consideration is capable of having its practical polarity reversed by changes in context.

It is hard to be sure quite how extreme a claim this is, partly because of the awkward modality in its characterization. But I shall not be discussing that matter much here. Perhaps I will have to admit that not all reasons are sensitive to context in this way—that there are a privileged few, including probably the intentional inflicting of undeserved pain, which necessarily constitute the same sort of reason wherever they occur. If so, I will have lost a battle but won the war. For the main aim of my particularist position is to break the stranglehold of a certain conception of how moral reasons function—the *generalist* conception under which what is a moral reason in one situation is necessarily the same reason wherever it occurs. Generalism need not be false of *every* moral reason in order to be largely false, and hence false as a general account of moral reasons and the way they work. And if it is false as a general account of such reasons, rational constraints on moral thought and action—in particular, accounts of what consistency requires in these areas—must not themselves be based on generalist assumptions.

It may be that my train of thought here is vitiated by being overconcerned with one specific opposing account of how reasons function—Ross's theory of *prima facie* reasons. But I might as well admit that I do have this theory constantly in mind, since it seems to capture so well the outlines of the position I am trying to dislodge. Omitting Ross's epistemology for the moment, the theory of *prima facie* reasons holds:

1. What is a reason in one case is the same reason in all.¹
2. Judgement is the attempt to determine the balance of reasons, so conceived.

Holism in the Theory of Reasons

In this section I argue in favour of particularism in ethics. In the past I tended to argue largely from example.² This persuades some people but not

¹ This is not fair to Ross: see D. McNaughton's 'An Unconnected Heap of Duties?' *Philosophical Quarterly*, 46 (1996), 433–47.

² See my *Moral Reasons* (Oxford: Blackwell, 1993), ch. 4.

others. Here my argument will be more theoretically grounded—though there will still be considerable use of examples as well.

As I said in the preamble, I see ethical particularism as merely one expression of an overall holism in the theory of normative reasons—that is, in the theory that discusses the reasons that favour one thing (action, belief) over another. Such an overall holism can be expressed as follows:

1. What is a reason in one situation may alter or lose its polarity in another.
2. The way in which the reasons here present combine with each other is not necessarily determinable in any simply additive way.

There are theoretical reasons and practical reasons, reasons for belief and reasons for action. My holism is intended to hold on both sides of that distinction. I start by trying to establish that theoretical reasons are holistic. We will quickly find that theoretical reasons are perfectly capable of changing their polarity according to context, without anyone making the slightest fuss about the matter. For instance, suppose that it currently seems to me that something before me is red. Normally, one might say, that is a reason (*some* reason, that is, not necessarily sufficient reason) for me to believe that there is something red before me. But in a case where I also believe that I have recently taken a drug that makes blue things look red and red things look blue, the appearance of a red-looking thing before me is reason for me to believe that there is a blue, not a red, thing before me. It is not as if it is some reason for me to believe that there is something red before me, but that as such a reason it is overwhelmed by contrary reasons. It is no longer *any* reason at all to believe that there is something red before me; indeed it is a reason for believing the opposite.

As I say, it seems to me that nobody ever thought of denying what I am claiming here. I know of nobody who has nailed themselves to an atomistic (i.e. non-holistic) conception of how theoretical reasons function. If generalism is taken to be the view that all reasons are general reasons, i.e. that if a feature is a reason in one case, it is the same reason in any other case, generalism is uncontroversially false of theoretical reasons.

Let us now turn to ordinary practical reasons. We will find just the same thing there. There are plenty of examples to persuade us that such reasons are holistic (or non-generalist, if you like). For instance, that there will be nobody much else there is sometimes a good reason for going there, and sometimes a very good reason for staying away. That one of the candidates wants the job very much indeed is sometimes a reason for giving it to her

and sometimes a reason for doing the opposite. And so on. Now examples would be of little use if there were some theoretical obstacle to taking them at face value. But again we should remind ourselves that nobody has ever really debated the question whether ordinary practical reasons are holistic or not. There should be no *parti pris* on this issue; so the examples, which are legion, should be allowed to carry the day without resistance.

Perhaps this is too quick. There is a theory-based reason for doubting my claim that practical reasons are holistic, one that derives from the common thought that practical reasons are grounded in desires of the agent in a way that theoretical reasons are not. What one wants should not affect what one judges to be the case, on pain of charges of bias or prejudice. But what one wants can perfectly well affect what one has reason to do. Indeed, many find it hard to conceive of our having any practical reasons at all if we had no desires. My own view on this matter, however, is that desires do not give us or ground our reasons. Reasons stem from the prospect of some good. If we have no other reason to do a certain action, wanting to do it will give us no reason at all; nor can wanting to do a silly action make it marginally less silly. (These are only the first moves in a long debate.³ I mention them here only to show the sort of way in which I find myself denying the possibility of grounding practical reasons in desires of the agent.) This view of mine is, of course, an independent input in the present debate. I mention it only to show that a certain motive for doubting the analogy I have been drawing between theoretical and practical reason is itself contentious.

It may be that here we come across the real motivation for generalism in the theory of practical reason—an adherence to the view that reasons for action are partly grounded in desires. For if we accept that view, and if we then think of desires as giving the desirer the same reason wherever the desire occurs, we will at least get the sort of generalism I discussed above. The right response to this, however, is to claim that even if all practical reasons are grounded in desires, the same desire need not always function as the same reason. Consider first the third-person case. That he wants power and she does not may be a reason to give the power to her rather than to him, as I have already said. (It may at the same time be a reason to give it to him, since according to me one feature can be a reason on both sides at once; but remember that here it is a reason not to give it to him, and that it need not always be such a reason.) Now consider the first-person case. Suppose that I am trying to train myself into indifference towards a girl. I want very much to spend time with her. But I also want not to have this

³ For the remainder of the debate, see ch. 2 of my *Practical Reality* (Oxford: Clarendon Press, 2000).

want, since she is permanently indifferent to me. It is better for me not to think of her at all. If I spend time with her, this will make things worse for me rather than better—so long as I have not yet succeeded in training myself into indifference towards her. Once I am indifferent towards her, I can spend time with her without loss. In this situation, it seems, my desire to spend time with her may be a reason for me not to do so.

Before carrying on to consider moral reasons, which have been claimed to be non-holistic, I want to step aside for a moment to ask whether I have not already made a mistake. There is a distinction between epistemic and what one might call constitutive reasons. An epistemic reason is a reason for believing something or other; a constitutive reason is a reason why something or other is the case. That the butler's fingerprints are on the murder weapon is a reason for believing that he did the deed, but no part of what makes it the case that he did it or of why it is true that he did it. That the hedgehogs are hibernating early is a reason for believing that we will have a severe winter, but not any part of what makes it the case that the winter will be severe. And so on. Now holism in the theory of reasons should concern itself with constitutive reasons rather than with epistemic ones. But I appear to have argued only that epistemic reasons are holistic, for my first example, or fulcrum, concerned reasons to believe that there was a red thing before me. It is, therefore, technically irrelevant.⁴

This is true, and I apologize for it. But matters can be redeemed. We should not suppose that all that I have shown is that epistemic reasons are holistic, it being left entirely open whether constitutive reasons are or are not. For many, possibly most, epistemic reasons are also constitutive. For instance, that an action involves the gratuitous inflicting of pain is held by many to make it wrong, but equally clearly functions as a reason to believe that the action is wrong. It is both an epistemic and a constitutive reason. Some epistemic reasons are not constitutive, and *perhaps* some constitutive ones are not epistemic; this is all that can be said. Now could it be the case that the epistemic ones are holistic but the constitutive ones are not? I think that this is inconceivable. The mere fact of the overlap between reasons of the two sorts should give us pause. But more importantly, can we suppose that the very logic of epistemic reasons is capable of differing at a very deep level from that of constitutive reasons? This supposition entirely undermines the sort of connection there needs to be between reasons why things are so and reasons for taking them to be so.

⁴ Thanks to Nick Zangwill for pointing this out to me.

I return, therefore, to the onward or outward spread of holism. So far we have it that theoretical reasons (constitutive ones) are holistic, and so are ordinary practical ones. Now could it be the case that moral reasons are quite different from others in this respect, being the only atomistic ones? This is what many have supposed, in supposing that moral rationality is based on the existence of a range of moral principles. Moral reasons, they have held, necessarily behave in regular (or rule-bound) ways, though other reasons see no need to behave in that way at all. About this I want to say that straight off it just seems incredible that the very logic of moral reasons should be so different from that of others in this sort of way. Consider here the sad fact that nobody knows how to distinguish moral from other reasons; every attempt has failed. How does that fit the suggestion that there is this deep difference between them? Not very well at all. Then of course there are examples to be considered, examples of apparently moral reasons functioning in a holistic way. I forbear to bore you with these. It just seems inevitable that moral reasons should function holistically in the way that other reasons do.

This certainly makes it hard to hold, as many do, that the very possibility of moral distinctions, of moral thought and judgement, is predicated on the existence of a range of moral principles. Moral principles, however we conceive of them, seem all to be in the business of specifying features as *general* reasons. The principle that it is wrong to lie, for instance, presumably claims that mendacity is always a wrong-making feature wherever it occurs (*pro tanto*, of course, not necessarily absolutely). It cannot be merely a generalization, a claim that lies are mostly the worse for being lies, for if all moral principles were of this sort the argument that moral thought and judgement depends on the possibility of moral principles would simply be the argument that such thought is impossible unless there is a considerable preponderance of normal cases over abnormal ones. I have never seen this argument made, and I doubt, what is more, whether it would be persuasive if restricted to ethics.

If moral reasons, like others, function holistically, it cannot be the case that the possibility of such reasons rests on the existence of principles that specify morally relevant features as functioning atomistically. A principle-based approach to ethics is inconsistent with the holism of reasons.

All the same, it might be argued, we have to admit that there are some invariant reasons—some features whose practical relevance is invariant. And surely I should allow this, because holism, as I expressed it, concerns only what may happen, not what must. It could be true that every reason may alter or lose its polarity from case to case, even though there are some

reasons that do not do this. If they don't do it, this will be because of the particular reasons they are. Invariant reasons, should there be any, will be invariant not because they are reasons but because of their specific content. And this is something that the particularist, it seems, should admit. It is like the claim that a man can run a mile in four minutes, that Sam Smith is a man, and that Sam Smith cannot run a mile in four minutes. These claims are compatible, and so are the claims that reasons are variable *qua* reasons though some reasons are (necessarily, given their content) invariant. The invariance, where it occurs, derives not from the fact that we are dealing here with a reason but from the particular content of that reason.

So can the particularist admit the existence of *some* invariant reasons? The obvious examples are things like the causing of gratuitous pain on unwilling victims. Surely, it is commonly urged, this is always for the worse, even if over all we might in some case be morally forced to do it. Well, the first thing to say is that admitting the possibility of some invariant reasons is a far cry from admitting that the very possibility of moral thought and judgement is dependent on our being able to find some such reasons. To support any such suggestion, we would somehow need to be able to locate a sufficient range of invariant reasons, ones that together somehow covered the moral ground entirely and themselves explained the nature and role of the variant reasons. This is quite a different matter from simply trying to refute particularism (which is merely an application of holism in the general theory of reasons to the moral case) by producing one counter-example of an invariant reason, which is normally what is going on.

Further, we should remember that the question whether reasons are atomistic or holistic is a very basic question about the nature of rationality, of how reasons function from case to case. It is, I suppose, conceivable that though the vast bulk of reasons function holistically, there are a few that function atomistically. But if this were true we would have a hybrid conception of rationality. There would just be two sorts of reasons, each with their own logic, and moral thought would be the uncomfortable attempt to rub such reasons together. It is *much* more attractive, if at all possible, to think of our reasons as sharing a basic logic, so that all are atomistic, or all holistic.

Let us consider, then, how the supposed invariant reasons function as reasons in the particular case. Take the well-known example of the fat man stuck in the only outlet from a cave that is rapidly filling with water from below. We and our families are caught in between the fat man and the rising water. But we have some dynamite. We could blow the fat man up and get out to safety. But the fat man is unwilling to be blown up (he, at least,

is safe from drowning); and, let us immediately admit, he is blameless in being where he is, and in being fatter than the rest of us. So what we propose to do involves the destruction of an unwilling and blameless victim. As such, we might say, this is *some* reason against lighting the fuse and standing back. The question I want to raise is whether this feature (that we are causing the death of an unwilling and blameless victim) is functioning as the reason it here is, in any way that is to be explained by appeal to the (supposed) fact that it functions in the same way in every case in which it occurs at all. It seems to me that this feature is the reason it is here quite independently of how it functions elsewhere.

Of course if the feature is genuinely an invariant reason, this fact, should we discern it, will be of use to us in any case where we might be in doubt as to the contribution it is making. We can say 'This is an invariant reason, it makes such-and-such a difference there, and so it must be making that difference here.' But suppose that we were to treat one of these supposedly invariant reasons as potentially variant, so as to deny ourselves the use of that inference. What sort of mistake would we have made? Would it be a failure of *rationality* to treat an invariant reason as potentially variant, or just a mistake of fact? I suggest that the invariance of the reason is an epistemic matter rather than a constitutive one. That the reason functions invariantly is a clue to how it is functioning here, but in no way constitutes the sort of contribution it makes to the store of reasons here present. In that sense, the invariance of its contributions is not a matter of the logic of such a reason, and failure to treat the reason as functioning invariantly is not a failure to understand how it functions *as a reason*. It is a perfectly good reason case by case without our worrying about how it operates elsewhere.

I conclude, then, that particularism should accept the possibility of invariant reasons, so long as the invariance is not a matter of the logic of such reasons, but more the rather peculiar fact that some reasons happen to contribute in ways that are not affected by other features. We can admit this without adopting a hybrid theory of rationality, so long as we treat the invariance of any invariant reason as an epistemic matter rather than as a constitutive one.

Holism in the Theory of Value

The next question concerns whether our holism in the theory of reasons spills over to generate a holism about *value*. This new holism, value holism, can come in various forms, just like the holism of reasons. In broadest

outline, in my hands it will amount to the claim that for any x that has value in some context, x may have a different value or none at all in other contexts, and if there is disvalue as well as value, x may have value in some contexts and disvalue in others.

Is there any prospect of accepting a holism of reasons and denying value holism? This is not a matter on which there is a long history of debate. Presumably we should approach the issue by thinking about the relation between reasons and values. Thomas Scanlon has recently revived the view that value is to be understood in terms of reasons.⁵ He defines a valuable object as one that has features that give us reasons to protect, promote, admire, respect, approve of it (etc.: this list is open-ended). On this view, it seems inevitable that a holism of reasons will generate a holism of value. There are, however, other views to be considered. One might allow that wherever there is value, there are reasons, but leave it open whether there can be reasons that are not directly connected with value.⁶ In this way one would make room for certain forms of deontology, without moving very far from the Scanlon position. Michael Slote and Roger Crisp suggest that there are some reasons that do not derive from values.⁷ One might also ask whether there are agent-relative values as well as agent-relative reasons; perhaps agent-relative reasons do not 'stem from' values at all. There is, then, a spectrum of views to be considered. On the most trenchant views, a holism of reasons must be matched by a holism of values. But even on the less trenchant views, it still remains possible that the holism of reasons must match that of values. Suppose we agree for a moment that *most* reasons are linked to values. ('Linked to' is the vaguest phrase I can think of.) And suppose that values are atomistic, i.e. insensitive to context. Immediately we have the problem of explaining how it can be that the reasons linked to those atomistic values are able to be holistic at all. Why are they not atomistic too? Of course, if there are some reasons that are not linked to values, and those too are holistic, we might think this some justification for supposing that the explanation of their holism will not make reference to any holism of values. And we could then argue that, by parity of reasoning, the holism of other reasons is to be explained in a similar

⁵ See Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1999), 95–7). Similar views were considered by W. D. Ross and adopted by A. C. Ewing. See Ross's *Foundations of Ethics* (Oxford: Clarendon Press, 1939), 278–83, and Ewing's *The Definition of Good* (London: Macmillan, 1947), 148–9.

⁶ I consider this possibility in my 'Should we Pass the Buck?', in A. O'Hear (ed.), *The Good, the True and the Beautiful* (Cambridge: CUP, 2000).

⁷ See M. Slote, *Goods and Virtues* (Oxford: Clarendon Press, 1983), ch. 5, and R. Crisp 'The Dualism of Practical Reason', *Proceedings of the Aristotelian Society* 96 (1996), 53–73.

non-value-related way. One trouble with this procedure, however, is that the preponderance of value-linked reasons seems to be enormous, at the least, and we are therefore in danger of letting the tail wag the dog. Another trouble with it is that the connection between reasons-holism and value-holism seems so plausible. It is far easier to explain the holism of the few non-value-linked reasons in terms of their relation to the holism of the many value-linked ones than to cast off the only obvious prop we have.

I take it, then, that reasons-holism does not entail value-holism, since it is at least possible that reasons-holism is to be explained in other ways. By far the most plausible picture is that, just as most reasons are linked to values, so their holism is linked to a holism of values. This leaves us with a strong incentive to be particularists in the theory of value. Of course this incentive needs to be supported by examples of values varying according to context. But these are not too hard to produce. The disvalue of pain may be affected by the question whether it is part of a merited punishment. This does not mean that the pain hurts less, but that the punishment is not as much the worse for involving the infliction of the pain as we might have been led to think by considering the disvalue of other pains that hurt as much.

If value-holism shadows reasons-holism, the two views have to be structurally similar. Now as far as reasons go, the holism that appeals to me holds:

1. What is a reason in one situation may not be the same reason in another; it may even change its polarity.
2. The way in which the reasons here present combine with each other is not necessarily determinable in any simply additive way.

By analogy, then, our value-holism should look like this:

1. A feature or part may have one value in one context and a different or opposite value in another.
2. The value of a complex or whole is not necessarily identical with the sum of the values of its elements or parts.

And this is therefore the form of value-holism that I adopt. In this I differ from G. E. Moore, despite that fact that we could both be called 'organicism' in the theory of value. Moore believed (2) above but not (1). He held that any feature or element necessarily retained the same value as it moved from context to context, but that it could contribute to a complex of which

it was a part a value other than the one that it had there. The whole, that is, could be more valuable because of the presence of a certain part than could be explained by the value of that part; a part can contribute more, or less, value than it actually has.⁸ I don't believe any of this.

What explains the difference between Moore and me? This difference needs explaining, since all the examples that impressed Moore are just the sort of examples that impress me. How then have we come to such different conclusions? The answer lies in the fact that Moore accepted without question a certain doctrine of supervenience. He believed that the intrinsic value of something supervenes upon its other intrinsic qualities, and so that where an element does not change its intrinsic qualities on moving from one complex or whole to another, it must retain the same intrinsic value. The examples of organic wholes that impressed him required him therefore to say that elements can contribute to a complex more, or less, value than they have got themselves there. I, however, do not accept Moore's doctrine of supervenience. I accept (or used to accept, and as far as the present debate goes continue to accept) a slightly but crucially different supervenience doctrine, that intrinsic value supervenes upon other qualities, but not that the intrinsic value of one object supervenes upon *its* other qualities. For me, it supervenes upon other qualities, including those of other objects.⁹

This may seem perverse. But there is a ready explanation of it, once one remembers my commitment to particularism. First, I distinguish between those features from which some value *results* (the good-making features, as we might put it), and other features whose presence or absence would have made a difference. The latter features are obviously relevant to the value, but they are not playing the same role as that played by the good-making features; they are not themselves part of what we might call the 'resultance base'. Given this distinction between roles, I can announce that intrinsic value is value that results from intrinsic properties of the object concerned, but also allow that that value can vary because of changes elsewhere, that is, in those properties whose presence or absence can make a difference to the ability of the intrinsic properties to generate the value that they do. The notion of supervenience draws less fine distinctions than that of resultance,

⁸ See G. E. Moore, *Principia Ethica* (Cambridge: CUP, 1903), 30.

⁹ For an early proponent of this view of supervenience, see A. C. Ewing, *The Morality of Punishment* (London: Kegan Paul, 1929), 166: 'It does follow from the conception of goodness or value that the value of something cannot be different except as the result of some other difference, but this difference need not necessarily lie in the thing itself, it may lie in something else. We cannot therefore say that the intrinsic value of any quality will always be the same under all circumstances.'

and all we can say is that intrinsic value supervenes upon other properties, not particularly upon intrinsic properties. And once we have said that, there is no reason to stop short of allowing that the 'other properties' can be properties of other objects, or relations between them and the first one, or whatever.

My doctrine, then, is a sort of global supervenience, since the supervenience base is cast so wide, while Moore's is a sort of local supervenience. Who is right? I do not see that there is any way of determining the answer to this question before we get down to arguing about particular cases. It is not, that is, going to be a logical question which of us is right. Nor is it going to be decided quickly by appeal to the notion of a reason, or to naturalism in metaphysics, or anything like that.

So the overall situation, as far as value-holism goes, is that if values are to track reasons, and if the structure of value-holism is therefore to be the same as that of reasons-holism, we have to abandon one traditional formulation of the doctrine of supervenience in favour of something less familiar. But to appeal to the traditional doctrine to defeat my form of value-holism would be to beg the question.¹⁰

Holism in the Theory of Choice

Let us suppose now that I am right on both counts: my reasons-holism is the truth, and my value-holism is the truth as well. What does this tell us about the possibility of a full ordering, in which everything has its place, and where for each A and each B, A is either better than B, worse than B, or roughly as good as B? Well, so far as what we have so far seen will take us, there might still be a full ordering of that sort, in which everything has its place in the table of values, from best to worst. For all the considerations we have so far adduced concern the way in which the value of a complex is determined from the values of its parts, and the way in which parts may change their values as they move from one complex (context) to another. Once the value of the whole is determined, however, it is not going to vary, and we can enter the whole in its proper place in the great ordering. This would mean, for instance, that transitivity is not threatened by particularism. For with everything in its own place, we are never going to get a situation in which A is above B, B above C, and C above A in the Great Order.

¹⁰ It seems to me that Susan Hurley does make this appeal in her *Natural Reasons* (Oxford: OUP, 1989), 235–7.

This pleasing picture is however not as secure as it initially appears, sad to say. Let us remember that some of the things that have value are actions, and actions are chosen out of a set of alternatives available to the agent at the time. Now to adopt the picture I have just described is to suppose that each alternative action has its place in the great ordering of values, a place that is not affected by the place or nature of other alternatives. And this seems to require that the value of an action is never affected by the question what the available alternatives are. Now this is a very attractive doctrine indeed, partly because it enables us to retain a plausible principle of rational choice which we might call the principle of the indifference of independent alternatives (IIA):

IIA: If in one situation I prefer action A to action B, it can never be rational for me to prefer B to A in other situations which differ from the first merely in the fact that further alternatives are available.

In simpler terms, if I choose A where my choice is between A and B, I cannot rationally choose B where my choice is between A, B and C. The availability of C may indeed alter my overall choice, but it cannot affect the relative ranking order of A and B that has already been established.

Though this principle is *very* attractive, I am not convinced that stubborn adherence to it is fully compatible with the broadening particularist perspective that I have been developing. For it is not obvious to me how one can prevent available alternatives from counting as part of the context within which an action is placed. And if one cannot do this, then the general particularist claim that context can make a difference to value as well as to reasons seems likely to take us to the view that the value of an action or choice can be affected by the alternatives that are available at the time.

There is a reply to this, however. The argument of the previous paragraph might have been merely that every alternative is an object, though not all objects are alternatives (to each other). Since every object may have its value affected by others, every alternative may have its value affected by other objects, including some that are alternatives to it. There could be no bar against this happening—no bar that ensured that only those objects that are not alternatives to this one are capable of affecting its value. But this, though true, does nothing to establish the controversial doctrine that is really what we are after. That doctrine is that when one object becomes an alternative to another, that change may make a *further* difference to the value of the second—a difference beyond that

made by the existence and nature of the first object. And this doctrine does seem very peculiar.¹¹

But not as peculiar as all that, perhaps. There may be examples of this phenomenon—of the arrival of a new alternative making a difference to the relative values of two existing alternatives, in ways that are not explained merely by the *existence* or possibility of the thing that becomes an alternative, but rather by its new status as an alternative. Suppose that we have two alternatives A and B, and that we prefer B to A. Our original question was whether we might be rationally led to prefer A to B if there appears a further alternative, C.¹² Any successful example of this must meet certain criteria. It must not be one in which we simply change our minds about our initial ranking of A and B, perhaps for the reason that the appearance of C as an alternative draws our attention to something that we had previously missed. If nothing changes, however, it is hard to see how the ranking of A and B can be reversed by the arrival of C. The question to bear in mind is whether the examples offered contain the right or the wrong sort of change.

Suppose that I have to buy a house in Reading, and have a choice between a smaller house within walking distance of the university, and a larger and more expensive one that requires a bus ride. I prefer the larger one despite the bus ride. Then a third house, even larger but also further away than the second, comes onto the market. I realize that if I buy the second house, I will always regret not having bought the third. With this in mind, I buy the first house. Is this rational? I suggest that it can be.

There is of course a change of information here, but it is not the *sort* of change that I tried to rule out earlier—a change that leads to a change in my initial ranking. If there were no further house available, I would still have chosen the second. I have not changed my mind about *that*. So we are

¹¹ I have expressed this doctrine in terms of a *change*. But that need not be the point. The question could equally well be phrased in terms of the difference between the case where C is not an alternative and the case where it is; can the difference between C's being an alternative and its not being one make a difference to the relative values of A and B? Here there is no talk of change. There is, of course, nothing wrong with examples that do involve change. It is just that I should be careful to avoid supposing that change is essential to the point.

¹² It is not, of course, strictly necessary for us to find an example in which the order of the initial choice is reversed. It would be enough if we found a case in which the relative values are altered, so that the one that we originally preferred we still prefer, but not by so much—or by more, perhaps. Then we could argue that this sort of change in relative preference is bound to lead, on occasion, to a change in ordering. But it is more striking to produce an example in which the ordering is reversed.

not dealing here with a simple revision of the initial choice in the light of new information.

Have I cheated in the description of the example? The obvious mistake would be a slide in the nature of the 'objects' of choice. Let us be sure that these 'objects' are my buying house A, my buying house B and my buying house C. House C was there all the time, but it was not on the market. The mere possibility of a further house C that is even larger than B, though still affordable, is not enough to cause me to prefer the nearer house. There has actually to be such a house C available to me before my continual regret at not having bought it can turn into a reason for me to choose the nearer one.

The simplest example I know, and perhaps for that reason the best, is found in Wilfred Thesiger's *Arabian Sands*.¹³ He wrote, about his travels across the desert on a camel: 'I would not myself have wished to cross the Empty Quarter in a car. Luckily this was impossible when I did my journeys, for to have done the journey on a camel when I could have done it in a car would have turned the venture into a stunt.'

So it does look as if there may be examples in which independent alternatives are not indifferent.¹⁴ Maybe, then, even if particularism does commit us to the existence of such examples, this is not a disaster. But I raised a question earlier that I have not yet answered. Does particularism itself constitute a reason for rejecting IIA? My first attempt to show that it does was a failure. I argued that particularists should not be surprised to see the nature of one thing making a difference to the value of another. But this was irrelevant. The real question was whether the rather special feature of 'being an available alternative' can make a difference to the relative values of two things. Particularism does not show that this must happen. Particularism is everywhere permissive rather than prescriptive; or perhaps we should say that it forbids some things and prescribes nothing but suspicion. In the present case it says that we should be open to the possibility

of such a thing and not make a fuss if some crop up. The only reason for supposing that there *cannot* be any examples would be the generalist claim that since the feature of 'being an available alternative' often makes no difference, it makes no difference anywhere.¹⁵

There remains a difficulty. There appears to be an argument that there could be no counter-example. If so, the situation is unstable. We would have an example, our holistic position that there could be examples, and an argument that there can't be. Here is the argument. Sadly, it is one to which there can be no particularist objection *as such*; it is not a switching argument, for instance.

Let us start with a supposed overall ranking, the Great Ordering. Everything has its place in the order. We can compare the values of different objects, and compare different objects in order to establish their relative placing, and these activities cannot themselves alter the values of the objects compared or their relative places; otherwise the very notion of relative value would be incoherent. One can compare the values of merely possible objects, e.g. possible courses of action or possible states of affairs. Could there be any difference between the values of merely possible objects and the values of those objects should they become real? No; for otherwise the activity of establishing the relative values of different possibilities would be incoherent. And this would make deliberation before action incoherent, if deliberation is the establishing of relative values of possibilities so as to decide which to make actual. Suppose then that I ask you to rank ownership of each of ten paintings. What you are ranking is a set of possibilities. Suppose then that I give you all the paintings, and ask you to rank the ten actual ownings. There can be no conceivable reason for a change in your ranking order (unless of course you have changed your mind). Now: could there be a difference between an order of preference and a ranking order for

¹⁵ An interestingly different avenue of approach, which I will not pursue here, starts from something that Derek Parfit is apparently happy to admit, namely an analogous claim concerning not value but 'ought'. Parfit's view seems to be that it is possible that one ought to do A if B is the only alternative, but that if C is also available one ought to do B. He denies, however, the claim that the values of A, B and C can be related in a structurally similar way. See his *Reasons and Persons* (Oxford: OUP, 1984), 429: 'Whether I ought to act in one of two ways may depend on whether it would be possible for me to act in some third way. . . . I then ask whether, compared with A+, A would have been better. The relative goodness of these two outcomes cannot depend on whether a third outcome, that will never happen, might have happened.' So in the evaluative realm Parfit is what one might call a choice-atomist but in the deontic realm he is a choice-holist. My own view is that this position is unstable. And my reason is fairly predictable: that if one admits that there are examples of choice-holism in the deontic realm, structurally similar examples will emerge in the evaluative one.

¹³ London: Longmans (1959), 260.

¹⁴ I am grateful to Lars Bergstrom for very helpful discussion of this and other potential examples. He referred me to M. Resnik, *Choices* (University of Minnesota Press, 1993), 40, for another example with the same structure. Here is a further one, which I owe to Eve Garrard. I have to choose between two men, Joe and Sam. Joe is dull but reliable, Sam is unreliable but exciting. I prefer Sam, because if I chose Joe I would always be missing the excitement that Sam would have given me. But then Sebastian comes along, who is even more exciting but yet more unreliable. I realize that if I chose Sam, I would always be missing Sebastian's excitements, even though Sebastian's unreliability is so terminal that he is ineligible. And this tells me that my reason for preferring Sam to Joe is no reason, in the new situation, since choosing Sam will not lead to my having no lost excitement to regret. So I choose Joe.

choice, where what one is dealing with is alternatives? No: there is no possible relevant difference between a preference order and a choice order. Suppose that instead of giving you all the paintings, I give you the money to buy one. You should buy the one that came top in your preference ranking. So the feature of 'being an alternative' cannot make a difference.

Matters are more complicated than this. The complications do not make a difference, but they are relevant to what happens later. There is an obvious difference between buying just one painting and ranking them all. The difference comes out when we consider a case where I give you the money to buy one painting, then enough to buy another, and so on until you have all ten (though you never knew in advance that I would give you the money for the next). The order in which you buy the paintings need not be the same as your original order of preference. To see this point, it is important to distinguish between two quite different preference orders. The first has ten slots, in each of which one is asked to put one item of the form 'I own picture n'. The second has ten slots, the first of the form 'I own picture n', the second of the form 'I own pictures n and m', the third of the form 'I own pictures n, m and p', and so on. There is no reason whatever why either of these two orders should be extractable from the other. The point is that if you already have, say, six of the paintings, you might rationally choose to add to those six a painting other than the one that came seventh on the list. To get a true analogue of the original ordering, when it comes to choice, we have to suppose that I give you enough money to buy one, but that just as you try to buy it someone else gets in first; so you should go for the second on your original list—but the same happens again, and so on down to the tenth. The order of choice should be the order of preference.

What we have, then, is an explanation of why a certain feature, 'being an alternative' cannot make a difference, and therefore of why particularism is compatible with a full ordering. In one sense (epistemically, perhaps) it is *possible* that 'being an alternative' can make a difference. But there is an argument (which is not a switching argument) that no instance of this could be found. So the situation seems to be that we have on the one side an example in which the feature does make a difference, and a weak general reason derived from our holism to expect this sort of thing to crop up, and an argument on the other side to show that it is impossible. Now this is not one of those situations in which there can be reasons on one side and reasons on the other, and we can just decide where the balance of probability lies, leaving the defeated reasons in place. If we go with the example, we have to show what is wrong with the argument on the other side.

Luckily this can be done. The property of being an alternative is incapable of making a difference to a ranking order already established because there is no relevant difference between overall preference and overall choice. And just as the ranking preference order may be affected by the list of things to be ranked, so that if we take something off the list, the rankings of the rest may change, so with the ranking of alternatives. But all that this shows is that *preference is like choice*, and like choice in the crucial respect that it deals with alternatives. Being alternatives is the same as being mutually exclusive. Not all preference rankings are rankings of objects conceived as mutually exclusive, as we have seen. But some are. And the same is true whether we are ranking existing objects or possible ones. So the explanation of why the feature of being an alternative cannot be the cause of a difference between a preference ranking and a choice ranking is that this feature is present on both sides.¹⁶ Our conclusion should be that being an alternative can make a difference to all three rankings: of possible objects, of actual ones, and of objects of choice.

The existence of persuasive examples should then move us without further resistance from value-holism to a sort of choice-holism, which holds that:

1. The value of one alternative can be affected by the nature of other available alternatives.
2. Assessing the relative merits of different alternatives is not the same as assessing the various alternatives one by one and then comparing the results.

I want to end this section by comparing what I have said here with something I wrote in *Moral Reasons*:

My daughter trod on a sea-urchin on holiday a few years ago, and we caused her considerable pain (not entirely with her consent) in extracting the spines from her heel. Was the pain we caused her something which made our actions worse than they would otherwise have been? Here is a switching argument which says that it was. Had there been available a painless method of getting the spines out, we would and should have adopted it. We would have been wrong to continue digging in her heel with a needle, because of the pain. Surely this shows that as things were our actions were the worse for the pain they caused?

¹⁶ Thanks to Eve Garrard and David McNaughton here. It would be wrong to say, in reply to this argument, that the feature we were originally discussing was that of being an available alternative, not that of being an alternative. The notion of availability merely takes us from possible choices (preferences) between mutually exclusive options to actual choices.

I don't think it does show this. What we should say about cases like these is that a feature which would have made this sort of difference had there been any alternative choice need not necessarily make it if there is no alternative. It seems to me quite consistent to say that as things stood our action was not the worse for the pain it caused, though that pain should have led us to choose another method had one been available.¹⁷

The idea, expressed in terms of reasons rather than, as above, in terms of values, was that the pain is not a reason against the action if there was no alternative, pain-free course of action available. It is not just that it is not sufficient reason; it is not any reason at all. I presented this thought as an application of a style of switching argument, whose general form is: if this action were less F, it would be better; so its being F must detract from its overall value. But it can be seen immediately that the example I gave goes further than is required for that purpose. My use of an example that hangs on a point about alternatives was more of a distraction than a help, since the general point I was trying to make was nothing to do with alternatives. Talk about available alternatives was intended more as an explanation of the supposed fact that, in the example given, the action was not the worse for the pain caused, even if it would have been better with less pain. Since it was not possible to do it with less pain, the pain caused does not make the action worse than it would otherwise have been. In possible world terms: even though, in the nearest world in which there was an available pain-free alternative, the action we did was wrong, wrong because of the pain it caused, and the worse for that pain, the actual action is not the worse for that pain. *This is just an application of the holistic thought that a feature can make a difference in one situation that it does not make in another.* Where this occurs, particularists admit that there must be an explanation of it; the explanation is that in the actual world, there was no alternative.

Perhaps, then, the situation is like this. Holism takes away from us one of our two main reasons for sticking to principles like IIA. If one is a holist, it is going to be hard to think that the question whether something was a real alternative cannot make a difference. If IIA expresses a form of generalism, holists don't have *that* reason to believe it. They might have the other reason, which is that if we lose principles like IIA (and all the rest), we lose what is really the only detailed account of the 'logic' of choice. The loss of IIA seems to be another nail in the coffin of the idea that there is such a logic.

¹⁷ Dancy, *Moral Reasons*, 65–6.

One might suggest, however, that the loss of IIA is far worse for generalists than it is for particularists. For generalists, the pillars of practical rationality would really be tottering.

Choice-holism and the Possibility of a Full Ordering of Values

Suppose then that we accept a general choice-holism to go with our reasons-holism and our value-holism. Have we totally lost any possibility of a full ordering? There are two ways in which we could hope to retain anything like a full ordering.

The first is to say that what we have established is only that the context of real choice (i.e. the actually available alternatives) can affect the value of an option. This result, we might say, is clearly disturbing. But it does not altogether disturb our full ordering. Choice-holism concerns itself with real situations, in which the question what alternatives are available is a serious practical one. As such, it is to be distinguished from any thoughts about the effects of *merely comparing* one option with another. The value of an option will not vary according to what we imagine as the possible alternatives to it; it will only be affected by what actually are the alternatives. And mere comparison is relevantly similar to imaginary choosing, we might say, so long as the purpose of the comparison is to establish relative value. So an object's value is not affected by the mere act of mental comparison with another; it is only able to be affected if the two objects become real alternatives for some agent.

Even this, of course, will do something to upset our full ordering, if we suppose that the very same thing can occur in more than one actual choice situation. But we might deny that possibility, supposing instead that objects of choice are incapable of recurrence. We cannot have the same action again, that is, and we cannot have any other choosable object again either. For the objects of choice are not repeatables. If I offer you a chocolate bar today and you refuse it in favour of a pint of beer, and I offer the very same bar to you tomorrow, the fact that it is the same bar does not show that you have the same choice again. You only have a similar one, and holists allow that objects that are intrinsically similar may yet differ in value because of their context.

Be that as it may, the position that this move is trying to defend is surely another unstable one. It holds that mere comparison of two objects A and B, which we can do at any time at will, is incapable of affecting the values of A and of B; but should they become actual alternatives, their values may

be affected. So I may compare A and B and prefer B, and yet when I have to choose between the two choose A without irrationality. This distinction between actual and imagined choice, or between the effects of choice and those of comparison, is surely unsustainable. And this means that the dream of a full ordering collapses entirely. For if I cannot compare two objects without being in danger of affecting their relative values by doing so, there is surely no sense left in which objects have their own place in the ranking order. The ranking order must mean that objects have their place on it whether one actually compares them or not; indeed, to compare them is just to establish their relative placings in the order. If one could affect those placings by the act of comparison, the notion of an order would be destroyed.

So much for the first way of defending the possibility of a full ordering. The second way involves us in redescribing each option in terms of the available alternatives to it.¹⁸ Instead of thinking of ourselves as having the three options of buying house A, buying house B, and buying house C, our three options are:

1. Buying house A when we could have bought houses B or C.
2. Buying house B when we could have bought houses A or C.
3. Buying house C when we could have bought houses A or B.

Now, we might say, these three options have an unvarying value, and occupy a fixed place in the full ordering. For if one of the three houses is taken off the market, or a fourth house enters the equation, we no longer have any of these three options, but either two or four new ones.

The difficulty that I see in this approach is that if it is to avoid the difficulty we have already exposed, that comparison of the value of two objects is relevantly similar to choice, we will have to relativize every item on the full ordering to all other items, first severally and then in pairs and so on up until each is relativized to all others at once. And there will be no way of predicting, from the value of an option that is relativized to degree *n*, what its value will be relativized to other degrees. Given this, the use of the full ordering will be limited indeed. Transitivity, for instance, must fail. For if we rank 'A when we could have had B' above 'B when we could have had A', and 'B when we could have had C' above 'C when we could have had B', it in no way follows that we should rank 'A when we could have had C' above

¹⁸ I have borrowed this manoeuvre from John Broome, *Weighing Goods* (Oxford: Blackwell, 1991), ch. 5, and 'Can a Humean be Moderate?' in R. G. Frey and C. Morris (eds.), *Value, Welfare and Morality* (Cambridge: CUP, 1993), 51–73.

'C when we could have had A'.¹⁹ It will be perfectly true, that is, that every relativized option occupies one and only one place in the ordering, without this doing anything to preserve the conception of rational choice that the ideal of a full ordering was designed to promote.

There is a further problem. Suppose that we have a full ordering of all relativized options. This locates each option with respect to every other. Suppose now that I ask of item 32 in the list how it compares in value with item 33. It need not be the case that my answer is that item 32 is more valuable than item 33. The option '33 when I could have had 32' is a different option from the simple option '33', no matter how internally complex option 33 may be—and the same goes for option 32. But if my ranking order does not even commit me to claims about the relative values of the items ranked, it is pointless.

Holism and Explanation

There remains one further matter that I think it worth bringing out. It seems to me that particularism commits one to a highly debatable doctrine in the theory of explanation. In *Moral Reasons* I was not so clear about this.

I start by considering the relations between two doctrines. One is the now familiar holism in the theory of reasons. The other is a doctrine in the theory of explanation, which has no agreed name that I know of. Here they are:

Holism: the ability of a consideration to stand as a reason for action can be affected by the context in which it occurs.

Non-guaranteeing explanations: an explanation can be perfectly good without being 'complete', where a complete explanation is one that is inconsistent with the non-occurrence of the explanandum. Where E and O both occur, the occurrence of E can explain that of O without guaranteeing it. Perfect explanations can be 'non-complete'. The idea that non-complete explanations are enthymematic is a mistake.

¹⁹ Of course, relative to one and the same three-way comparison, transitivity must be preserved—or at least nothing that I have said gives us any reason to dispute that. If we do dispute it, we will probably do so for quite different reasons, i.e. those stemming from comparisons in which many different criteria are operating at once. For a recent rehearsal of such considerations, see Larry Temkin, 'Rethinking the Good' in J. Dancy (ed.), *Reading Parfit* (Oxford: Blackwell, 1997), 290–345.

The question that I want to start with is whether I was right or wrong to think, as I once did, that my doctrine of reasons-holism was effectively the same as this doctrine in the theory of explanation. If it is, we have uncovered another awkward consequence of particularism.

One preliminary first. Some philosophers maintain that though guaranteeing explanations are not required when what we are explaining is an action, they are required when we are explaining events. A scientific explanation, then, will have to be a guaranteeing one.²⁰ My own view, however, is that we would need a lot of persuasion to say any such thing. The only reason for adopting it, I think, is a sense that we would do well not to demand something of action-explanations that we have no prospect of achieving: in science, by contrast, where prospects are better, we should not be satisfied with anything less than perfection. The proper riposte to this is that action-explanations are as good as any explanation needs to be, and that, so far as the purposes of explanation are concerned, there is no reason to think there is either need or room for anything better. We should therefore accept the possibility of non-guaranteeing explanations on both sides or on neither.²¹

The concept of a non-guaranteeing explanation requires that of an enabling condition. If there is a non-guaranteeing explanation F of an event E, there must be an event O such that:

1. The occurrence of O is not part of the explanation of E, and
2. If O were to fail, we would have a situation (F, not-E).

In such a situation, O would be an enabling condition for the explanation that F gives us of E. The reason why there must be such things as enabling conditions if there are to be non-guaranteeing explanations is that if every candidate for the role of enabling conditions were to turn out to be part of the explanation of the event-type E, all explanations would be guaranteeing ones (when 'complete'). If we think that there are *any* non-guaranteeing explanations, then, irrespective of whether we think that no explanations are guaranteeers, we must be able to make sense of the notion of an enabling condition—a condition whose satisfaction is required for the explanation, but which is not itself a part of that explanation. And this sets us something of a challenge.

²⁰ For a recent example of this view, see Philippa Foot, 'Does Moral Subjectivism Rest on a Mistake?', *Oxford Journal of Legal Studies*, 15, 1 (1995), 1–14.

²¹ My views on the nature of explanation in general bear interesting similarities to those of Nancy Cartwright. See her *How the Laws of Physics Lie* (Oxford: Clarendon Press, 1983).

But we face the same challenge in the theory of reasons for action, once we adopt holism there. This is not perhaps too surprising antecedently, since presumably the reasons favouring an action are reasons that explain that action's rightness (or whatever moral status it has). If we start from one case where there are reasons R1–Rn making the action right, but allow that changes elsewhere may affect the ability of those reasons, not merely to make the action right, but even to be reasons at all, we will again say that features over and above the reasons must be functioning as enabling conditions. They enable the features that are reasons to be the reasons they are in this case, without themselves being among the reasons why the action is right.

To give a very simple example: suppose that ought implies can. Then if I cannot do the action, the features which, were I capable of doing it, would be reasons why I should do it, are incapable of playing that role. But we should not conclude that my ability to do it is one of the reasons why I should do it. It is a condition that enables the reasons why I should to be the reasons they are, but is not itself among those reasons. Allow this, and you will probably allow the next: that I have the opportunity to help is something without which the reasons why I should help would not be those reasons. Her need is not a reason for me to help her if I have no possible opportunity of doing so. It is only a reason for me to seek an opportunity, which is different. But that I have an opportunity to help is not itself among the reasons for doing so. Another similar example: if I were not alive, the reasons that there are for me to help the needy would not be able to be the reasons they are. But this does little to show that among the reasons why I should help the needy is the fact that I am alive.

So we see the same structure both times. Both holism and the claim that explanations need not guarantee are committed to making sense of the notion of an enabling condition. One might think²² that there are forms of holism that don't have this effect. We might try to adopt a sort of weak holism without adopting any distinctive doctrine in the theory of explanation.

Weak holism: the ability of a consideration to stand as a reason can be affected by what other features are present as reasons (but not by other things).

With this in hand, we might hope to avoid the need to talk about enabling conditions at all, and avoid the need to allow non-guaranteeing

²² And I did think so, until Eve Garrard showed me that I shouldn't.

explanations. But we can only do this if we can add to our weak holism a sort of:

Holistic generalism: taking all the reasons here present together, they guarantee the rightness of the action; i.e. where present together elsewhere, they will always have the same effect (non-causally, of course).

This holistic doctrine is certainly compatible with the denial of a different form of generalism:

Atomistic generalism: each reason has the same tendency as a reason, no matter what the context in which it is found.

But in fact weak holism cannot be coherently combined with holistic generalism. If the ability of a consideration to stand as a reason can be affected by what other features are present as reasons, why is it that the presence of further reasons in a second case, in addition to all those present in the first, is incapable of making any difference to the original reasons? Surely the official statement of weak holism says that a new reason can upset others whether we take them one by one or all together. So the combination of weak holism and holistic generalism is not a sustainable position.

Moving now to the theory of explanation, we find the same thing. There might seem to be a form of holistic generalism available there, that is compatible with the demand that all explanations be guaranteeing explanations. We might say, that is, that though no individual feature has its own explanatory potency, which it carries with it from case to case, regardless of changes in other explanatorily potent features, still the entire complex of features that together explain the event constitutes a guaranteeing explanation. It is a guaranteeing explanation because that complex of features could not have occurred unless the event to be explained had occurred (or been going to occur).

But this attempt to make room for guaranteeing explanations within a holistic picture suffers from the same incoherence that we found in the theory of reasons. No explanation is given of why the arrival on the scene of a new element is capable of affecting the explanatory contributions of individual features, but not of a set of such features. Why is it that a whole set of explanatory features is necessarily invulnerable to the sort of difference that a new feature can make to individual elements in that set? There is just no answer to this question. So there is no way for a holist to avoid making

room for the concept of enabling conditions, and denying the possibility of guaranteeing explanations.²³

One useful consequence of these thoughts is that the incoherence of the combination of weak holism and holistic generalism gives us an answer to one leading reply to particularism.²⁴ This reply is that my examples only concern what we might call *contributing* reasons. It is true that contributing reasons do vary in their polarity from one context to another. But it is not true of *complete* reasons. These remain the same, as reasons, regardless of changes elsewhere.

There are two ways in which I can respond to this. The first is, as above, to point out that there is no obvious explanation of why a new reason in a new case is supposedly able to change the behaviour of one reason but not of all of them at once. Why is it that the whole pack of them is immune to change, when no individual one is? Note that this question is asked within the constraints of weak holism; it would be a different matter to maintain that considerations that are not reasons in the new case are able to make changes in this way.

The second way in which I can respond to this attack is to challenge its notion of a complete reason. This cannot be identical with all the reasons present in the particular case, since we have already seen that there is no justification for the view that the pack of all such reasons is invulnerable to changes brought about by the presence of a new reason in a new case. It must therefore be something greater than that; it will presumably contain all the enabling conditions, as well as the absence of disabling conditions (specified one by one). But, first, this collection is ceasing to look like a reason at all, and, second, there is beginning to be a prospect that a complete reason will expand indefinitely.²⁵

²³ The same may be true about causal statements: a sufficient cause need not be a guaranteeing cause. But to try to argue this would take us too far away from present concerns.

²⁴ I mentioned this reply earlier, in discussing holism in theoretical reason. A recent example is to be found in Jonathan Bennett, *The Act Itself* (Oxford: Clarendon Press, 1995), 80.

²⁵ Jonathan Lowe pointed out to me that a valid deductive argument is something we could reasonably call a complete reason, since it is monotonic, i.e. no addition of premises will affect the validity of the argument. In this, deductive reasoning differs from probabilistic reasoning. So there is one model of a complete reason that I cannot undermine. This model is, however, not applicable to moral reasoning, unless there are absolute moral principles; reasoning that runs in terms of *prima facie* principles is non-monotonic. There is certainly available a notion of deductive reasoning in ethics. The only question is whether the reasoning that takes us from premises specifying the features that make an action right to the conclusion that it is right is ever deductively valid. (I mean by this form of words to exclude such premises as 'all actions of this sort are right', 'this feature is the only relevant one' and 'this feature is a *pro tanto* reason'.) I think not.

What I really want to say, of course, is just that my interests are in the ordinary notion of a contributing reason, not in the concocted notion of a complete reason, which in my view is designed merely to save a dubious philosophical theory from refutation. But this may seem to be a definitional sulk; so I support it by attacking the notion of a complete reason that is designed to save generalism from counter-example.