

PEKKA VÄYRYNEN

PARTICULARISM AND DEFAULT REASONS

Accepted 19 December 2003

ABSTRACT. This paper addresses a recent suggestion that moral particularists can extend their view to countenance default reasons (at a first stab, reasons that are pro tanto unless undermined) by relying on certain background expectations of normality. I first argue that normality must be understood non-extensionally. Thus if default reasons rest on normality claims, those claims won't bestow upon default reasons any definite degree of extensional generality. Their generality depends rather on the contingent distributional aspects of the world, which no theory of reasons should purport to settle. Appeals to default reasons cannot therefore uniquely support particularism. But this argument also implies that if moral generalism entailed that moral reasons by necessity have invariant valence (in the natural extensional sense), it would be a non-starter. Since generalism is not a non-starter, my argument forces us to rethink the parameters of the generalism-particularism debate. Here I propose to clarify the debate by focusing on its *modal* rather than extensional aspects. In closing, I outline the sort of generalism that I think is motivated by my discussion, and then articulate some worries this view raises about the theoretical usefulness of the label 'default reason'.

KEY WORDS: default reasons, generics, moral generalism, moral particularism, moral principles, normality

1. INTRODUCTION

One feature of ordinary moral experience is that many particular facts immediately elicit a preliminary moral opinion that often remains stable on reflection. Examples would be regarding it as a moral reason for some response or other – such as action, attitude, or emotion – that someone is in danger, that this action would get a person something she needs, or that it would constitute lying. Insofar as our stable intuitive response is to treat such facts as moral reasons, our response would seem to treat many reasons as having a fair degree of generality. This much of ordinary moral thought might be taken to support generalist theories of moral reasons,

¹I borrow the phrase “a suitable provision of principles” from Dancy (2001). I take it to pick out an appropriately comprehensive set of substantive moral principles such that the types of considerations those principles identify as reasons are reasons (when they are) somehow in virtue of falling under those principles. But I cannot develop this interpretation here.



according to which moral reasons depend on some suitable provision of general moral principles that identify reasons.¹ The stability of our response, and the fact that we find its stability intelligible, suggest that it isn't arbitrary. Showing that it has a principled basis would be an excellent way of showing that it isn't arbitrary.

But is it the only way? Moral particularists claim that what considerations function as moral reasons depends fundamentally only on the nature of the particular case. They are notoriously keen to point out difficulties in modeling the way that constellations of particular facts which we regard as morally significant (say, that Jill is freezing or that Jack had his shoes stolen) might be captured in terms of general moral principles (say, a principle about helping those in need). Why not contrast arbitrary responses simply with those supported by reasons, and give a particularist account of our response? Most particularists grant that many reasons have a fair degree of generality, and a particularist account of our response should seek to explain how that can be so.²

I want to explore the notion of a *default* or *presumptive* reason, which particularists might invoke to provide such an account. Such a reason is usually characterized as – and this is how I'll use the term until further notice – a consideration which is a reason unless something prevents it from being so (Dancy, 1999, p. 27), or which is a pro tanto reason unless undermined (Cullity, 2002, p. 188). To unpack this, let's think of a reason (for an agent *S*) to ϕ as a consideration that *favors* (*S*'s) ϕ -ing. And let's focus on those statements of reasons which express non-moral – and perhaps broadly empirical – facts or considerations, such as that *someone is in distress* or that *I promised to do it*. Then the initial notion introduced here is that of a non-moral consideration that carries a defeasible presumption that it favors bringing about certain things. (As we'll see, however, people who talk about default reasons often have in mind something richer.) Such reasons would naturally seem to invite from us a stable response that assigns to them some significant degree of generality. Thus the appeal to default reasons might help the particularist to bolster her case for the descriptive adequacy of her view.

My main aim in this paper is negative: I'll argue that if default reasons are to be at all relevant to the generalism-particularism debate, they must be construed as lacking precise implications concerning their extensional gen-

²To take just one example, Jonathan Dancy takes it as more or less a datum that some considerations are "more commonly important than others; and therefore some properties should be viewed as more 'central' than others" (1993, p. 67). See also Little (2000, p. 294ff.).

erality, by which I mean the ratio of cases in which reasons are undermined relative to those in which they aren't (rather than the generality or specificity of the reasons themselves or their the statements). As I'll show, their extensional generality depends on the sorts of contingent worldly matters which cannot *ground* moral reasons. Default reasons cannot therefore be recruited to provide any *further* support for particularism. (In particular, they aren't responsible for the stability of our responses noted at the outset.)

This result has wider ramifications. If the extensional generality of reasons depends on contingent worldly matters, then generalism is pretty much a non-starter if we take *it* to have any fairly precise extensional implications. This is what happens if we take it as the view that at least all basic or non-derivative reasons are *invariant*, for on the natural reading invariance is an extensional notion. Since I'm a generalist, I must reject this standard take. So I have a secondary, positive aim: I'll argue that the debate is better construed in terms of its *modal* than its extensional aspects. I want to explain how the sort of generalism we may expect to get as a result not only merits attention on its own, but also improves our grip on the notion of a default reason – even as it turns out to question the theoretical usefulness of the *label* 'default reason'.

As a final preliminary, I want to note that modal concerns are not all there is to the generalism-particularism debate. In fact, I believe that we must also consider its explanatory aspects and more robustly metaphysical aspects, which cannot be captured in purely modal terms. But I lack the space to address these issues here.³

2. DEFAULT REASONS AND NORMALITY CLAIMS

Let a *practical norm* specify a relationship of favoring between a non-moral consideration and a response. Candidates for genuine practical norms include (1)–(5):⁴

³But see my "A Theory of Hedged Moral Principles" (unpublished).

⁴I interpret the downward arrow as representing the favoring relation characteristic of reasons. The examples and the notation are drawn from Cullity (2002, pp. 184–185). I should add that Cullity's formulations are almost certainly inadequate (although I believe they are good enough for the present purposes). They omit time and agent specifications, and perhaps they are best regarded as (enthymematic) instantiated forms of the relevant norms. It is also not clear whether ' φ ' is supposed to stand for an act-token or an act-type in both of its occurrences in a norm-formulation. Resolving these and related issues would require an excursion into action theory and philosophy of language that isn't possible here. Thanks to Vivienne Brown for pressing these concerns.

- (1) φ -ing would be enjoyable
 ↓
 φ -ing
- (2) φ -ing would get her X
 she needs X
 ↓
 φ -ing
- (3) φ -ing would be forcing him against his will
 ↓
 not φ -ing
- (4) φ -ing would be lying
 ↓
 not φ -ing
- (5) I have promised to φ
 ↓
 φ -ing

In a recent paper, Garrett Cullity (2002) argues that a particularist can treat the reasons set out in practical norms like (1)–(5) as *holistic* reasons whose moral valence may vary, depending on the circumstances, but which also are *pro tanto* unless undermined.

Put roughly, the holistic claim is that reasons can vary from one context to another. That φ -ing would be enjoyable is a reason to φ , but it might be no reason at all when φ -ing would be enjoyment of something bad.⁵ Similarly, when a person needs something for a malicious enterprise, its maliciousness means that (2) doesn't give a reason to give it to her. These reasons can be undermined by intervening considerations, such as the maliciousness of an enterprise, which defeat the favoring relation that a practical norm sets out between a reason and that for which it is a reason. The possibility of being undermined accounts for the holistic nature of the relevant reasons.

What about the idea that the considerations set out in (1)–(5) are *pro tanto* reasons unless something prevents them from being so? There are different options here. One might think the explanation lies in the connec-

⁵Or, perhaps, enjoyment that depends on the badness of what it is an enjoyment of.

tion between those considerations and ‘thick’ properties such as kindness and respect. Cullity’s account turns on their association with a good person’s normative orientations: each gives us “the core of a virtue” (2002, p. 187). For example, sometimes we *ought* to coerce people to abandon evil ends. But in these cases the normative orientation of coercion isn’t contrary to (3), which is a norm that a respectful person should recognize. What is contrary to (3) is being the sort of person who enjoys domination, sees other people as mere means to his ends, and the like. This association with virtue is why the relevant reasons aren’t merely *prima facie*, and why being properly responsive to reasons requires recognizing that considerations set out in (1)–(5) *are* reasons unless undermined. (Cullity, 2002, p. 188.)

The usual generalist response (but not one that I generally endorse – see §5) is to argue that (1)–(5) cannot be correct practical norms if the reasons they set out can be undermined: this only shows that our real reasons are more complex and have yet to be specified in sufficient detail.⁶ So, why think that default reasons can have the simple, unqualified contents listed in (1)–(5)? Cullity argues as follows. First, we must distinguish reasons from conditions for the presence of reasons (cf. Dancy, 2000, p. 153). When (say) the fact that ϕ -ing would be enjoyable gives me a reason to ϕ , but there wouldn’t have been a reason for me to ϕ if it had been enjoyment of something morally bad, it doesn’t follow that part of the reason for me to ϕ must be that ϕ -ing would be morally unobjectionable (cf. Cullity, 2002, p. 180). Second, given that reasons are considerations that motivate us insofar as we are rational and aware of them and the other relevant facts,⁷ a reason for me to ϕ must be something I could invoke in explaining my ϕ -ing to myself, insofar as I am rational and factually aware (Cullity, 2002, p. 179). But, third, it is in general fallacious to reason that if A wouldn’t have happened in the presence of B , the absence of B should figure in a rational explanation of why A happened. Rather, the explanation it is rational to give of any phenomenon is relative to “background expectations of normality” (Cullity, 2002, p. 178). So (concludes Cullity), the expectations it is rational for an agent to have will constrain the content of her normative reasons.⁸

⁶See e.g. Shafer-Landau (1997), Sinnott-Armstrong (1999), Raz (1999), and Hooker (2000a).

⁷This is close to what Korsgaard (1986) calls “the internalism requirement”.

⁸There are questions here that Cullity ignores. For example: to what extent does being rational involve being responsive to reasons? Whether it would be rational for S to ϕ may well be relative to S ’s situation in ways that whether something is a reason for S to ϕ isn’t. I’ll set these questions aside.

To illustrate, suppose I'm having some fun in circumstances C and this raises no questions of permissibility for me. If it is a rational expectation about C that having fun in C is normally morally unobjectionable, then its being morally unobjectionable will not normally be part of a rational explanation to myself of what I'm doing, and hence not part of what motivates me insofar as I'm rational (Cullity, 2002, p. 179). Then the rational self-explanation of what I'm doing can feature simply the fact that this is enjoyable. Since the absence of factors that would undermine the status of this fact as a reason needn't be part of my reason, my reason for doing what I do can be simply the consideration that doing it is enjoyable.

This line of argument has its attractions. But it is available to anyone who distinguishes reasons from the conditions for their presence (and accepts Cullity's premise about the nature of rational explanation), and later I'll suggest that the distinction is available to certain forms of generalism (see §6). What specifically interests me here is not so much the conclusion of the argument as the relation that the argument forges between default reasons and background expectations of normality. For it commits Cullity to claiming that default reasons entail corresponding *normality claims* about reasons: if something is a default reason to φ , then it is *normally* a reason to φ .

The argument for this commitment is simple: Any default reason to φ is non-undermined only when rational background expectations of normality don't raise questions about the permissibility of φ -ing. Circumstances where they don't are 'straightforward' or 'uncomplicated' (Cullity, 2002, pp. 179, 180). If (as Cullity also claims) circumstances are normal insofar as they are straightforward (2002, p. 186), then, for any default reason p to φ , the fact that p is normally a reason to φ .⁹

In what follows I'll first consider the notion of normality in a bit more detail and then, drawing on those considerations, argue that the notion of a default reason doesn't as such imply that default reasons are normally reasons. As we'll see, the reasons why this is so have broader implications for the generalism-particularism debate.

3. NORMALITY: A PRIMER

Normality is usually characterized as a complex relational characteristic of *worlds*, dependent on their *general shape*: a world w is normal relative

⁹This is clear in the special case of a virtuous person: since her background expectations will reflect the fact that her attitudes/enterprises have no bad content, there is nothing in them to undermine her reasons.

to w only if w' shares the fundamental laws of w and exhibits contingent initial conditions that aren't too bizarre from the point of view of w .¹⁰ Let p be 'This is a dry piece of wood'. Heat the piece in a normal world and it will catch fire, if it is combustible there, as it is in the actual world. But when w is a p -world, it needn't be p -normal, that is, normal relative to itself as regards its being a p -world.¹¹ It may be p -abnormal, in virtue of some specific condition or other. In the actual world, a given dry piece of wood that is combustible may not burn even though thoroughly heated. Perhaps there isn't enough oxygen. So, although whether a given world is normal depends on what regularities hold there, a world may not always be normal in its own terms with respect to its specific conditions: things may not, on occasion, behave there as they tend to behave there.¹²

The idea that what is normal depends on the general shape of the world is reflected in our background assumptions. I expect that if I thoroughly heat a dry piece of wood, it will catch fire. Similarly, when it perceptually seems to me that p , I tend to judge that p . Built into this tendency is the expectation that my perceptual conditions lack perturbing factors, and therefore don't undermine the conclusion that as things seem so they are (and as they are so they seem). These are the sort of perceptual conditions that we think of as normal. So our background expectations of normality commit us to substantive assumptions about our environment.

There is a further respect in which normality is a relational notion: a world may be normal or abnormal, in virtue of some condition or other, in one respect but not others. Tweety, a bird, may be normal for a bird regarding being feathered but not regarding flying. Tweety will be abnor-

¹⁰Cf. Asher and Morreau (1995, p. 314), Pelletier and Asher (1997, p. 1164), and Priest (1999, p. 202). The target notion so characterized in Morreau (1997, p. 200ff.) is *regularity*. But that term is often used, in the sense that is relevant here, interchangeably with 'normality,' 'typicality,' etc. (Pelletier and Asher, 1997, p. 1129).

¹¹Therefore, as Bonevac (1998, p. 42) notes, p -normal worlds relative to w must not be construed as the closest p -worlds relative to w .

¹²In other contexts, we would need further distinctions. For example, a p -world is *unnecessarily abnormal* regarding p if there is a normal alternative to it – if there is another p -world that is normal (see Morreau, 1997, p. 202). That would be a world which shares the general shape of our world and in which the piece of wood does catch fire. Similarly, a p -world is *necessarily abnormal* regarding p if there is no other such p -world. Such a world would be one in which dry pieces of wood aren't combustible in virtue of its general shape.

mal for a bird regarding flying if he is a penguin.¹³ In the actual run of things, both (6) and (7) are true:

(6) Birds fly.

(7) Birds that are penguins don't fly.

Sentences such as (6) have some notable characteristics. First, they are *modal* statements: their truth values do not depend (just) on actual matters of fact. Second, they are *generic* statements: they report a general property of individual instances of a kind, which transcends particular facts.¹⁴ (6) makes a claim not about the closed class of existing birds, but about (some set of) (realistically) possible birds. Third, generics *tolerate exceptions*: (6) can be true even if not all birds fly. In general, the truth of a generic claim allows that not all instances of the kind have the property predicated of them.

It is attractive to explain the exception-tolerating character of generics by saying that generics imply certain normality claims. Evaluated relative to w , (6) predicates the property of flying only of those birds which are found in worlds that are normal with respect to w and which aren't abnormal with respect to the property of flying, but instead fly in the appropriate circumstances.¹⁵ Many exceptions (such as penguins) can themselves be expressed by true generics – witness (7).

This sort of relation between generics and normality would help to explain why generics support inferences such as this:

¹³These claims are to be understood as implicitly relative also to the actual world. There should be possible worlds where Tweety is a perfectly normal bird and yet cannot fly, these simply being possible worlds where it isn't true that birds can normally fly (Asher and Morreau, 1995, p. 312).

¹⁴I should note that this characterization can be denied in the case of kind-referring uses of *definite descriptions*, such as an utterance of the generic 'The tiger has stripes.' The referent of this sentence could be claimed to be a unique contextually salient tiger. Further, just one type of generics is characterized here. Another type of generic statement reports a property of the kind itself; consider 'The potato was first cultivated in South America.' I'll set this sort of generic statements aside in this paper. On the entire spectrum of generics, see e.g. the papers in Carlson and Pelletier (ed.) (1995).

¹⁵As Pelletier and Asher (1997, p. 1162) note, the relevant set of normal worlds won't be fixed independently of the topic of the generalization: 'is a bird' will determine, presumably by the nature of birds, certain respects (such as flying and being feathered) as the relevant ones.

- (8) Tweety is a bird.
 Birds fly.
 Therefore, Tweety flies.¹⁶

The status of (8) as a reasonable inference depends on the substantive assumptions about the general shape of our environment to which accepting (6) commits us, and thus on assumptions about what counts as normal. Of course, we might reason as we do in (8) and yet, given evidence that Tweety is a penguin, retract the conclusion that Tweety flies – but without retracting any of the original premises or without saying that the original inference was unreasonable.¹⁷ Our background expectations of normality reflect this as well. To assume that Twyla is normal for a bird is to assume, roughly, that Twyla is as normal in certain respects as is consistent with Twyla's being a bird. If Twyla is a penguin, and birds that are penguins don't fly, we cannot consistently assume that Twyla is normal regarding flying, and we'll retract the conclusion that Twyla flies if we realize she is a penguin.

4. REASONS, NORMALITY, AND NASTY WORLDS

Normality is a rich notion, then. But what I have said gives us enough to go back and ask whether default reasons entail corresponding normality claims. I'll distinguish a few construals of such claims, and argue that on none of them does being a default reason entail normally being a reason.

¹⁶Implicit here is the alternative option of thinking that generics such as *birds fly* specify *default rules of inference* (Reiter, 1980). The problem with this is that rules of inference cannot be true or false; they only transmit truth, more or less intact, from the premises to the conclusion. But *birds fly* seems *true*.

¹⁷Thus reasoning involving generics is *non-monotonic*. A bit more precisely, an inference is non-monotonic if the set of premises Γ generates conclusion ϕ but the premises $(\Gamma \cup \psi)$, for some ψ not in Γ , don't generate conclusion ϕ . The contrast is with monotonic logics, such as the standard first-order logic, in which an argument that is valid *stays* valid no matter what or how many premises are added. Priest (1999, p. 200) suggests that non-monotonically valid inferences are truth-preserving in all 'normal situations.' On this proposal, ϕ must hold in all models of Γ that are as normal as Γ will let them be. Throw in ψ , and ϕ doesn't necessarily hold in all models of $(\Gamma \cup \psi)$ that are as normal as $(\Gamma \cup \psi)$ will let them be. The proposal handles (8) nicely, and something like it lies at the core of many current accounts of non-monotonic inference. For a proposal along similar model-theoretic lines which is more ecumenical regarding the selection of presupposed models, see Morgan (2000, pp. 347–352).

Let's begin with Cullity's claim that default reasons are "sufficient for the existence of a reason for action in most circumstances, where the background is uncomplicated" (2002, p. 180). This claim about the content of background expectations of normality could be read either (i) as stating what happens in most (but perhaps not all) straightforward circumstances, or (ii) as stating what most circumstances are like, namely straightforward. The latter reading is clearly the intended one. (However, it turns out not to matter for my argument which one we pick.) But does the notion of a default reason entail that most circumstances *are* straightforward?

Let's return to generic statements. Suppose we read (5) as saying that promising to φ is a default reason to φ . Then it at least entails the generic statement (9):

- (9) Promises give reasons to keep them.

As a generic claim, (9) tolerates exceptions, as it should if the reasons it specifies can be undermined. A parallel treatment with (6) would model this by saying that (9) predicates the property of being reason-giving of (to put it roughly) those promises that aren't abnormal with respect to *that* property. (9) is meant to be compatible with claims that set out tolerable exceptions to it, such as (10):

- (10) Manipulated promises give no reasons to keep them.

The compatibility is desirable because (arguably) if I'm tricked or coerced into making a promise, this undermines the status of *I promised to φ* as a reason for me to φ .

But does (9) entail that most cases of promising are straightforward? It does so only if in most cases the status of *I promised to φ* as a reason is *not* undermined. But compare (9) and (6). The truth of (6) doesn't require, on the extensional level, that most birds fly. (6) could have been true even if it had so happened that most birds were penguins, emus, kiwis, ostriches, etc., or had clipped wings. Also consider (11):

- (11) Turtles die young.

Most turtles die in infancy as they rush to safety in the sea. But although (11) is therefore true if prefixed with 'most' or 'in a significant number of cases,' it is false if read generically. What we regard as true is something like (12):

- (12) Turtles live to a grand old age.

The reason for this assignment of truth values is that in normal circumstances turtles are long-lived, presumably because of the nature of the species. Most turtles in fact have their lives cut short, but only because of interferences that, so far as being a turtle goes, are abnormal with respect to the property of being long-lived.

Examples like these support what is a well-worn point in the literature: it is neither necessary nor sufficient for the truth of a generic '*Fs are Gs*' that most *Fs* are *Gs* (see e.g. Pelletier and Asher, 1997, p. 1132; Drewery, 2000, p. 6). It suggests that no purely extensional analysis of generics can be correct and, by implication, that 'normal' must be given some non-extensional meaning. But then (9) does *not* imply that most promises give reasons to keep them. The actual extension of promises that give reasons has almost nothing to do with whether (9) is true or not. (9) allows that *most* promises may be manipulated, that promises *usually* don't give reasons to keep them, and so on. Nothing in the notion of a default reason implies that complicated cases are rare. Thus we cannot assume that if default reasons entail normality claims, those claims bestow upon default reasons any relatively definite degree of extensional generality. Therefore we must reject Cullity's construal of normality claims.

A weaker option is that rational background expectations assume simply that circumstances *normally* are straightforward, so that they *normally* raise no questions about the permissibility of our actions. Suppose *p* is a default reason to φ . Then, on this view, *p* is normally a reason to φ only if such expectations *are* rational. Now they may in fact be rational. But whether they are rational depends on a variety of additional factors. What I deny is that *p*'s being a default reason *as such* makes it rational to assume that cases where *p* holds are normally straightforward.

We can assume that circumstances are normally straightforward only if we can also assume it to be part of the general shape of the world that promises aren't manipulated, enjoyment isn't enjoyment of something bad, and so on. No such assumption, however warranted, follows simply from the notion of a consideration that is a reason unless undermined. For the possibility of such reasons doesn't show that we couldn't have been so unfortunate as to live in what I'll call a Nasty World.

A Nasty World is a world where – whether by nature, nurture, or both – human interaction is coercive and fraudulent and human psychology sadistic, deceitful, and perverse. The endeavors of the denizens of such a world would provide no rational basis for a general expectation that promises, enjoyments, and other considerations *are* normal with respect to reason-giving, since they would *normally be undermined*, and so would normally not give reasons. Manipulative situations, for example, would

be in the relevant sense normal with respect to promise-making – these are the sorts of situations in which promisors in a Nasty World tend to find themselves, thanks to the general shape of their world.¹⁸ Since manipulated promises, at least *qua* promises, are abnormal with respect to reason-giving, it follows that in a Nasty World the fact *I promised to φ* wouldn't normally give me a reason. So, again, being a default reason doesn't entail normally being a reason.

We are now in a position to draw some morals. One is that, insofar as promises are normally manipulated (and, perhaps, we have some awareness of this sorry fact), it will be no part of rational background expectations of normality that the circumstances of promising are straightforward. Then it won't do to argue that the reasons had by the denizens of Nasty Worlds can have the simple contents set out in (1)–(5) by appealing to background expectations of normality. But, intuitively, a relatively simple consideration *can* be a reason unless undermined even when contingencies make it rational to expect that it normally is undermined.¹⁹ The fact *I promised to φ* looks to be capable of functioning as a reason for the denizens of Nasty Worlds to do as promised, even as their normal promise-making situations would be such as to normally undermine the reason-giving status of promises.²⁰ Intuitively, such a fact can function as a reason so long as no undermining conditions obtain. So, considerations such as *I promised to φ* can be, relative to *w*, reasons unless undermined even if promises made in situations that in *w* are normal for promise-making don't in *w* normally give reasons to keep them. These claims cannot be vindicated by appealing to normality claims.

Another moral we can draw is that there is an important respect in which it would be misleading to say that, relative to a Nasty World, making a promise gives *by default* a reason to do as promised, even if promises do

¹⁸As noted in §3, generics appeal to normality conditions for both situations and objects (here, promises). See also Krifka *et al.* (1995, p. 31ff.).

¹⁹Perhaps when rational agents in Nasty Worlds fail to respond differentially depending on whether a reason is undermined, this is a mistake *within* rationality, insofar as one needn't be irrational for failing to know when a reason isn't undermined in circumstances where it normally is undermined.

²⁰To test your intuitions about this, consider (à la Cullity) what response a good person's normative orientations would yield in such cases. It seems clear to me that an appropriately informed good person would take her having made a promise as a reason to keep it. The separate question why people would bother to make promises at all in Nasty Worlds, though perhaps practically pressing, is tangential here.

so in the actual world. If promises normally are manipulated in Nasty Worlds, this is probably due to their having contingent initial conditions that are sufficiently far-fetched from our point of view. In general, whether promises normally are manipulated is contingent on the general shape of the world (which is itself a contingent matter). *If* by a default assumption we mean one that automatically entitles us to proceed in a certain way, absent interfering factors, then promises fail to have a default status in Nasty Worlds. It won't do to proceed on the assumption that promise-making is reason-giving. In particular, there is no intuitive sense in which *I promised to ϕ* entitles me to presume that I have a reason to ϕ , absent evidence to the contrary. So, even if a consideration functions as a default reason in our official sense, this doesn't entail that it functions as an *epistemic default* with respect to its being reason-giving. It might do so in the actual world, but it well might not in Nasty Worlds; yet it might be true in both worlds that it is a reason unless undermined.

Epistemic defaults are no doubt important to our moral knowledge in many ways that I cannot explore here.²¹ All that I want presently to note is that being an epistemic default in the above sense is distinct from being a reason unless undermined because it depends on further contingent facts about the general shape of the world. I needn't mind if considerations that are reasons unless undermined are *called* default reasons (this has so far been our official sense of the term), so long as we handle properly the rea-

²¹For instance: much of our moral knowledge might depend on whether the considerations we treat as reasons are largely normal with respect to their reason-giving force. We normally take promising to ϕ to provide a reason to ϕ because we take it implicitly for granted that people normally aren't manipulative. If such background assumptions were largely false without our being aware of this, even our confident judgments about what we have reason to do would largely go astray. This might support the following idea about the epistemic force of normality claims: if p is normally a reason to ϕ , then p is a *good indicator* that there is a reason to ϕ . If people normally *were* coercive, then my having promised to ϕ wouldn't be a good indicator that I have reason to ϕ . That p is a good indicator of q needn't mean that p warrants the belief that q , for whether p is a good indicator of q doesn't depend on available information in the way that having warrant for some belief in given circumstances does. Yet being a good indicator is an epistemic notion. This idea has some interesting implications that I cannot pursue here. For example, insofar as relevant normality claims have no definite extensional implications, p 's being a good indicator of q cannot require much by way of extensional relations between p and q . But I take it that the issue of what it *does* require is no less than the issue of what grounds epistemic principles.

son why they are distinct from epistemic defaults.²² Perhaps they deserve the label on some other grounds. But we have yet to see what those grounds might be.

5. LESSONS FOR GENERALISTS AND PARTICULARISTS

I turn now to the generalism-particularism debate. If the extensional generality of reasons is a contingent matter, how should we view the phenomenon with which I started, namely that our typical response toward many particular facts treats them as moral reasons with a significant degree of generality and remains stable on reflection?

If my argument is on the right track, our responses can assign to reasons a definite degree of *generality* only on contingent assumptions about what happens in the corresponding portion of cases. The truth of such assumptions depends on the world's distributional aspects, just as whether a significant number of birds fly in a world depends on the distribution of, for example, doves, hawks, and ducks over penguins, emus, and kiwis in that world. Should our responses be mistaken, *no* theory of reasons need have trouble explaining our error by reference to mistaken distributional assumptions.²³ Similarly, explaining the *stability* of our responses need require only the assumption that circumstances are normally of a certain sort. In the actual run of things, we take a whole slew of assumptions about our surroundings implicitly for granted: we assume that people normally aren't manipulative, normally have done nothing that justifies killing them,

²²There may naturally be other views of epistemic default reasons that avoid the complaints I voice here by eschewing the assumptions about epistemic defaults which generate those complaints. Mark Lance and Margaret Little perhaps offer such a view when they argue, in "Defeasibility and the Normative Grasp of Context" (unpublished), that cases where one isn't justified in taking appearances at their face value demand explanation in a way that unproblematic cases don't, that the explanation appeals to unproblematic cases where one *can* rely on appearances, and that this is because the evidential valence of an appearance that *p* (that is, being such as to favor believing that *p*) bears also on situations where appearances are misleading or unreliable (because of the drugs one has ingested, or whatever). I have some doubts about whether moral reasons exhibit this sort of explanatory asymmetry (see §7 and n. 35). I take no stand on epistemic reasons here.

²³Accordingly, Roger Crisp isn't entirely fair when he writes (in response to claims like those cited in n. 2 above) that "if particularism is true, we should perhaps be a little surprised that certain considerations dominate our ethical thinking, whereas others do not" (Crisp 2000, p. 36 n. 42).

and so on. These are (probably rational) contingent assumptions, here concerning the general shape of the actual world rather than its specific distributional aspects.

It is, I think, clear that neither of these two sorts of contingent assumptions are entailed by the normative connection between reasons and the actions for which they are reasons, which figures in practical norms of the general form (13):

$$(13) \ p \\ \downarrow \\ \varphi\text{-ing}$$

Such norms serve merely to restrict the connection between non-moral and moral possibilities. They don't restrict the non-moral possibilities themselves or, therefore, the distributional aspects or the general shape of the worlds where they hold.

It would then be very odd if either generalism or particularism by itself had any definite implications concerning these contingent matters. They are in the first instance broadly metaphysical doctrines about the role of moral principles regarding the status of certain facts as moral reasons – about whether the status of (say) *I promised to φ* as a moral reason to φ depends on whether (5) expresses, or is underwritten by, some substantive principle concerning promise keeping. Neither view restricts the distributional aspects or the general shape of the world. So, insofar as reasons of a given sort imply such restrictions, appeals to reasons of that sort provide no *further* support for any theory of reasons.²⁴ Since normality claims incorporate such restrictions, the implications of claims about reasons which carry normality assumptions are tangential to the generalism-particularism debate.

This conclusion will meet resistance. Surely generalism has definite extensional implications: it claims that genuine reasons are *invariant* reasons! My argument, however, implies that the generalism-particularism debate doesn't turn on the existence of invariant reasons. Invariance is an *extensional* notion: something is an invariant reason to φ if it *in fact al-*

²⁴What I say here applies also to Margaret Little's attempt to support particularism by arguing that it supports generalizations about "general epistemic warrant," such as that "stabbing is presumptive of cruelty" (2000, p. 294). As Little notes, appeal to such claims "involves projecting beyond any individual case to the kinds of context one thinks *likely to be encountered*" (2000, p. 295; my emphasis).

ways is a reason to φ . But then it must be a mistake to reject the holism of reasons on the grounds that some non-moral considerations have invariant valence. For if the invariance of a reason is an extensional notion, then by my argument the existence of invariant reasons depends on such contingent facts about the world as its distributional aspects and general shape. Then it must be equally a mistake to reject generalism on the grounds that there are no or very few invariant reasons.²⁵ It is beside the point to debate the status of invariant reasons, which is what is often going between generalists and particularists.²⁶ Indeed, it would be hard to see what all the fuss should be about if the debate were really about the contingent shape of the world.

My argument recommends instead that we construe the generalism-particularism debate in terms of its *modal* aspects. Reasons holism is compatible with the existence of (merely) invariant reasons because it is best construed as a modal thesis about what *may* happen: any consideration that here is a reason to φ *may* elsewhere be no reason at all (or even a reason to not φ), depending on the surrounding circumstances. This allows that some reasons that *can* be undermined are *never in fact* undermined, and so happen to be invariant (see Dancy, 2000, pp. 136–137). Of course, given holism, the invariance of a reason would be a kind of world-historical accident; its status as a reason wouldn't in any way depend on this accidental general truth. The point is just that *if* the surrounding circumstances always happen to line up in a certain way, the moral valence of a reason may, at the end of the day, turn out never in fact to vary. My argument leaves this possibility open while explaining why it is tangential to the generalism-particularism debate, since *how* things happen to line up in a world depends on its distributional aspects and general shape.

How is a modal construal of the debate supposed to help generalism avoid entailing the existence of invariant reasons? Construed modally, holism entails that there is *no necessary connection* between being a rea-

²⁵Similarly, insofar as the notion of regularity is taken to carry more or less definite extensional implications, Dancy misses his target when he announces skepticism about “the need for or even the possibility of the features that generate such properties operating in *regular* ways, that is, those supposed by the generalist” (1999, p. 26; emphasis added).

²⁶For particularist arguments that focus on invariance, see e.g. Dancy (2000, pp. 136–137) and Cullity (2002, p. 173). For generalist discussions with that focus, see e.g. Shafer-Landau (1997, p. 586), who defines an “invariably morally relevant” consideration *non-modally*, as a consideration that always *is* relevant. This is too weak. Hooker (2000a, p. 10) makes essentially the same mistake.

son, on the one hand, and being invariable, on the other – where an invariable reason is one whose valence *cannot* vary.²⁷ To see the distinction between invariance and invariability, take the claim that every war is wrong. A de facto pacifist would mean this in the sense that every war, in the circumstances that *actually* prevail, is wrong. She doesn't take the claim to be necessary. But a radical pacifist does. She means the claim in the sense that every war is wrong, *whatever* the circumstances. The de facto pacifist takes war to be a merely invariant, the radical pacifist an invariable, wrong. This is a modal difference. Then the orthodox assumption that generalism is incompatible with holism would commit generalists to claiming that the mere fact that something is a reason entails that its moral valence is invariable. But surely reasons would have to be invariant if they had to be invariable! I have two comments on this objection.

First, dialectical considerations suggest there must be *something* wrong with the objection. I argued that the degree of extensional generality that reasons have in a given world, as well as certain forms of normality their behavior may exhibit there, depend on contingent facts about that world's distributional aspects and general shape. If generalism entailed the existence of invariant reasons in the relevant extensional sense, assumptions about what these facts are would have to bear on generalism and particularism in a curiously asymmetric way. Suppose these facts were such as to bestow some significant degree of generality on the behavior of reasons. Particularism is compatible with this claim because of its insensitivity to what these facts are, so the claim's truth wouldn't uniquely support generalism. Yet its falsity would have to be uniquely pernicious to generalism, if generalism entailed the existence of invariant reasons. This picture of the dialectic must be wrong.

Second, what I think is *specifically* wrong with the objection is that it relies on an excessively strong notion of an invariable reason. It is only if invariable reasons are supposed to function invariably *without qualification*, or across all instances, that they are also invariant in our extensional sense. This notion still runs modal claims about reasons together with claims that have definite extensional implications. But my argument implies that to construe the parameters of the generalism-particularism debate properly, invariable reasons are to be construed, as far as possible, so as to avoid such implications. The question is whether a viable notion of that sort is available to the generalist. Since reasons that are invariable

²⁷I owe this way of putting the thesis to Michael Ridge (personal communication). Note that, as stated, the thesis allows that some reasons might be invariable all the same.

without qualification cannot be undermined, the generalist needs to deny that moral principles determine reasons that are invariable in this sense.

The literature on both sides largely agrees that this move is a pipedream for the generalist. This is because of the orthodoxy that generalists cannot allow the holistic claim that the very valence of a reason may vary depending on the circumstances. One sign of the powerful hold of the assumption is that generalists typically defend their principles against particularist counterexamples *by complication*. To take a toy example, when the generalist says (claiming to affirm a principle) “That an act would be enjoyable is a reason for doing it” and the particularist retorts “Not if the pleasure is sadistic,” the usual response is “All right, that an act would be a case of non-sadistic enjoyment is a reason for doing it.” But if what is being affirmed is a revised take on an invariant reason, then the response should often seem beside the point. Suppose, on the one hand, that it turns out that there in fact always *is* a reason to go for non-sadistic enjoyment. The particularist can explain this by saying “The reason to do it is simply that it would be enjoyable. But as it turns out, whenever an enjoyable act isn’t sadistic, no further conditions for its presence fail.” Suppose, on the other hand, that there is no reason to go for some non-sadistic enjoyments. For the generalist, this means the existence of counterexamples, while the particularist only needs to give up a contingent factual claim. But whatever further complications the generalist offers in response to the counterexamples, the particularist can replicate her retort if she likes – *ad nauseam*.

Although I am painting with an obviously broad brush here, the dialectic in this toy example should seem dubious (cf. Little, 2000, p. 284). It is also unlikely to resolve the debate either way. For one thing, it nearly inevitably drives the debate to question-begging claims about what exactly functions as a reason and what (if anything) functions as a condition for the presence of a reason. (To make the dialectic neutral with respect to this distinction, I had to use the less committal expression ‘what there is a reason to do’ above.) But if the main issue concerns unqualified invariability, the generalist hardly has any other way to try to secure the maximal extensional generality her reasons must have, but to individuate reasons themselves more finely in the dubious way just illustrated.

The way out is to deny that generalist theories of moral reasons must affirm a necessary connection between being a reason and being invariable without qualification. The option I favor is to construe generalism as claiming that reasons needn’t be invariable within the universal domain, but only in some appropriately restricted domain. A reason that is invariable in this sense isn’t thereby invariable without qualification, for its moral valence is allowed to vary outside the relevant domain. It can clearly be a

reason unless undermined, since reasons that are invariable in this sense must be distinguished from the conditions for falling in the relevant domain. But this is just a special case of conditions for the presence of reasons. Hence generalism, thus construed, would also be compatible with the holism of reasons, contrary to the orthodox assumption that I have been berating.²⁸

I believe that principles of pro tanto moral reason are most plausibly seen as claiming to identify just this sort of reasons. I can believe this only if there is a credible generalist theory of norms of the general form (13) which provides some informative account of considerations that are reasons unless undermined and of the domains within which they are supposed to be invariable, but without restricting the distributional aspects or the general shape of the world. Generalism of this sort wouldn't entail invariant reasons. Its ambitions might well be limited to generic claims like (9), which have no definite extensional implications. To clarify and support the dialectical progression of this section, I want now to outline a view of moral principles that supports just such a form of generalism, and then see what further light it casts on the issue of default reasons.

6. OUTLINING MODERATE MORAL GENERALISM

The view of pro tanto moral principles I want to describe has three basic elements.

A. FAVORING. Just like norms of the form (13), a given principle specifies a relationship of favoring between some non-moral consideration p and φ -ing.

B. HEDGES. If the scope of the favoring relation between p and φ -ing were fixed simply in terms of p , any exception to p 's being a reason to φ would be a counterexample to the corresponding norm. The norm can tolerate exceptions only if its scope is partly determined by further *hedging conditions* whose absence would undermine the status of p as a reason to φ . So, for example, we should think of (5) as hedging the favoring rela-

²⁸The largely dialectic argument for the compatibility of generalism and the holism of reasons that I offer here complements the argument that I offer for the same conclusion in my "Moral Generalism: Enjoy in Moderation" (unpublished). In that paper, I also defend in a bit greater detail some of the dialectic claims I make above, and argue that even a stronger form of holism than the one I operate with here can be *explained* from purely generalist premises. See also McKeever and Ridge (2005).

tion between *I promised to ϕ* and ϕ -ing by conditions such as that the promise isn't manipulated. This would be to think that pro tanto moral principles set out what I'll call *contextually unanimous* reasons: roughly, reasons whose moral valence is *invariable under some more or less inclusive set of hedging conditions*, but may vary when such a condition fails to hold.²⁹ Such reasons lack definite implications concerning the distributional aspects of the world because the frequencies with which the relevant considerations function as reasons depend on contingent worldly facts about just how widely the relevant hedging conditions hold.

C. **NORMATIVE BASIS.** Any theory of moral reasons should have an account not just of what facts function as moral reasons, but also of what makes them do so. It is one thing to say that it is a fact that p , quite another to say that the fact that p gives someone a moral reason to act in a certain way. Any theory should provide an account of the latter type of normative fact (p itself may not have to be a normative fact). Given the structure of my account, this means I need to explain why certain facts favor certain actions *only* under the relevant hedging conditions (and why those conditions are relevant), and why they *must* do so under those conditions. Saying "because in the absence of those conditions those reasons are undermined" is not enough to show that there is any determinate, informatively specifiable truth about what type of conditions can play the hedging role or that that the domains within which the relevant facts are supposed invariably to favor the relevant actions are not morally *ad hoc*. If (5) gives a moral principle, it should be *no accident* that *I promised to ϕ* gives me a reason to ϕ and that certain conditions but not others can undermine the reason. I should be able to offer a principled account of what it is to be a moral reason and what it is to be a hedging condition with respect to a given reason.

Take moral reasons first. The fact that some fact p is a reason to do something needn't always hold (when it does) solely because of the content of p (although it might). From what other source could reasons derive their normative force? The general answer I have in mind is, roughly, that what makes a consideration a reason is at least in part that responding on the basis of it would promote or honor values such as honesty, kindness, and justice, or other moral categories such as rights. This is a first-stab proposal about what makes a consideration a moral reason. If considerations

²⁹The term 'contextual unanimity' was introduced into discussions of probabilistic causality by John Dupré in the 1980's. The notion I introduce here under that name is the normative analogue of the notion one finds in that context. For more details, see my "Moral Generalism: Enjoy in Moderation".

are reasons in virtue of falling under moral principles, then the proposal makes it natural to see a principle as somehow involving a *normative basis* such that the consideration which a principle identifies as a reason is a reason only because of its relation to that basis. Fixing the normative bases of moral principles fixes what favors and fails to favor the actions those principles recommend. As we vary a given principle's normative basis (as different substantive moral theories may do), what the principle recommends may vary as well.

Moving on to hedging conditions, it should be equally natural to think that the domain within which a given consideration *is* a reason is similarly fixed by the relevant principle's normative basis. For example, many egalitarians believe that the fact that some people are worse off than others gives us a reason to remove the inequality. But they typically believe this only when the worse off are worse off through no fault or choice of their own. This indicates that they regard principles of equality as being based on the importance of comparative fairness. When some people are worse off than others through their own fault or choice, this prevents the fact that they are worse off from being a reason to remove the inequality, precisely *because* comparative disadvantage of this sort is not unfair. Similarly, the moral standing of fidelity to one's commitments and the value of sustaining useful social practices number among the possible normative bases for (5), but neither assigns value to the keeping of manipulated promises. Here the explanation might be that to do something under manipulation isn't as such to exercise the virtue of fidelity, and surely the practice one sustains in keeping manipulated promises isn't a useful one.³⁰

This suggests that what makes a feature of a particular situation count as a hedging condition with respect to a reason is that its absence would be such as to prevent responding on the basis of that reason from promoting or honoring the relevant normative basis. Thus only certain sorts of circumstantial variation may lead to variation in moral valence. The point generalizes because there is nothing special about my particular examples. They simply illustrate the sorts of general mechanisms by which tolerable exceptions to moral principles arise and may be explained. Then fixing the relevant principle's normative basis fixes not only what counts as a reason but also what counts as an undermining condition regarding the reasons the principle identifies.

³⁰Essentially the same point applies to several other candidates, such as that promise-keeping gives due weight to expectations one has created, or that it gains others' trust. Which candidate one accepts depends on one's substantive moral theory. I use these examples merely for illustrative purposes.

To summarize: I propose a conception of pro tanto moral principles as *hedged principles* which involve: (A) a relationship of favoring between some non-moral consideration and the response for which it is a reason, (B) hedging conditions that limit the scope of that relationship, and (C) a normative basis that both underwrites the normative relationship in (A) and determines what conditions fall under (B).³¹ I have tried to highlight some of the intuitive and explanatory merits of this view in sketching it. I want now to run a manifesto on further points that the view scores.

As desired, the view is silent on how widely the hedging conditions hold; this is left to depend also on the world's distributional aspects and general shape.³²

As designed, the view respects the holism of reasons.³³ It draws a distinction between reasons and conditions for the presence of reasons. Hedged moral principles allow the moral valence of reasons to vary when the hedging conditions aren't satisfied.

As should be possible (already by holism), some reasons may turn out to be invariable without qualification all the same. In my terms, one way this might happen is that a consideration that functions as a reason somehow (not necessarily conceptually) has its normative force "built into" it, so that the reason gets to bear the required relation to a normative basis without further ado by containment. Perhaps the example Cullity (2002, p. 182) gives of an "invariant" reason – that inflicting suffering on others for your own enjoyment always counts against an action – is like that.

As advertised, the view predicts that the current orthodoxy in the generalism-particularism debate is false. For if (as I claim) the view is moderately generalist, then not all forms of generalism rule out holism.

As you may have guessed, the reason I claim that the view is moderately generalist is that, although it doesn't entail that reasons are invari-

³¹This complex structure may not be transparent in our ordinary *formulations* of moral principles. I don't think this is a problem: we have good reasons to distinguish principles from their linguistic formulations.

³²It is, of course, a contingent fact that a consideration *does* in fact function as a reason in a given case, since it is contingent *that* the hedging conditions hold. This is a perfectly harmless form of contingency.

³³Holism isn't strictly entailed, however. Technically the view allows that being a reason entails being invariable without qualification. In my terms, the claim would be that, for every moral principle, the set of hedging conditions associated with it is empty. (Of course, if that were true, then the introduction of hedged moral principles would be harder to motivate.) The resulting reasons would be a limiting case of contextually unanimous reasons.

able without qualification, it does yield contextually unanimous reasons. Explicating hedging conditions in terms of normative bases in the way I do ensures that if there are true hedged principles, the reasons they identify function invariably within their domains. Insofar as those domains are determinate, the restrictions that hedged principles place on which circumstantial variations may lead to variation in moral valence will be no less recognizably generalist. The relevant invariability won't be a merely epistemic matter (contra Dancy, 2000, p. 137). Because of the role that normative bases play in fixing reasons and hedging conditions, it will be essential to the status of a consideration as a moral reason in a particular case that it be the same reason invariably within a certain domain. In not making moral reasons fundamentally context-dependent, the view avoids collapsing into particularism. Moral reasons will still depend on some suitable provision of substantive moral principles.

Needless to say, this moderate form of generalism faces important challenges. In particular, one might worry that the apparently open-ended nature of undermining conditions – the open-ended variety of forms that manipulation can take, of bad things that one can enjoy, and so on – makes it impossible to specify them informatively. I have a great deal to say about this challenge elsewhere.³⁴ Here I can only assert that the open-ended nature of hedging conditions doesn't show that the membership in that category is indefinite and therefore not amenable to an informative account. (Such an account needn't come in the form of a finite list.) The limited point I would like to make in favor of the view here is that it is nicely motivated by my recommended focus on the modal aspects of the generalism-particularism debate, since both detach the debate from issues about the distributional aspects or the general shape of the world.

7. ANOTHER LOOK AT DEFAULT REASONS

Let me close by returning to my ostensible topic: Is there a viable notion of a default reason that matters for the generalism-particularism debate?

Defaultness is a property of the valencies that considerations have (Kirchin, 2003, p. 70 n. 22). The notions of defaults we have so far weeded out had to go because they involved extensional, normality, or epistemic assumptions that are independent of generalist and particularist theories of moral reasons. But in the literature one also finds the distinct idea that

³⁴See my "A Theory of Hedged Moral Principles" (unpublished).

some considerations need no further conditions to enable them to function as moral reasons, though other conditions may, by their presence or absence, neutralize or turn the direction of their valence (see Kirchin, 2003, p. 61; Dancy, forthcoming). Call these *moral defaults*. Metaphorically, the idea is that moral defaults “start out” as reasons – they “arrive switched on,” as Dancy puts it – but may be “switched off” in certain sorts of contexts.³⁵ So when a fact is a moral default, this is akin to being a default setting of a computer program: the reason-giving status kicks in automatically when the fact obtains, and stays put unless something changes the setting.

The problem I want to raise here is that the moderate generalist account of moral reasons outlined in §6 finds little to recommend interest in moral defaults explicated in this sort of *asymmetric* fashion. The reasons figuring in hedged moral principles – what I have called contextually unanimous reasons – do depend for their status as reasons *negatively* on the absence of undermining conditions. But they *also* depend for their status as reasons *positively* on (whatever our substantive moral theory designates as) the relevant principles’ normative bases, for these are what make them reasons. In that case, the considerations that we find in hedged principles will not generally “arrive switched on” somehow on their own, as the above notion of a moral default looks to require. But then the interest and the usefulness of moral defaults to debates concerning contextually unanimous reasons – on which I have argued the success of generalism to hang – will be severely limited.

I conclude that people talking about moral defaults have in mind a notion that is richer than our initial notion of a reason that can be undermined, but the richer notion has neither a clear rationale nor a clear bearing on the central issues at stake in the generalism-particularism debate. As with epistemic defaults, it doesn’t much matter if we stick the *label* ‘default reason’ onto our initial notion. But the label ‘defeasible reason’ strikes me as perfectly adequate and less liable to mislead.

³⁵Dancy frames this account of moral defaults also in explanatory terms: “[I]f a default reason-giving feature does give us a reason in this context, there is nothing to explain; we only have something to explain when such a feature doesn’t provide a reason” (Dancy, forthcoming). This doesn’t seem equivalent to the formulation in the text. If the cases where the feature does function as a reason are the utterly familiar case, then in a sense there is nothing to explain. But that sense concerns the pragmatics of explanation that by itself shows nothing about the metaphysics of reasons. Denizens of Nasty Worlds might be in an epistemic predicament where there is something to explain. More generally, even if the feature carries moral valence “on its own,” that fact may be both unobvious and explainable. Also, the view sketched in §6 can allow this kind of explanatory asymmetry while denying the metaphysical asymmetry I challenge in the text below.

To be sure, there are ways of relaxing the above account of moral defaults which make it compatible with moderate generalism. For instance, we might allow the modal profile of moral defaults to depend also on the character of the world – more weakly than in the case of epistemic defaults, but even so. There may be worlds, even more far out than Nasty Worlds, whose general moral truths differ from ours because the systematic connections between moral and non-moral possibilities differ from those that hold in worlds like ours. Perhaps the denizens of some worlds evolved into really good lie-detectors, in which case lying wouldn't deceive and there arguably would be no basis for regarding lying as wrong-making in such a world. Then we could say that lying has a default negative status, but only relative to certain sorts of worlds.³⁶

But again I fail to see what the label 'default' is supposed to add. Insofar as it connotes asymmetry, it is still liable to mislead. Talking simply about defeasible reasons seems perfectly adequate. Suppose a true moral principle identifies the fact that p as a reason to φ . What scenarios like the above highlight isn't the possibility that p may be undermined as a reason to φ , but the possibility that p may lack even that defeasible status in far-out worlds where different principles reign, and hence not have it across all worlds. This expands our account of the conditions under which the relationship between p and φ -ing may break down. But the core generalist commitment regarding defeasible reasons of this sort remains untouched: p must function as a reason under certain (informatively specifiable) conditions, or else the relevant principle is false. This commitment expresses what it is for p to be a contextually unanimous reason, which is just the kind of reason for which my construal of the generalism-particularism debate in terms of its modal aspects told us to look. This strikes me as all the more reason to recommend generalists and particularists alike to focus on the merits of contextual unanimity as a constraint on at least the basic moral reasons.

³⁶Note that even if lying isn't wrong-making in a lie-detector world, the different hedged moral principles that hold in different worlds can still have their normative bases in common (something about deception, perhaps). They differ on the considerations that enter into the principles and the conditions under which those considerations function as reasons. Because the view I favor is sensitive to this sort of circumstantial variations, it can explain how and why the principles operative in far-out worlds might be *intelligibly* different from ours by indicating how they are related to ours. It can therefore explain how and why the principles *we* accept might intelligibly have been different. Note also that even if the modal profiles of hedged moral principles have this sort of limits, they may nonetheless serve to express necessary moral truths about given worlds. Hooker (2000b, pp. 231–234) discusses the general issue (in the context of two-tier theories).

ACKNOWLEDGMENTS

Many thanks to the participants of the BSET 2003 Conference at Queens University, Belfast, for lively and instructive feedback on an earlier version of this paper, both in session and after-hours. I hope they had as much fun as I did. Special thanks to Vivienne Brown, Jonas Olson, Piers Rawling, and Soran Reader for pushing especially hard and to Michael Lacewing for backup. I'm also grateful to Maggie Little and to an audience at Lewis and Clark College for discussions, and to Michael Ridge, Jussi Suikkanen, and an anonymous referee for BSET for comments, on earlier versions of the paper.

REFERENCES

- Asher, N., and Morreau, M., What Some Generic Sentences Mean, in Carlson and Pelletier (eds.), 1995, pp. 300–338.
- Bonevac, D., Against Conditional Obligation, *Noûs* 32 (1998), pp. 37–53.
- Carlson, G.N., and Pelletier, F.J. (eds.), *The Generic Book*, Chicago, IL: University of Chicago Press, 1995.
- Crisp, R., Particularizing Particularism, in Hooker and Little (eds.), 2000, pp. 23–47.
- Cullity, G., Particularism and Presumptive Reasons, *Proceedings of the Aristotelian Society Supplementary Volume* 76 (2002), pp. 169–190.
- Dancy, J., *Moral Reasons*. Oxford: Blackwell, 1993.
- Dancy, J., Defending Particularism, *Metaphilosophy* 30 (1999), pp. 25–32.
- Dancy, J., The Particularist's Progress, in Hooker and Little (eds.), 2000, pp. 130–156.
- Dancy, J., Moral Particularism, in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2001 Edition). URL: <<http://plato.stanford.edu/archives/sum2001/entries/moral-particularism/>>.
- Dancy, J., *Ethics without Principles*, Oxford: Oxford University Press, forthcoming.
- Drewery, A., Laws, Regularities and Exceptions, *Ratio* 13 (2000), pp. 1–12.
- Hooker, B., Moral Particularism: Wrong and Bad, in Hooker and Little (eds.), 2000, pp. 1–22. [Hooker 2000a]
- Hooker, B., Reflective Equilibrium and Rule Consequentialism, in B. Hooker *et al.* (eds.), *Morality, Rules, and Consequences*, Lanham, MD: Rowman & Littlefield, 2000, pp. 222–238. [Hooker 2000b]
- Hooker, B., and Little, M. (eds.), *Moral Particularism*, New York, NY: Oxford University Press, 2000.
- Kirchin, S., Particularism, Generalism, and the Counting Argument, *European Journal of Philosophy* 11 (2003), pp. 54–71.
- Korsgaard, C. Skepticism about Practical Reason, *Journal of Philosophy* 83 (1986), pp. 5–25.
- Krifka, M. *et al.*, Genericity: An Introduction, in Carlson and Pelletier (eds.), 1995, pp. 1–124.
- Little, M., Moral Generalities Revisited, in Hooker and Little (eds.), 2000, pp. 276–304.
- McKeever, S., and Ridge, M., What Does Holism Have to Do with Moral Particularism?, *Ratio* (2005) (forthcoming).

- Morgan, C.G., The Nature of Nonmonotonic Reasoning, *Minds and Machines* 10 (2000), pp. 321–360.
- Morreau, M., Fainthearted Conditionals, *Journal of Philosophy* 94 (1997), pp. 187–211.
- Pelletier, F. J., and Asher, N., Generics and Defaults, in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*. Amsterdam: Elsevier/Cambridge, MA: MIT Press, 1997, pp. 1125–1177.
- Priest, G., Validity, in A. Varzi (ed.), *European Review of Philosophy, Vol. 4: The Nature of Logic*. Stanford, CA: CSLI Publications, 1999, pp. 183–206.
- Raz, J. The Truth in Particularism, in *Engaging Reason*. New York, NY: Oxford University Press, 1999, pp. 218–246.
- Reiter, R., A Logic for Default Reasoning, *Artificial Intelligence* 13 (1980), pp. 81–132.
- Shafer-Landau, R., Moral Rules, *Ethics* 107 (1997), pp. 584–611.
- Sinnott-Armstrong, W., Varieties of Particularism, *Metaphilosophy* 30 (1999), pp. 1–12.

Department of Philosophy
1238 Social Sciences and Humanities Building
University of California
Davis, CA 95616, USA
E-mail: pekkav@ucdavis.edu

