

Contours of the Practical Landscapeⁱ

David McNaughton and Piers Rawling

Preface

While both Piers and I have learned much from Jonathan Dancy, I have had the privilege and pleasure of having Jonathan as a colleague and personal friend over many years. I was appointed as a Lecturer at Keele in 1970 and Jonathan joined the department a year later. We were of an age, and quickly became friends and allies. From the outset, Jonathan was highly impressive. He was not only extremely clever and very quick thinking, but willing to challenge orthodoxy. Philosophical debates tend to be conducted on certain agreed assumptions in terms of which the problem, and its possible solutions, are structured. It was Jonathan's willingness to question such assumptions that made his thinking so innovative and stimulating. Conversations with him – of which there were a very large number – were thus always liable to take unexpected, and even slightly unsettling, turns.

Jonathan was a profound influence on my philosophical development in incalculably many ways, but I want here to mention two crucial respects in which I am deeply indebted to him. In my early years I was a convinced, if somewhat reluctant, ethical non-cognitivist, and disciple of Richard Hare. Jonathan's incisive critique of the case for non-cognitivism persuaded me that moral realism was not only defensible but highly plausible – a position I have never since abandoned. For about ten years, he and I jointly taught an exciting and demanding final year course on meta-ethics and ethical theory. Jonathan's Socratic method of tutorial teaching kept everyone, including me, on our toes. I have rarely had to think so hard, nor enjoyed it more.

Early publication was not especially encouraged, still less required, in those days, nor were we given much by way of guidance on how to publish. My technique for writing papers – think of a philosophical problem I found deeply perplexing and try to solve it – was not a recipe for success. Indeed, I began to feel a failure and seriously considered leaving the profession. Jonathan, who had already published *An Introduction to Contemporary Epistemology*, as well as two ground-breaking articles on moral particularism, encouraged me to write an introductory book on ethics, based on the course we had jointly taught for many years. As very many of the ideas that had emerged from that teaching were primarily Jonathan's, this was a deeply generous gesture. Not only did he suggest the project, he read every draft of *Moral Vision* and made extensive and extraordinarily helpful suggestions for its improvement. His encouragement and guidance at this crucial stage changed the course of my professional life. It is the sort of debt that one cannot repay, but for which one can simply remain profoundly grateful.

I

Introduction

Jonathan Dancy is renowned both for his moral particularism and his reasons holism. Moral particularism, in its original formulation, is, roughly, the thesis that there are no moral principles that are true, finitely storable, and knowable. Moral reasons holism is the claim that any feature might count as a moral reason, and that (almost) any feature that can count as a moral reason in favour of (or against) an action that bears it might be morally irrelevant, or even count morally against (for), such action in other circumstances.

How are particularism and holism related? That depends, of course, on the details of specific versions of the doctrines. But here's the general idea. Particularists certainly reject strict principles – say: lying is never morally permissible. They also, however, reject weaker principles, such as: the fact that your utterance would be a lie always counts against it morally speaking (even in cases in which lying is morally permissible). (Dancy (1993, p. 61) gives the example of certain games that require lying.) Such weaker principles can be expressed in terms of reasons: the fact that your utterance would be a lie is always a moral reason not to utter it (even when this reason is outweighed by reasons on the other side). The moral reasons holist's denial of this is of a piece with the particularist's rejection of the weakened principle against lying. But aren't both these positions easily defanged? Even if the fact that an act would be a lie doesn't always count against it, perhaps this is true only under certain conditions – specify these conditions and we can then safely say that, in their absence, lying counts against. However, both the particularist and the holist might adhere to a doctrine of uncodifiability: such conditions can't be specified (at least, not non-normatively – see below).

Moral reasons are a type of practical reason, and holism can, of course, be extended to practical reasons generally. One of our concerns will be with the relation between such holism and what one might call the contours of the practical landscape. Like us, Dancy sees a potential difficulty here (2004, p. 111):

Holism maintains that anything whatever might make a practical difference, or provide a reason, if the circumstances were suitable. It sees no difference, apparently, between such features as being very damaging to one's health and the number of leaves on a tree. It sees no difference between the causation of unwanted and undeserved pain and whether one sets out with the right foot or the left. If there are differences between these things, it can only be that one of them matters more often than the other. But we all think that this is not all there is to it. The stubborn intuition is that though holism may be right in stressing the possibility of exceptions to all moral rules, still there are rules, and there must be rules for there to be exceptions. What is more, the exceptions are not statistical exceptions

but moral exceptions. These exceptions bear their exceptionality on their face, as one might put it.

In fact, holism, in a more sweeping form, not only maintains that ‘anything whatever might ... provide a reason’, but also that any consideration, such as ‘being very damaging to one’s health’, can switch “valence”, as it is sometimes put, in analogy to the chemical notion. That some act would severely damage your health is ordinarily a reason against performing it (negative valence); but the radical holist claims that there are circumstances where this consideration would be irrelevant (zero valence), or even a reason in favour (positive valence). For many considerations, the idea of switching valence raises no concerns: that it’s cold outside is usually a reason for you to wear your coat, but if you’re trying to impress your Muscovite hosts in midwinter it may be a reason not to. For others, though – the fact that an act would, say, inflict undeserved harm – the idea of variance may well seem problematic: inflicting undeserved harm, it might be thought, has invariantly negative valence.

One proposal to solve the contouring difficulty for the holist, then, is to limit the holism. This is one thing we have proposed in the past (McNaughton and Rawling, 2000), and we’ll revisit it here (section II), in part to contrast it with Dancy’s own proposal (section III). But we shall also explore a somewhat different approach to structuring practical reason. In the case of valence switching, the strength of a reason switches from positive to negative as circumstances vary. Moving from positive to negative strength may be viewed, however, as just a decline in reason strength in which, as it were, the reason’s strength passes through zero. But changes in the strength of a reason can, of course, occur in the absence of a switch in valence – more generally a reason’s strength simply varies in magnitude as circumstances vary. In section IV we explore the extent to which a systematic account of reason strength can be given by looking to the notions of welfare and value, and the distinction between them.

II

Universal reasons

Donald Davidson (1980) famously argues that practical reasons are mental states with causal powers. But we adopt a different usage, according to which practical reasons are facts – the fact that it’s cold outside is a reason for you to wear your coat.

It’s important to note that there are two facts lurking here. Your reason is the first: it’s cold. But there’s also a second: the fact that the first fact is a reason. We have, then, a two-tier view of practical reasons. At tier one are the reasons; at tier two are the facts that the tier one facts *are* reasons. Experience tells us that it’s easy to muddle this distinction, so perhaps it helps to appreciate that the two tiers give rise to different possibilities of error: you might be mistaken about the

weather (tier one error); or you might fail to realize that cold weather is reason to wear a coat (tier two error).

In some contexts 'reason' refers only to a cause, as in, 'metal fatigue was the main reason for the bridge's collapse'. For Davidson, 'reason' refers, in psychological contexts, to a cause that rationalizes – if, say, his reason for calling was that he wanted to persuade her to go, and he believed he could do so by calling her, then, according to Davidson, these mental states caused his calling. But they also rationalize his behaviour. And rationality, unlike causality, is a normative notion: *ceteris paribus*, we *should* be rational. Our notion of a reason is also normative: if it's cold, and this fact is a reason for you to wear your coat, then (roughly speaking), in the absence of countervailing reasons, you should wear your coat.

Rationality, and reason (in our sense), are part of a broader family of intertwined normative notions. What is normativity? We have no analysis of it in other terms (we take seriously Butler's dictum that 'everything is what it is, and not another thing' (1970, 14)); all we can do is give illustrations of normative notions, and their relations. On our usage, for example, a normative requirement is not made so by the fact that most people *do* heed it (if they do), but by the fact that they *should* heed it – they have overriding *reason* to heed it. Or consider another pair of normative notions: the evaluative and the moral. To say that a potential act of yours would make the world better is, on Scanlon's (1998) 'buck-passing' account of the good, to say that you have a reason of a certain kind to perform it. If an act would be morally wrong, then (on many views, at least – but see section VII below) you have overriding reason not to perform it. The list of other normative notions includes harm, benefit, innocence, desert, justice, fidelity, gratitude, reparation; and these too are inter-related. W.D. Ross (1930, p. 21), for example, takes justice to be concerned with the distribution of benefits and harms in accord with desert; and desert, of course, is tied to innocence. The normative is also tied to the non-normative. Supervenience holds, for instance: innocence, say, cannot be lost without a change in non-normative circumstances. And there may be non-normative items on lists of benefits, harms, and valuable (or not) features of states of affairs. Not all such items are non-normative, however; and normativity also enters when we ask what unifies the lists.

According to what we'll call the "simple view" of reasons, normative notions do not enter into their content. On this view practical reasons (tier one facts) are non-normative facts. It is tier two facts that are normative – i.e., have normative content. (Note that adverting to content is crucial here – normative facts can be *picked out* using non-normative vocabulary, at least on occasion: for instance, 'the first fact David thought of when waking today' might pick out the fact that he has reason to go into the office.) A practical reason (on this view) is not itself a normative fact; it is the fact that you have it that's normative.

But is this view correct? Suppose an act you're contemplating would cause someone else undeserved harm. This, we think, is a reason against performing it, the normativity of the notions of harm and desert notwithstanding. We reject, then, the simple view of reasons, and allow that there are normative tier one facts – i.e. what we call 'normative reasons' (usage varies). And these may help with the contouring difficulty: there are some normative reasons that count against in all circumstances, and some that count in favour in all circumstances. In our view, we've just seen one of the former: that your act would cause someone else undeserved harm counts against it in all circumstances.ⁱⁱ

Our discussion so far, however, has neglected a logical possibility: that a reason might count both for and against the very same act – might have both positive and negative valence – in the very same circumstance. Then it might be true of some reason both that there is no circumstance in which it doesn't count against (say) and that there is a circumstance in which it also counts in favour. The standard holist response to any purported example of, say, a reason that always counts in favour, is to come up with a case in which it counts against. But, if a reason can count in both directions at once, then merely coming up with a case in which a reason counts against does not suffice to rule out its always counting in favour. So, are there any cases of reasons counting in two directions at once?

A previous example will serveⁱⁱⁱ. You find yourself in Moscow. It's a cold winter's day. This is a reason to wear your coat. On the other hand, not wearing your coat on such a cold day would impress your hosts (something you have reason to do, let's say). Thus the fact that it's cold is, here, both a reason for and against wearing your coat. Now, of course, there's a story to be told – we've just told part of it – and it might be maintained that if we spell everything out we'll see that no *full* reasons count in two directions at once. On the one hand, that it's cold and not wearing your coat would be uncomfortable is a reason for you to wear it. On the other, that it's cold and not wearing your coat would impress your hosts with your endurance of discomfort is a reason for you not to wear it. But matters here are somewhat tricky.

Dancy (2004, pp.97-99; see also p.123) criticizes attempts, by such authors as Bennett and Raz, to defend the notion of a 'complete' reason – if reasons are complete, holism is false, since the more you pack into the content of a reason, the more difficult it becomes to switch its valence. On Dancy's (2004, see e.g. p.38ff) view, reasons are to be distinguished from 'enabling' (and 'disabling') conditions. Reasons for Dancy, we'll say, are "compact", and our previous example stands: the fact that it's cold is the reason; but it is differentially enabled. What enables it to be a reason to wear your coat? That doing so would prevent discomfort. What enables it also to be a reason against wearing your coat? That not wearing your coat would impress your hosts.

Here's how we see the logic of the situation. One component of a thoroughgoing holism is:

(H) For any possible fact F, act phi, circumstance C, and agent A:
 If F is a reason favouring (against) A's phi-ing in C, then there is a circumstance C*, distinct from C, and an agent A* such that F is a reason against (favouring) A*'s phi-ing in C*.

But if the reasons are not “restrained”, they will “swallow” the circumstances, rendering (H) false for, as it were, the wrong reasons: it would not be falsified by a reason that cannot switch valence, but by the fact that a reason F will include all its circumstances, so that wherever F holds, C obtains, and hence there is no C* distinct from C in which F holds. This route to falsifying (H) is blocked by taking a compact view of reasons. The compact view, however, also allows for reasons that count in both directions at once – so that (H) is not inconsistent with the claim that some reasons count in favour (or against) in all circumstances.

So now let's return to a reason that, we contend, counts against in all circumstances: the fact that your act would cause someone else undeserved harm. Are there cases where it also counts in favour? Suppose someone threatens that he will inflict undeserved harm on several people unless you inflict it on one (it doesn't matter what harm, or on whom – it just has to be unjustly inflicted). That your act would cause someone else undeserved harm now seems to be a reason in favour of performing it. But it also remains a reason against. There is, of course, a difference in the means-end status of the act when considered, as it were, from the two different directions. Avoiding the infliction of undeserved harm is always an end in itself. However, in this case, its infliction is also a means to preventing the infliction of even more of it, so that it counts here in both directions at once.

That your act would inflict undeserved harm is an example of what we'll call a “universal” reason against – a universal reason being a reason^{iv} that counts in one direction (but not necessarily in *only* one direction) in all circumstances. There are other universal reasons; and all such, we think, are normative (but not all normative reasons are universal). How much, however, do these universal normative reasons really help with the contouring difficulty? One thought is that, in the case at least of moral reasons, there is always some universal reason present. Take lying, for example. It is not always the case that you have moral reason not to lie – we certainly agree with Dancy that there are circumstances in which the fact that your utterance would be a lie does not tell against uttering it morally speaking (he cites (1993, p.61) the example of certain children's games in which lying is part of the fun). But when you do have moral reason not to lie, then, the idea runs, the lie would fall under a normative description that points to a universal reason against it. For example, lies can harm innocents, or violate fidelity (cf Ross, 1930, p.21). And the fact that your act would harm an innocent or violate fidelity is always a reason against it. When there is a moral reason not to lie, there is, then, we claim, some universal reason also operative – a normative reason that counts against in all circumstances. (It is, however, merely

a contingent statistical fact about our world that most lies fall under such further normative descriptions.)

We have suggested a similar approach before (McNaughton and Rawling, 2000), to which Dancy responds as follows (2004, p.119):

I think it fair to say that McNaughton and Rawling never offer any reason for supposing that variance is impossible without invariance. This is, of course, what is really needed. Otherwise the dispute will descend to examples, when for the particularist the point really is why we need the invariant at all.

The particularist here is not quite the person we characterized in our opening paragraph – someone who denies that there are any moral principles – rather, she is an advocate of Dancy’s updated version of the position, who claims that ‘the possibility of moral thought and judgement does not depend on the provision of a suitable supply of moral principles’ (ibid. p.7). This position can, of course, be extended to practical reasons generally: the transcendental holist, as we’ll call her, maintains that the possibility of practical thought and judgement does not depend on there being any universal reasons. Do we have an argument against transcendental holism?

Well, following Wittgenstein (1953), we certainly don’t think that appealing to rules is helpful in explicating normativity *per se*. And we don’t have a sort of transcendental argument to the effect that variance is impossible without invariance – Dancy is correct in this. Do we, then, have any argument to the effect that invariance is ‘needed’ – or, to put it in our terms, that universal reasons are needed?

In short, no. That is because we do not think that universal reasons are needed for any purpose. For instance, we certainly don’t adhere to a model of practical reasoning according to which you should reach conclusions about what to do by a process of syllogistic deduction, or some such, so universal reasons are not needed as premises in deductive arguments. Rather, we simply see universal reasons as an integral part of the practical landscape. Ultimately (as we’ll shortly see) we see the strength of your reason to perform some act as dependent on the amount of value (good or bad), benefit and/or harm it would contribute – and, in the case of benefit and harm, to whom. And universal normative reasons are universal because they concern the ways in which your act would contribute in these respects. In the case of lying, for instance, we cited harm to innocents and fidelity. The first of these refers to harm directly – harm to an innocent, which is bad. And breaches of fidelity, we think, are harmful to the victim of the breach in a particular way. (There is an issue, we acknowledge, about how interesting these universal reasons are – we see it, for example, as built into the notion of fidelity that one always has reason not to breach it. However this, obviously, does not exhaust the content of the notion.)

In our view, then, reasons holism reaches its limit when confronting universal reasons, which allow for principles such as: breaches of fidelity always count against. This is not to deny, of course, that competing reasons might outweigh this consideration, so that you have more reason to breach than not. We'll return later to the topic of weighing reasons; we turn now to Dancy's own approach to the contouring difficulty.

III

Dancy's View

Whereas we propose that there are universal reasons, Dancy proposes instead that there are merely 'default reasons' (2004, pp.112-113), where:

A default reason is a consideration which is reason-giving unless something prevents it from being so. The idea is that some features may be set up to be reasons, in advance as it were, although it is always possible for them on occasions to fail to give us the reasons they are set up to give. One can express this idea more or less metaphorically. More metaphorically, one could say that some considerations arrive switched on, though they may be switched off if the circumstances so conspire, while others arrive switched off but are switched on by appropriate contexts. Less metaphorically, one could say that if 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. With other features it is the other way around; if they do provide reasons there is something to explain, and if they don't, there isn't.

In some ways this approach echoes that of W.D. Ross in the following passage (1930, p.138):

Pleasure seems, indeed, to have a property analogous to that which we have previously recognized under the name of conditional or *prima facie* rightness. An act of promise-keeping has the property, not necessarily of being right but of being something that is right if the act has no other morally significant characteristic (such as that of causing pain to another person). And similarly a state of pleasure has the property, not necessarily of being good, but of being something that is good if the state has no other characteristic that prevents it from being good. The two characteristics that may interfere with its being good are (a) that of being contrary to desert, and (b) that of being a state which is the realization of a bad disposition. Thus the pleasures of which we can say without doubt that they are good are (i) the pleasures of non-moral beings (animals), (ii) the pleasures of moral beings that are deserved and are either the realizations of good

moral dispositions or realizations of neutral capacities (such as the pleasures of the senses).

One difference, of course, is that Ross is speaking of rightness and goodness rather than reasons. And another is that, in his discussion here of the goodness of pleasure, Ross lists what he takes to be the characteristics that may 'disable' pleasure's goodness. Dancy, as we read him, would deny that any lists of enablers or disablers can be given in advance of consideration of particular cases (see, e.g., *ibid.* p.122).

While we do not see the gulf between our view and Dancy's as that wide – he would seem to agree with us, for example, that certain normative reasons have a special status (when it comes to thick concepts, for example, he wants 'centrality', though not invariance: *ibid.* p.122) – nevertheless we have some doubts about his view.

How is the remark that 'if a default reason-giving feature does give us a reason in this context, there is nothing to explain' to be interpreted? What if an inquirer falsely believes that a disabler is present (i.e., falsely believes that something is preventing the default reason from being a reason)? To such an inquirer, the fact that the default reason is, as it were, operative *does* require explanation. Dancy contends (*ibid.* p.113), for example, that the fact that an act would be just is a default reason in favour of it. On his account, then, it is only in the presence of disablers that there is anything to explain concerning justice's valence: when justice counts in favour there is nothing to explain. But suppose you falsely believe that there is a disabler present. Then there *is* something to be explained to you – namely, that there is no disabler present, so that justice counts favourably here.

This difficulty might be avoided, of course, by appeal to an account of explanation that is not inquirer-relative. Dancy rejects (*ibid.* p. 46), however, one non-inquirer-relative account – what he dubs 'the completeness of a full explanation' view – in favour of a view that parallels his account of reasons and their enablers. On this view (*ibid.* pp. 45-49) the explanans explains the explanandum even though it would not do so were certain enabling conditions not met, where these conditions stand outside the explanation. Dancy suggests, for instance, that a 'full causal explanation of an event might be thought of as one that specifies a sufficient set of events as causes' (*ibid.* p. 46), with the relevant laws playing the role of enabling conditions for such an explanation as opposed to being a part of it. Or, in the case of explaining an action, your having the belief, say, that it's cold outside is part of what enables the fact that it *is* cold outside to explain your wearing your coat – your having the belief is not itself part of the explanation.^v

How might this account of explanation be applied to the explanation of some fact's counting as a reason in favour of an act? Consider our case of its being

cold outside favouring your wearing your coat. Perhaps part of the explanation of this case of favouring is the fact that your wearing your coat would prevent your discomfort – one of the facts, as we have seen, that enables the fact that it's cold to be a reason for you to wear your coat. If so, then an enabler of the reason contributes to explaining its reason-giving force. If Dancy's default reasons required no enablers, as suggested by his claim that a 'default reason is a consideration which is reason-giving unless something prevents it from being so' (ibid. p.112), then there would be no enablers to explain the default reason's reason-giving force – just as expected if there is nothing to explain.

On Dancy's account, then, it looks as though the claim that there is nothing to explain when a default reason counts in the default direction might be equivalent to the claim that defaults require no enablers. But if so, trouble looms, since Dancy suggests that all reasons require enablers (ibid. pp. 38-40). For example, the fact that you are able to phi is a 'general enabler' (ibid. p. 40) of the reasons that favour phi-ing. One response to this might be to point out that ability doesn't help explain reason-giving force – it is not what we might call an 'explanatory enabler'. So it could be that default reasons are distinct in having no *explanatory* enablers. But this returns us to the lack with which we began: we need a non-inquirer-relative account of explanation to elucidate the notion of an explanatory enabler.

Dancy also discusses (ibid. pp. 113-117) an approach to the contouring difficulty, due to Mark Lance and Margaret Little (ms), that looks to defeasible generalizations. But he remains unconvinced that their approach yields a 'way of moving beyond the notion of a default reason' (ibid. p. 117). At this point, then, rather than continuing to canvass sufficient conditions for a fact to be a practical reason, perhaps we should look for other ways of structuring the practical landscape. Can we formulate necessary conditions for something to be a reason? Can we classify reasons into types? What contributes to their strength?

IV

Welfare and the good

Whereas, say, Bernard Williams (1979) explores the idea that among the necessary conditions of reasonhood are certain features of an agent's motivational psychology, we propose to look in a different direction – at the claim that, roughly, you have reason to perform some act only if benefit or good would ensue. If this is correct, then a universal practical question is: what benefit or good would my proposed act contribute (or what harm or bad would it prevent)? Of course, the notions of good and benefit, it might be said, are so close to the notion of a reason itself as to make all this rather uninteresting. But, if so, we will have succeeded in part of our enterprise, since this close relationship is part of what we hope to establish.

What about types of reasons? We see reasons as falling into three and only three categories – the personal, the special, and those associated with considerations of value. Roughly, “personal reasons” are reasons you have to benefit yourself; “special reasons” are reasons to benefit those to whom you stand in special relationships of various sorts; and “value reasons” are reasons you have to promote the general good. The strength of your reason to perform some act, on our view, then, is a matter of how much so acting would benefit you or those to whom you stand in some special relationship, and/or how much good it would do. We turn now to look at some details, and some contrasting views.

In acting you modify the state of the world. Some states are better than others; and some states are better *for you* than others. The 'better than' relation ranks states in accord with their goodness or value; the 'better for x than' relation ranks states in accord with how beneficial they are to x – that is, in accord with x's level of well-being in them (we use 'benefit', 'well-being', and 'welfare' interchangeably). Terminology can be confusing here: some authors distinguish between impersonal and personal value, but as we use the term, value is always impersonal; it is welfare that is personal. Each is measured on an objective scale – it is not the case, for example, that x's welfare scale varies in accord with the perspective of the evaluator. And care must be taken not to conflate welfare scales with the value scale.

Consider, for instance, the following case from Ross (1930, pp.34-5):

Suppose ... that the fulfillment of a promise to A would produce 1,000 units of good for him, but that by doing some other act I could produce 1,001 units of good for B, to whom I have made no promise, the other consequences of the two acts being of equal value ...

Ross sees this case as a potential counterexample to the view that right acts are 'those productive of the best possible consequences' (ibid. p.34) – i.e., consequentialism. The consequentialist, however, might reply by distinguishing between welfare and the good. Ross seems to suppose that providing x units of welfare to someone (i.e., x units of good *for them*) produces x units of good simpliciter. But the consequentialist can deny this and claim that units of welfare are independent of units of good. She might, for instance, hold the following view of the example: keeping the promise to A has value x; providing 1,000 units of welfare to A has value y; providing 1001 units of welfare to B has value z. She then asks which is greater: x+y or z? And the answer is determined by the particulars of the case. (Ross's oversight is made explicit in his discussion (ibid. p.35) of a variant of the above example, when he starts by speaking of a disparity in the provision of 'units of good for A', and ends by speaking of 'a disparity of good' simpliciter.)^{vi}

We are not consequentialists, but we think that consequentialism can cover more of the practical territory than some of its opponents suppose. According to what

we'll call 'simple consequentialism', you only have reason to perform a particular act if doing so will produce a better state of the world than doing nothing (where the value of the act itself counts as part of the value produced by it); and what you have most reason to do is perform the act that will maximize value^{vii}. The content of your reasons, however, need not mention value – the simple consequentialist could adopt something akin to Scanlon's (1998) 'buck-passing' view, and claim that although value determines reason strength, it need not enter reason content. Your reason to put on your coat is that it's cold; the strength of this reason is a matter of the value of the act relative to other alternatives available to you. But what determines the value of an act?

The hedonistic utilitarian holds, perhaps, the simplest view: pleasure is the only good (and pain the only bad). This might be construed as denying that welfare is relevant to value. Alternatively, however, the claim that pleasure is the only good can be seen as emerging from the conjunction of: (1) pleasure is the only thing beneficial; and (2) welfare is the only good. In addition, on this approach, the hedonistic utilitarian holds that (3) pleasure is always beneficial, and welfare is always good.

The simple consequentialist need not, of course, be a hedonistic utilitarian. The latter has no room for the thought that one distribution of welfare is better than another; the simple consequentialist, on the other hand, can incorporate, say, the view that justice – in the sense of distributing welfare in accord with desert – is itself good (see, for instance, Ross, *ibid.* pp.26-27, 138). This adds to the list of goods, of course, in denial of (2); and (3) fails also, since a benefit going to someone who deserves harm, for example, is bad.

(In passing, we would note that is important for our purposes below that in order even to *raise* the issue of whether one distribution is better than another, there must be something to be distributed – namely, welfare.)

Setting aside the issue of (1), let's turn to consider reason strength. According to simple consequentialism, since each of us only has reason to do things that increase the good, you only have reason to pursue a benefit (for yourself or others) if conferral of the benefit in question will increase the good; and the strength of your reason to act in pursuit of a benefit is proportional to the amount of good the act will produce. You might, say, have most reason to pursue your own welfare in some circumstance, on this view, but only if the state you produce in that pursuit is the best (in the sense of maximizing value) you can achieve^{viii}. Welfare is relevant to reasons here, but only indirectly: while welfare is relevant to the value of a state, the strength of your reason(s) to produce that state is proportional only to its value.

We now have three notions in play: value, welfare, and reason strength. As we have just seen, the simple consequentialist sees reason strength as a matter only of value. But, of course, there are other possibilities. According to the

normative egoist, for instance, value is irrelevant to reason strength – rather, the strength of your reason to do something is a matter only of how much it would benefit you. We try to occupy an intermediate position between these two extremes. Like the advocate of egoism, we see welfare as playing a direct role in our practical reasons. And like simple consequentialism, we also see the good as playing such a role.

Here's one component of our position: you have reason to perform some act only if, in comparison to doing nothing, either so acting will give rise to benefit (for someone or something) or good, or both. If your contemplated act is of no benefit or good (relative to doing nothing), you have no reason to perform it.^{ix} And, *pace* simple consequentialism, we think that there are occasions on which you have more reason to perform an act that is worse (in terms of value) than some other; and you do not always have most reason to maximize the good.

On our view, the strength of your personal and special reasons (respectively, reasons to benefit yourself and reasons to benefit those to whom you stand in some special relationship) may exceed whatever contribution the conferral of such benefits would make to the general good. You may have reason to benefit someone to whom you bear no special relationship, but the strength of such a reason, we claim, is a function only of how much good would be accomplished – when benefitting such people you should, other things equal, distribute the benefits so as to do the most good. When it comes to personal and special reasons, however, we contend that their strength can outstrip value.

What are the special relationships we have in mind? On the one hand, there are the ties you have with your friends and family and so forth. On the other hand, there are the ties you can also have with strangers in virtue of such things as making promises^x, accepting benefits, or inflicting harm (see Ross, 1930, p.21). Consider, for instance, Ross's 'duties of reparation' (ibid.). In our framework, these become special reasons to benefit those whom you have unjustifiably harmed – that is, you may have a reason to benefit someone that you have unjustifiably harmed the strength of which is greater than the value of the act of reparation would warrant. The simple consequentialist could incorporate reparational thoughts along the following lines: the world goes better if wrongdoers make reparation to their victims themselves. But we're not sure that this goes far enough – for example, it would require that if you could ensure more such reparation by failing to make reparation yourself, that is what you should do.

We see the simple consequentialist as confronting at least a *prima facie* dilemma (among other difficulties). Either she acknowledges the existence of benefits and harms or she doesn't. If she doesn't, she is in the position of having to deny even the possibility of a debate over distributional concerns (see above). If she does acknowledge their existence, she has to counteract the plausible thought that, on occasion at least, we have personal and special reasons the strength of which is disproportionate to the value of the acts in question. For instance, a state in

which I benefit may well be equally valuable to one in which you do. Let us suppose it so. Then, as far as value goes, you have as much reason to benefit me as you do to benefit yourself. But does such parity of value rule out the possibility that you have more reason to benefit yourself than to benefit me? Or consider a case in which your receiving some benefit would make the world worse, might you not have some reason to pursue it? (Consider the escape of a justly imprisoned felon.)

V

Constraints

Vis-à-vis personal and special reasons, then, we disagree with the simple consequentialist. But what about those moral restrictions that are now standardly known as constraints – rejected by consequentialists but accepted by many traditional deontologists? (See, e.g., Alexander & Moore, 2008.) A constraint is a prohibition against harming people, even in pursuit of good ends – even, indeed, to prevent a greater amount of the very kind of harm that is prohibited by the constraint in question. Proponents of constraints differ in how stringent they take them to be. Some think them absolute: Roman Catholic moral theology, for example, has traditionally held that one may never intentionally kill an innocent person – even to prevent others killing many more innocents. Other deontologists have held that, though constraints are always a significant consideration, they may be overridden, especially if that is the only way to avoid catastrophe. Constraints that are seen, in this latter fashion, as having some threshold beyond which the bad consequences of adhering to them dictate that we should violate them, are known as threshold constraints. How, if at all, do we accommodate constraints within our framework?

First, we need an account of constraints in terms of reasons. While, in our view, there is no sharp division of practical reasons into the moral and the non-moral, examples can be provided that are clearly on one side or the other. Your reason to choose a peach over an apple – that the former is sweeter – is non-moral. Your reason to give to Oxfam, on the other hand – that doing so will reduce innocent suffering – is moral. Or suppose that you promised to repay a debt on Thursday; this fact is a moral reason to do so. What about your reasons to favour your friends? Some object to the idea that any of these are moral on the grounds that there is something less than ideal about doing things for friends out of a sense of obligation. But that is to confuse the issue of reason with that of motivation: it's quite possible to act on a moral reason – to refuse to betray a friend, say – out of affection.

The moral reasons just mentioned fall within two categories: some are associated with promoting the good, and some with special ties. And some personal reasons may also be moral. But the advocate of constraints can be seen as claiming that there is a further category of moral reasons. Suppose, for

instance, there were a constraint against killing the innocent. Then, whether this constraint be absolute or threshold, there would be at least one possible occasion on which the strength of your moral reason not to kill an innocent stranger would be greater than that which would correspond to the disvalue of the killing. On such occasions, your killing the innocent stranger would be bad, but not doing so would be worse (in the sense of being more disvaluable – more innocents would be killed by others, say); yet the constraint would dictate that you have more moral reason not to kill. So the strength of your moral reason not to kill an innocent stranger does not vary only with the badness of doing so – call such a reason a “constraining reason”.

Our view is that the strength of your reason to perform a given act is a matter only of how much benefit or harm to yourself or those to whom you stand in special relations, and/or how much good or bad, would result^{xi}. Thus we leave no room for constraining reasons. The badness of harm cannot do the job: the constraint *violator* in the previous paragraph would ensure that less bad comes about. And special relations do not help either. Constraints are independent of relations such as friendship, and the only relation to strangers that might appear relevant is the tie that grounds reasons for reparation. But that tie results from the infliction of *unjustified* harm. Even if you harm someone in the course of doing greater good, the advocate of constraints will still see this as unjustified. But that is to beg the question against the simple consequentialist (and us, in this case), who sees the doing of good as justifying the harm.^{xii}

Constraints, then, from our perspective, would require a further primitive concerning reason strength – they cannot be accommodated by appeal to welfare or value. We see little prospect of a plausible rationale for including such an added extra.^{xiii} But arguing this case in detail here would take us too far afield.^{xiv}

VI

Weighing Reasons

We have been speaking of the determinants of reason strength, but have said neither how reasons are to be weighed against one another, nor how our views on reason strength relate to holism.

Concerning the determinants of reason strength, matters are notoriously complex, both epistemically and metaphysically. Is there, for instance, a fact of the matter concerning rankings of the options available to you in terms of value and welfare? If so, how do we know the rankings? How do the rankings relate to reason strength? Some cases are clear-cut – for instance, you have more reason than not to get your shoes wet in order to save an infant. Other cases are much trickier: under what conditions, say, do you have more reason than not to

sacrifice your own life? We haven't space to comment further on such issues here, however – rather, we'll return to holism.

Our view appears consistent with it: reasons have strength only within a context, since an agent's situation determines how beneficial and/or valuable an act would be relative to her other options. Vary the context and reason strength may vary too – valence switching being an extreme case of this. However, Dancy takes the holistic conception of weighing reasons a step further. He considers (2004, p.190) what he dubs the 'kitchen scales conceptions of rationality'. It makes two claims:

1. The weight of a rational element or reason is not affected by the weight, nor indeed even by the presence, of any other rational element.
2. Once one has assessed the separate weight of each element, evaluative judgement consists of adding up the pros and cons to see which side is weightier.

Dancy rejects both components. We agree with him vis-à-vis the first. Holism entails its denial (although one need not be a holist to reject it: merely acknowledging that a reason's strength can vary with circumstances suffices). But we're not inclined to abandon the second component. One might take the view that holism involves an interactive effect that rules out, in a given circumstance, separating the reasons from one another and assessing their weights individually. Alternatively, on what one might call a 'Moorean' holism about reasons (after Moore's approach to assessing the value of a whole – see, e.g., Moore, 1903/1966, p. 28), although the weights of reasons can be separately assessed, their overall weight is not the sum of their individual weights. Both these approaches entail the rejection of (2). But there is a third holist approach – the one we favour – that sits quite comfortably with accepting (2) while rejecting (1). The idea is simply that, although the strengths of reasons are influenced by their surroundings, nevertheless, in any given circumstance, one's reasons for action have strengths that can be (roughly, at least) summed and compared to yield verdicts about what one has most reason to do.

But here's a case that might tell against our account of holism in favour of the interactive view. People are often somewhat incredulous at the very idea of particularism – 'I'll give you a moral principle: inflicting pain on babies for fun is wrong!'^{xv} We agree. Contrast this case with the following: driving your car for fun is wrong. Why? Because, it might be claimed, it harms the innocent by contributing to the pollution of the air they breathe. In each case the agent will derive enjoyment from his actions, and we'll suppose, for the sake of argument, that his enjoyment constitutes a benefit and a reason for him to act in both cases^{xvi}. But each also has a reason not to act (whether or not he sees this): it'll harm the innocent. So what distinguishes the two cases? Aside from differences in how much harm each might involve, in the torture case the sadist takes enjoyment *in* inflicting the harm. In the car case, you can weigh the enjoyment on

the one hand against the harm on the other – for a given level of harm, perhaps if the enjoyment is sufficiently great you have more reason to drive than not. But the torture case has a different structure: increasing the amount of enjoyment for a given level of harm *decreases* the strength of the reason for the sadist to torture – assuming that increasing the pleasure the sadist takes in inflicting the suffering makes the situation worse, and to such a degree that it outweighs the increased benefit of the extra pleasure. There is an interactive effect: the enjoyment cannot be considered independently of the harm.

So how can we reconcile this case with our non-interactive holism? Consider the reasons confronting the sadist. We're assuming that the sadist's benefit (in the form of enjoyment) weighs on one side of the reasons scale; and we claim that the sum of (1) the badness of the harm he inflicts, and (2) that of the enjoyment he takes in inflicting it, weighs on the other – the interactive effect is situated, as it were, *within* (2). But aren't we counting the enjoyment twice? Yes, but in different ways. It's important, as we've seen, to distinguish sharply between benefits and harms on the one hand, and value on the other. Here we have a case where benefit is weighed against badness on the reasons scale: whatever the benefit's reason-giving weight, it is far outweighed by the sum of (1) and (2) (and the greater the enjoyment, the larger the disparity).

VII

Conclusion

All we have proposed in the latter part of the paper is a framework for structuring the practical landscape. We have said, for example, little about what is valuable or beneficial, nor about the details of weighing reasons. But the framework does have some bite. We categorize reasons, for instance. And our framework leaves out constraints. We also place further bounds on holism beyond our suggestions concerning universal reasons. Recall that '[h]olism maintains that anything whatever might make a practical difference, or provide a reason, if the circumstances were suitable' (Dancy, 2004, p.111). The circumstances being 'suitable', we claim, requires that the act in question be of benefit or value.

Our attempt to systematize practical reason succeeds, then, in some respects. But it is limited. No doubt Dancy will see our view as far too systematic. We see ourselves, however, as in sympathy with his general approach. And, more broadly, we take it that he shares with us a view of morality as always reason-giving, and thus rejects the view of morality implicit in the following passage from Singer (1999, pp. 308-309)^{xvii}:

[Nagel and I] were discussing 'Famine, affluence and morality', and Nagel was unable to accept that morality could be so demanding. But eventually it emerged that he was assuming that if morality did demand that we give so much to famine relief, then there must be overriding reason to do so. I

was making no such assumption. On my view, I could recognize that if I were totally committed to doing what I ought to do, I would give away my wealth up to the point indicated in my article; but at the same time I may, without any irrationality, choose to be less than totally committed to doing what I ought to do. My own interests, or those of my family, may counteract the demands of morality to some degree, and I may think it reasonable to give in to them, while recognizing that it is morally wrong for me to do so. Once Nagel and I realized that we held these distinct understandings of morality, the practical difference between Nagel and myself over the demandingness of morality became less acute.

On Singer's view, then, morality and practical reasons are, to some extent at least, in different camps. Singer sees morality, perhaps, as akin in some respects to the law – it is true, for example, that you do not always have overriding reason to obey the law of the land. Indeed, it may be asked whether you have any reason to do so, and Singer would seem to invite the general question, “Do I have reason to be moral?” Within our framework this question becomes, “Do I have reason to do what I have moral reason to do?” – a non-question if there are moral reasons. So for us the real question here is: “Are there any moral reasons?” – a different question from the first. For example, some Hobbesians might claim that, at least on occasion, we each have reason to be “moral” because it is in our self-interest^{xviii}. But they might deny that there are any moral reasons on the grounds that egoistic reasons to be moral are not moral reasons.

Dancy would join with us, we take it, both in rejecting the legal model of morality (see, e.g., 2004, p. 83), and affirming the existence of moral reasons.

References

- Alexander, L. and Moore, M. (2008) 'Deontological Ethics', *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, E. N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/ethics-deontological/>
- Butler, J. (1729/1970) *Fifteen Sermons* (ed. T.A. Roberts) London: SPCK
- Copp, D. (2006) *Oxford Handbook of Ethical Theory* Oxford: Oxford University Press
- Dancy, J. (1993) *Moral Reasons* Oxford: Blackwell
- Dancy, J. (2004) *Ethics Without Principles* Oxford: Clarendon Press
- Davidson, D. (1980) *Essays on Actions and Events* Oxford: Clarendon Press
- Gauthier, D. (1967) 'Morality and Advantage' *Philosophical Review* 76(4), 460-475
- Hare, R.M. (1963) *Freedom and Reason* Oxford: Clarendon Press
- Hooker, B. and Little, M. (eds.) (2000) *Moral Particularism* Oxford: Clarendon Press.
- Kagan, S. (1984) 'Does Consequentialism Demand Too Much?' *Philosophy and Public Affairs* 13, 239-54
- Lance, M. & Little, M. (ms). 'Mad dogs and Englishmen: Moral Valence, Defeasibility, and Privileged Conditions'
- McNaughton, D. and Rawling, P. (2000) 'Unprincipled Ethics' in Hooker and Little, 256-275
- McNaughton, D. and Rawling, P. (2006) 'Deontology' in Copp, 424-458
- Moore, G.E. (1903/1966) *Principia Ethica* Cambridge: Cambridge University Press (page nos. refer to the 1966 paperback edition)
- Parfit, D. (2011) *On What Matters* Oxford: Oxford University Press
- Portmore, D. (2009) 'Consequentializing' *Philosophy Compass* IV(2) 329-347
- Ross, W.D. (1930) *The Right and the Good* Oxford: Clarendon Press
- Ross, W.D. (1939) *The Foundations of Ethics* Oxford: Clarendon Press

Scanlon, T. (1998) *What We Owe to Each Other* Cambridge, Massachusetts: Harvard University Press

Williams, B.A.O. (1981) 'Internal and External Reasons' in Williams, 101-113.

Williams, B.A.O. (1981) *Moral Luck* Cambridge: Cambridge University Press.

Wittgenstein, L. (1953) *Philosophical Investigations* G.E.M. Anscombe and R. Rhees (eds.), G.E.M. Anscombe (trans.), Oxford: Blackwell.

ⁱ Many thanks to Brad Hooker, Maggie Little, and various audiences, for helpful comments.

ⁱⁱ Rejecting the simple view means abandoning a simple account of the supervenience of practical reasons on the non-normative: on the simple view, there cannot be a change in what you have reason to do without a change in the non-normative facts because your reasons are among the non-normative facts. If we allow normative facts at tier one, this, obviously, no longer holds. However, supervenience still holds provided that, to take the current case as an illustration, harm and desert have non-normative supervenience bases – that is, there cannot be two cases that differ in terms of harm or desert without some non-normative difference. This seems to us obviously true.

ⁱⁱⁱ Fans of *Yes, Minister* will recall another possible case – Sir Humphrey's oft repeated response to the Minister's proposals: 'That would be very courageous, Minister.'

^{iv} Perhaps the advocate of compact reasons would maintain that the reason here is only that your act would cause someone harm, with lack of desert featuring in a different role. But for simplicity we'll ignore this somewhat flexible boundary between reasons and other considerations from here on since it doesn't affect our line of argument – rather than "universal reason", sticklers might prefer the term "universal consideration".

^v As Dancy acknowledges, his account here is 'highly contentious' (ibid. p. 46). One obvious difficulty (which, to be fair, he does his best to address) is how the account is to cope with cases of false belief – e.g., explaining your wearing your coat on a hot day in the mistaken belief that it's cold: that it's a *hot* day appears unlikely to do the trick.

^{vi} G. E. Moore is one of Ross's targets. Moore's theory, Ross claims, 'says, in effect, that the only morally significant relation in which my neighbours stand to me is that of being the possible beneficiaries of my action' (ibid. p.19). We're not sure that's the case. But even if it were, it does not follow that Moore has to equate the value of a state with the sum of the benefits it includes. And Moore does not do this (although he does not adopt the approach in our previous paragraph either: see e.g. Moore, 1903, p.214 ff).

^{vii} This definition of simple consequentialism differs from earlier definitions we have proposed – e.g., in McNaughton and Rawling, 2006. Also, it should be noted that we are setting aside recent moves to "consequentialize" all moral theories. Consider, for example, Portmore (2009). His leading idea is that any theory that determines the deontic status of an act 'by how its outcome ranks relative to those of the available alternatives on *some* evaluative ranking' (ibid. p.330; italics ours) is a form of consequentialism. The evaluative ranking here can be, for example, egoist, so that egoism is a form of consequentialism on this account (pp.334-335). For Portmore, the appeal

of consequentialism, as he defines it, rests on the thought (roughly) that outcomes can be ranked in accord with what we have reason to prefer, and we should perform the act at the top of the ranking. Disputes then arise over what we in fact have reason to prefer – is it, for instance, what would be good for me, or what would be good simpliciter? For us, by contrast, the appeal of consequentialism, as we define it, is the thought that the good simpliciter plays a central role in practical reason.

^{viii} Note that on our account, consequentialism is not a doctrine solely concerned with what one is morally required to do – your reason to pursue your own welfare here, for example, need not be a moral one. Parfit (2011, vol. I, p. 168) attributes something like our view of consequentialism to Sidgwick – at least to the extent that he sees Sidgwick’s version of consequentialism as concerned only with impartially assessed reasons. But Parfit then goes on to say that ‘this kind of Consequentialism may be better regarded, not as a moral view, but as ... an external rival to morality’ (ibid.). We disagree, but this is a result of many further differences between our view and his that we have not the space here to address.

^{ix} Two aspects of this claim might initially appear puzzling. First, why the comparison to doing nothing? This is to accommodate cases in which, if you act either harm or badness will result, but if you do nothing even more harm will arise, or things will go even worse. We do not want to rule out your having reason to act in such unfortunate circumstances. (Implicit in our view, then, is the thought that reducing harm counts as producing benefit; and reducing badness counts as producing good.) Of course, the notion of ‘doing nothing’ is tricky, and we certainly don’t want to enter the debates about acts and omissions. What it would be to do nothing will, however (we hope), be clear in any given case. Second, why not strengthen the claim from ‘only if’ to ‘if and only if’? Well, suppose you could benefit a justly imprisoned felon by helping him escape. His benefit notwithstanding, you may have no reason to. However, if an act would do some good (in comparison to doing nothing), then you do have some reason to perform it.

^x Some might claim that you have reason to keep a promise even though it will benefit no one and not do any good (and the same might apply to refraining from stealing). If this is correct, it is a counterexample to our view. Consider, for instance, a confidential death-bed promise to do something posthumously for the promisee that, as things turn out, will dishonour the promisee’s memory. Is there room on our view for the thought that, even if you have most reason not to keep the promise, nevertheless you do have some reason to do so? It is open to us to maintain that the very keeping of a promise, regardless of its consequences, can be either beneficial to the promisee or valuable or both. We’re not sure about this. But in our view, of course, whether you do have some reason to keep the promise is precisely a matter of whether it will benefit anyone (or thing) or do some good.

^{xi} What is it for an act to be morally obligatory on our account? It’s complicated, so for present purposes let’s make the following unrealistic simplifying assumptions: the agent knows what her practical reasons are, what they are reason to do, whether or not they are moral reasons, and their relative strengths; and there’s exactly one act that she has most reason to perform, and exactly one act that she has most moral reason to perform. We then tentatively suggest that you are morally required to phi if and only if:

- (1) phi-ing is what you have most reason to do overall;
- (2) the preponderance of your reasons to phi are moral; and
- (3) you have more moral reason to phi than you have moral reason to do anything else.

^{xii} What about personal reasons? Admittedly, violating a constraint might harm the violator (consider the psychological trauma, for example). But such harms to the violator do not provide her with *moral* reasons not to violate, and thus are not, presumably, what the advocate of constraints has in mind as grounding them.

^{xiii} In many ways our view is similar to that of Ross (1930, 1939). But there are also many points of difference. For example, we see practical matters in terms of reasons, moral and otherwise; Ross, by contrast, sees practical reason as comprising only moral obligations. This leads to many further differences – for example, Ross (1930, pp.24-26) claims that we have a ‘duty to produce pleasure for ourselves’; we think only that you often have a reason to pursue your own pleasure. Constraints may constitute another point of difference. Ross speaks of the ‘duty of non-maleficence’ (ibid. pp.21-22) in a way that may imply he thinks of it as (what we would now call) a constraint against injuring others. However, one of Ross’s main concerns in contrasting non-maleficence with beneficence is to emphasize that the former is ‘a duty of a more stringent character’; but this point can be accommodated without any appeal to constraints once we distinguish (as Ross does not do as sharply as he might: see section IV) between benefits and harms on the one hand and value on the other: the bad of injury outweighs the good of benefit.

^{xiv} For instance, one of the issues that we lack the space to address here are arguments to the effect that admitting personal reasons without constraints yields counterintuitive results – see, e.g., Kagan (1984) and McNaughton and Rawling (2006).

^{xv} David Copp pointed out to us that the principle we have in mind might be more precisely worded as: inflicting pain on babies is wrong if the only potential reason in favour of doing so is your own enjoyment.

^{xvi} Note that we are neither committed to the view that enjoyment is always beneficial, nor to the view that benefits always provide reasons.

^{xvii} Thanks to Doug Portmore for alerting us to this passage.

^{xviii} Although early Gauthier (1967, p.470) denies this: ‘The individual who needs a reason for being moral which is not itself a moral reason cannot have it’ (p.470).