

## Unprincipled Ethics<sup>1</sup>

David McNaughton and Piers Rawling

The term 'ethical particularist' has sometimes been used, in a broad and loose way, as a label for anyone who expresses hostility to the view that a decision about what we ought to do in some particular case can be mechanically 'read off' from a general moral principle or principles. Rather, it is urged, a correct moral verdict can only be reached by paying close attention to the individual case -- to what differentiates it from other cases as much as what it has in common with them. As well as an understanding of the correct moral principles, we need fine judgement, sensitivity and even something approaching a perceptual capacity to appreciate the saliences of the circumstances in which we find ourselves. Particularism in this broad sense, which claims that a grasp of moral principles is insufficient for the correct moral appreciation of the particular case, has won many adherents in recent years. We will call this view, with which we agree, moral verdict particularism. It is a position explicitly held by intuitionists<sup>2</sup> (and no doubt by some other moral theorists) as can be seen from Rawls' classic definition:

Intuitionist theories, then, have two features: first they consist of a plurality of first principles which may conflict to give contrary directives in particular types of cases: and second, they include no explicit method, no priority rules, for weighing these principles against one another.<sup>3</sup>

---

<sup>1</sup>The authors would like to thank Roger Crisp, Chris Daly, Jonathan Dancy, Andre Gallois, Eve Garrard, Brad Hooker, Philip Stratton-Lake, Stephen Wilkinson, members of the Philosophy Department at Manchester University, and participants at the NEH Institute, *Ethics: Principles or Practices?* And Piers Rawling would like to thank the University of Missouri Research Board, and the University of Missouri-St Louis Research Awards for financial support.

<sup>2</sup> We are not using the term 'intuitionism' to cover all moral theories that take the epistemological view that intuition tells us what is right.

<sup>3</sup> J. Rawls, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, 1971), 34.

But the term particularism is also used to refer to more controversial claims, one of which is the radical claim that there are *no* moral principles. It is not that general principles are insufficient to guide us in our consideration of the particular case -- they simply do not exist. It is this more radical view we consider here. It arose, in part, as a response to Hare, who coined the term 'particularist'.<sup>4</sup> We shall examine this response to Hare, and go on to argue that it does not give us reason to abandon all conceptions of a moral principle (as Dancy sometimes does).<sup>5</sup> We urge a more moderate position that acknowledges the truth and importance of certain weak moral principles.<sup>6</sup>

#### I TWO CONCEPTIONS OF A MORAL PRINCIPLE

As Hare defines particularism,<sup>7</sup> it is, at least in its 'extremest' form, the rejection of the supervenience of the moral upon the non-moral -- that is, the rejection of the notion that if two circumstances are alike in all non-moral respects, then they must be alike in all moral respects. As Dancy uses the term, however, it means something different. Since Dancy sets up his position in conscious opposition to classical intuitionism, of the kind championed by W. D. Ross, we need to briefly expound Ross' notion of a *prima facie* or (better) *pro tanto* duty.<sup>8</sup>

Suppose the fact that an act would be just is always a reason in favour of doing it. Then justice is a *universally right-making* feature of acts. Or, as Ross puts it, there is a *prima facie* duty to be just. This is not to say that all just acts are right (there can be occasions when justice is outweighed by other considerations), but, rather, the fact that an act would be just always counts for that act, deontically speaking. Similarly, if the fact that an act would

---

<sup>4</sup> R. M. Hare, *Freedom and Reason* (Oxford: Clarendon Press, 1963), 18

<sup>5</sup> J. Dancy, *Moral Reasons* (Oxford: Blackwell, 1993).

<sup>6</sup> For a classification of particularisms, see pp. 36-41 of R. Audi, 'Moderate Intuitionism and the Epistemology of Moral Judgement', *Ethical Theory and Moral Practice*, 1 (1998), 15-44.

<sup>7</sup> Hare (*Freedom and Reason*, 18-19).

<sup>8</sup> W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930).

be a betrayal of an innocent friend always counts against it, then the betrayal of innocent friends is universally wrong-making (which, to reiterate, is not to say that it is universally wrong). To employ a partial analogy with the chemical notion of valence, we have the following pair of principles: justice has a universally positive valence; whereas the betrayal of innocent friends has a universally negative valence. (A valence of zero corresponds to moral irrelevance in the circumstance). Such principles, of course, are not sufficient for determining the deontic status of an act, since they may conflict (we will refer to such principles as *weak*).

Dancy's particularism (which is clearly consistent with supervenience) combines moral verdict particularism (the claim that there are no absolute principles for determining the overall deontic status of an act) with the more startling claim that there are no weak principles (that is, there are no properties, apart from the thin moral properties right, wrong, etc., that have universally and counterfactually invariant valence). We dub this latter claim, with which we disagree, moral valence particularism.

It is important to note that these claims are independent of one another. Hare's 'universal prescriptivism' commits him, among other things, to the 'thesis of universalizability'.<sup>9</sup> Where 'M' abbreviates any moral term (such as 'good', 'right', 'just', or 'courageous'), we have:

(U) If an act or circumstance is M, then any relevantly similar act or circumstance is M (where relevant similarity is a matter of sharing some set of descriptive (i.e., non-moral) properties).

When we judge an act, say, morally proscribed, we do so because it possesses non-moral properties, say,  $P_1$ - $P_n$ . And we are then rationally committed to the following *strong* moral principle:

(P) Any action with non-moral properties  $P_1$ - $P_n$  is morally proscribed.

Hare thinks that all moral judgements rest on principles of this form. Thus his position is inconsistent with moral verdict particularism. But this does not commit him to the view that  $P_1$  (say) has invariant moral valence.

Dancy sometimes speaks<sup>10</sup> as if universalizability commits one to thinking well of arguments such as the following. Suppose one accepts that a person's lack of direct benefits from the local schools (she has no children) is

---

<sup>9</sup> R. M. Hare, *The Language of Morals* (Oxford: Clarendon Press, 1952) and *Freedom and Reason*.

irrelevant to the issue of whether or not she should contribute to them, then (we have heard it argued) one must accept that a person's lack of direct benefits from the departmental coffee fund (she does not drink coffee) is irrelevant to the issue of whether she should contribute to it. But universalizability entails no such commitment: it is quite consistent with (U) to maintain that a lack of direct benefit can have varying valence.

It is only the extremest of anti-particularists who would maintain that the valence of all properties is invariant; and we doubt that there are any such extremists. It is certainly not a view held by the intuitionists. Urmson, having quoted Rawls' definition with approval, gives a second definition in terms of reasons.

Some fact will be a primary reason for acting in a certain way if that fact's obtaining is always a reason for acting in that way, though not necessarily a sufficient reason. Some fact will be a secondary reason for acting in a certain way if that fact's obtaining brings about some fact which is a primary reason for acting in that way.<sup>11</sup>

If the fact that an act is just, to revert to our earlier example, is a primary reason for performing it, then justice has universal positive valence, and we have the weak principle:

(J) In so far as an act is just, this counts in its favour.

But there are plenty of examples of what Urmson would call secondary reasons -- considerations with varying valence. The fact that some act would be illegal is often a reason against doing it, but we need not assume that the fact that an act is illegal will always count against doing it. It may be in virtue of other things that are generally true of illegal acts that the fact that an act is against the law is a strike against it. There might be special cases where its

---

<sup>10</sup> Dancy (*Moral Reasons* 88-90).

<sup>11</sup> J. O. Urmson, 'A Defence of Intuitionism', *Proceedings of The Aristotelian Society* ns, lxxv (1975), 112. The notion of bringing about some fact is not clear. We try to say a little more in connection with Ross' theory in part III.

being illegal is morally neutral or even counts in favour of an act, as perhaps was true in the case of the Nuremberg laws in Nazi Germany. Or, to take a different example, that someone is strong can sometimes be a morally acceptable reason for giving them a job, and sometimes not. Justice may well be thought to require that we discriminate between job applicants only where the properties in question are relevant to the job. But which properties are relevant to the job will vary from case to case. Endorsement of (U) is consistent with sharing these judgments.

To put matters in a different way: relevant similarity in Hare's sense can be interpreted holistically. That is, the valence of non-moral features in a principle such as (P) can be sensitive to context --  $P_1$  in the context of  $P_2$ - $P_n$  might have negative valence, but it might have positive valence in another context. However, unlike Dancy, Hare does think that there are some terms (other than the thin moral terms) that have constant valence: the 'secondarily evaluative' terms, such as 'courageous'. These are, roughly speaking, coincident with Williams' 'thick' moral terms.<sup>12</sup> And, modulo certain wrinkles that will appear in section III, we concur with Hare that thick moral terms are univalent.

However, the significance of this claim is very different on our view from its significance on Hare's. When we say that justice has universally positive valence, we mean that it always counts in favour of an act that it would be just; and that just acts are commendable in so far as they are just. Justice is an important feature that must be weighed up with the other considerations present in determining what to do. Hare simply means that to declare an act just is, on pain of violating the current prescriptive meaning of the term 'just', to commend it for being so -- it is part of his project to avoid commitment to properties such as commendability.

Indeed, Hare's model of morality is not one based on the notion of weighing reasons against one another. The weighing of reasons is a crucial aspect of intuitionism and other theories that endorse moral verdict particularism, and comports with our ethical intuitions. Reasons for action are not overriding, of course: there will

---

<sup>12</sup> B. Williams, *Ethics and the Limits of Philosophy* (London: Fontana Press/Collins, 1985), 129, 140.

typically be reasons for and against any action. On our view, (rational) agents confronted with (difficult) choices seek out reasons for and against their various options and weigh them up.<sup>13</sup>

Hare's model is very different. We noted above that the valence of any non-moral feature can vary from context to context on Hare's view. But the view is also consistent with non-moral features having no moral valence at all. The model simply has it that the bearer of some set of non-moral features has some particular moral property; it is no essential part of the model that the features have (context dependent) moral valences taken one by one in such a way as to facilitate the evaluation of their individual contributions to the whole.

Among other things, the contribution of thick properties is rendered opaque. On our view it is helpful, in deciding whether to perform an act, to determine whether it would be, say, just. Hare, on the other hand, seems concerned only to look to non-moral properties in making such a decision. On his view, the following principle:

(Q) Any action with non-moral properties  $Q_1$ - $Q_m$  is just

can tell you that an act is just, but it is only some set of the act's non-moral properties that feature in the determination of its rightness (and, as far as we can tell, this set may or may not include various of properties  $Q_1$ -

---

<sup>13</sup> In the case of reasons, not only is there no weighing algorithm, but, of course, agents can fail to know the facts that constitute their reasons (or know the facts and fail to realize that they are reasons). The relation between an agent's reasons and her obligations is, we think, complex -- and we shall not discuss it further here, beyond noting the failure of a naive equation of what an agent ought to do with what she has most reason to do. This is brought out by the following adaptation of an example due to H. A. Prichard (in 'Duty and Ignorance of Fact' in *Moral Obligation and Duty and Interest* (Oxford: Oxford University Press, 1968, 29). Do you have a reason to slow down before entering a main road? Legal issues and example-setting aside, if there is no traffic coming, you have no reason to slow down. Yet we concur with Prichard in thinking that you should slow down, even if there is no traffic coming.

Q<sub>m</sub>).<sup>14</sup> So we have the possibility of commending an act in so far as it is just, but not commending it overall. The relevant imperative here would appear to be something along the lines of: 'In so far as an act is just, do it'. But it is unclear what it would be to follow a command so qualified. It is, of course, not the conditional command: 'If an act is just, do it'. Rather one might be commanded to do the act in so far as it is just, but not do it overall -- which makes little sense.

There is a certain irony here: for the particularist (such as Dancy) who denies even that any morally thick properties have invariant valence, a similar issue arises. On this view, having determined that an act is just, we then have to ask the further question: does justice count for or against here, or is it morally irrelevant?

So far we have classified Dancy's position as being opposed to classical intuitionism, as defined by Rawls. If intuitionism is understood in a rather wider sense, however, as the view that there is an irreducible multiplicity of morally relevant considerations that have to be weighed to reach a moral verdict, then even a radical particularism such as Dancy's can be classified as an intuitionist theory. It will be simpler, for purposes of exposition, to think of the question of whether there are properties with invariant moral valence as a family squabble within intuitionism. All intuitionists are moral verdict particularists, but they differ in their attitude to moral valence. We shall distinguish three positions. The first variety of intuitionism claims that there are non-trivial cases of universally and counterfactually invariant valence, and these all involve thick moral properties. We shall call this thick intuitionism.<sup>15</sup> (We speak of counterfactual invariance because the claim is not merely that a feature happens to have

---

<sup>14</sup> We have our doubts as to whether Hare, in *The Language of Morals*, can maintain his imperative model of prescription here. (Q) does not entail, of course, that any action with non-moral properties Q<sub>1</sub>-Q<sub>m</sub> has any particular thin moral property. Rather, the prescriptive meaning of 'just' presumably entails merely: 'In so far as an act is just, commend it' as distinct from 'Commend all just acts'

<sup>15</sup> We have not the space here to discuss the issue of whether all morally thick concepts have invariant valence. Blackburn, 'Through Thick and Thin', *Proceedings of The Aristotelian Society* suppl. vol. lxvi, (1992), .285-299, for example, denies that 'lewd' is always used with disapprobation -- someone might declare: 'What we need here is an injection of lewdness into the proceedings.' And lewdness might be considered morally thick. We agree that the

invariant valence in all actual acts, but also that it would count in the same direction in any act possessing it.) The straightforward cases are along the lines of principle (J). But there are less straightforward cases that we shall discuss in section III.

The second variety of intuitionism adds some (but not all) non-moral properties to thick intuitionism's list of features with invariant valence. To mark this abundance we shall call this fat intuitionism. And the third variety of intuitionism makes a sweeping subtraction: it claims that the only cases of invariant valence are trivial, such as the claim that murder is wrong, if murder is construed as wrongful killing. On this view, which we shall call thin intuitionism, there are no primary moral reasons in Urmson's sense. This is Dancy's moral valence particularism.

Contemporary particularism arose, in large part, as a response to Hare. We turn now to the arguments of McDowell and Dancy against Hare's conception of morality, and examine their impact on the cases for the varieties of intuitionism.

## II PARTICULARIST ARGUMENTS

Perhaps the first question that arises for Hare's conception is that of how to justify an initial moral judgement. As we saw above, having once proscribed an act because it has non-moral features  $P_1$ - $P_n$ , (U) dictates that we proscribe any future act with those features. But, as Dancy points out, nothing is said about how to justify the initial proscription.<sup>16</sup>

And there is a further difficulty.<sup>17</sup> How are we to rule out the possibility of an act instantiating  $P_1$ - $P_n$  to the required degree together with some countervailing property,  $P_{n+1}$  (a 'defeater'), the combination of which requires us

---

term is culturally encrusted, but we do not see it as having the stature of justice. Furthermore, the relevant issue for our purposes is not that of the usage of the term for purposes of expressing disapproval or otherwise. The issue is whether or not it counts against an act that it is lewd.

<sup>16</sup> Dancy (*Moral Reasons*, 82-3).

<sup>17</sup> Dancy (*Moral Reasons*, 80-81).



to alter our attitude to one of approbation? The obvious response to such an instance would be to add the absence of the countervailing property to  $P_1$ - $P_n$  to yield

(P') Any action with non-moral properties [ $P_1$ - $P_n$  and not  $P_{n+1}$ ] is morally proscribed.

But this is a trivializing manoeuvre (as Hare realises<sup>18</sup>), unless there is reason to declare some stopping place: if all non-moral properties are included, then no two numerically distinct acts share all the relevant properties. And it is precisely the burden of Dancy's plaint here that no such stopping place can be determined in advance. Even if there are principles such as (P) that stop short of listing the entire supervenience base, we have yet to find any, let alone any that are of practical use.

This argument supports all three types of intuitionism, since it attacks a vision of morality that all intuitionists reject (they are all moral verdict particularists). However, it also places a burden upon the advocates of fat or thick intuitionism. The thin intuitionist is radically holistic about moral reasons: she claims that the pertinence and force (i.e., both the sign and the magnitude of the valence) of *any* reason contributing to a moral judgement in a circumstance are dependent in part upon the other reasons there present. Thus the addition of  $P_{n+1}$  above might reverse proscription by modifying the valence of other factors. The challenge to the fat or to the thick intuitionist is: what reason have we to suppose that the valence of any consideration (non-moral or thick) is invariant across all possible cases? Why cannot the positive valence of a consideration in a particular case be reversed (or equal zero) in a new case? We will address this question further in sections III and IV, where we continue our argument for thick intuitionism. But we note here that we are not claiming that (say) justice has the same degree of relevance in all cases -- it can vary in import depending upon other features of the case. But its import never falls to zero or turns negative.

---

<sup>18</sup> Hare (*Freedom and Reason*, 12).

We consider now a pair of arguments due to McDowell.<sup>19</sup> McDowell sees strongly principled views such as Hare's as attempting to ground moral judgement and justification in the following of rules, and draws upon Wittgenstein to argue that this strategy is wrong-headed. To take the standard example of the rule for adding 2, we declare that 1002 is the result of adding 2 to 1000; but what justifies this? Adverting to a further rule is otiose: we would then be required to justify the claim that this further rule is being followed -- and if justification must be grounded in rules, we clearly have a vicious regress. So, if the motivation for adverting to moral principles is to ground moral justification thereby, such views are ill-motivated: there is no such grounding to be had.

Consider the question: 'How do I know that this act is morally proscribed?' In the case where the act is 'relevantly similar' to some previous act that was morally proscribed, Hare's answer is that there is a known strong principle of the form of (P) in force. This advertence to principle, on McDowell's interpretation of Hare, supposedly grounds the present judgement in apodictic fashion. But to make this grounding claim is to misconceive rules as antecedently present 'rails' extending from old judgements to new. (No sequence can dictate further entries.) According to McDowell, Hare supposes that strong moral principles are required in order to ground our moral judgements. But this they cannot do.

There is a response to this challenge, however. The rule-following considerations do no more than undercut a motivation for the appeal to strong principles in ethics. Mathematics *is* principled in a strong sense -- (ungrounded) justification in mathematics is rule-bound. Provided we have mastered the relevant subject matter, we can, in mathematics, derive theorems from axioms in accord with derivation rules (for example). And, according to Hare, we can derive and justify judgements (provided, again, we have mastered the subject matter) from strong moral principles and current conditions. There is the obvious problem of how we arrive at the first judgements of each type. But supposing this overcome, (ungrounded) moral justification would require judging in accord with the strong moral principles.

---

<sup>19</sup> J. McDowell 'Non-cognitivism and Rule-following' in S. Holtzman and C. Leich (eds.) *Wittgenstein: To Follow A Rule* (London: Routledge & Kegan Paul, 1981). Repr. in McDowell *Mind, Value, and Reality* (Cambridge, Mass.: Harvard University Press), 198-218.

According to the opponent of strong moral principles, of course, morality is not rule-bound in this way. But since the Wittgensteinian rule-following considerations apply to mathematics, care must be taken not to assume that they tell the whole story against moral principles.<sup>20</sup> They tell against the possibility of ultimate grounding in both the mathematical and moral spheres. Indeed, if the motivation for Urmson's view that primary moral reasons have unvarying valence is to ground some notion of justification thereby, then this motivation too falls prey to them. But suppose one sees the dispute between Urmson (or Ross) and Hare simply as a matter of debating what (ungrounded) justification *is* in the moral sphere. On this account, they can both acknowledge Wittgensteinian 'vertigo', and we must measure the merits of their positions by other yardsticks.

We turn now to another, related argument of McDowell's. Implicit in Hare's view (recall (Q)) is the notion that our moral classifications coincide neatly with our non-moral classifications: acts can be classified into groups on non-moral criteria, each group being such that the same moral judgement applies to all its members. Against this McDowell takes another Wittgensteinian line:

... supervenience [of the moral on the non-moral, in the present case] leaves open this possibility ... : however long a list we give of items to which a supervening term applies, described in terms of the level supervened upon, there may be no way, expressible at the level supervened upon, of grouping just such items together. Hence there need be no possibility of mastering, in a way that would enable one to go on to new cases, a term that is to function at the level supervened upon, but is to group together exactly the items to which competent users would apply the supervening term. Understanding why just those things belong together may essentially require understanding the supervening term.<sup>21</sup>

Suppose Hare were correct and we could codify moral judgements in the manner of:

---

<sup>20</sup> Thanks to Bert Dreyfus for alerting us to this point.

<sup>21</sup> McDowell (*Mind, Value and Reality*, 202).

(Q) Any action with non-moral properties  $Q_1-Q_m$  is just,

there would be a list comprising arbitrarily many principles of this form:

(Q) Any action with non-moral properties  $Q_1-Q_m$  is just,

(Q') Any action with non-moral properties  $Q'_1-Q'_p$  is just,

(Q'') Any action with non-moral properties  $Q''_1-Q''_q$  is just,

and so on.

In order to apply the term 'just', we would need to see what all the groupings  $[Q_1-Q_m]$ ,  $[Q'_1-Q'_p]$ ,  $[Q''_1-Q''_q]$  and so on, have in common with one another, and be able to proceed into the future. Why, for instance, is  $[Q_1-Q_m]$  grouped with  $[Q'_1-Q'_p]$ ? We have seen that the possibility of defeaters gives us reason to doubt that there are (useful) principles such as (Q), but even if there were, what reason have we to suppose that the grouping of acts as just is transparent or makes any sense at the non-moral level? Why would we have one grouping rather than another if that were so? The point of the groupings is left obscure on Hare's view. (Compare the case of tin openers: at the non-functional level, it makes no sense to group together such a physically diverse set of objects under a single rubric.)

As Dancy puts it: although the moral supervenes upon the natural, 'the subvenient base may be naturally shapeless'.<sup>22</sup> The onus is placed upon Hare to explain in what sense justice can have a 'shape' on his view, according to which it is a non-moral construct with a prescription contingently attached.<sup>23</sup> The thick intuitionist does not bear this burden so heavily. She can acknowledge a property of justice visible, because of its invariant valence, at the moral level. The valence is not a merely contingent attachment, it is an essential feature of justice. And this helps explain why justice shows up in the moral perspective.

---

<sup>22</sup> Dancy (*Moral Reasons*, 79). This is not strictly accurate. It is actually the set of instances of a moral property that may be shapeless at the natural level.

<sup>23</sup> Hare (*Freedom and Reason*, 187-9).

At least Hare does embrace the univalence of justice. The complaint is that he generates it in the fashion of an afterthought (it is as if we could group things into tin-openers and others without looking to the function of opening tins). The thin intuitionist denies this univalence, and this might yield its own set of problems when it comes to appreciating the 'shape' of justice. Justice is a moral concept, and we suggest that understanding it, qua supervening term, requires an apprehension of its essential connection to the right. But the thin intuitionist denies that there is such an essential connection. With its variable valence, justice is in this respect, for her, on a par with the non-moral properties. Why, then, is justice shapely from, and only from, the moral perspective?

The fat intuitionist has non-moral properties with invariant valences. And such invariance could lend more 'shape' to moral practice from a non-moral perspective than McDowell's argument might lead us to expect.

In order to examine further the relative strengths of the three varieties of intuitionism, it is helpful to have a model of thick intuitionism before us. We take as our starting point Ross' ethical theory. As Urmson rightly points out<sup>24</sup> Ross' famous list of prima facie duties is intended as a list of primary reasons.<sup>25</sup> Ross' provisional list (he tries to reduce it yet further a little later in the chapter) contains duties of reparation, gratitude, fidelity, justice, self-improvement, beneficence and non-maleficence. These are all terms for thick moral concepts; it is thus tempting to regard Ross as a thick intuitionist. If so, as Dancy readily admits, Ross' generalism is immune to the shapelessness argument. 'There is no suggestion in Ross that those who lack the relevant moral concepts should still be able to see the natural shape of those concepts.'<sup>26</sup> Ross' theory, so understood, appears able to take on board the other two arguments employed in favour of particularism. As we have seen, the rule-following argument only undermines the notion that rules can provide apodictic justificatory grounding. The argument does nothing to show that there cannot be rules or general principles.

Moral particularists often support their view by appeal to holism in the theory of reasons. Reasons holism claims that reasons are context-dependent; what counts as a reason in one context may not count as a reason in another. Particularists hope to establish holism by appeal to particular examples.<sup>27</sup> But what these examples establish is only that there are many considerations whose valence can and does change with context. This is unsurprising and should, as we have already seen, be uncontentious. It is hard to see, however, how appeal to a few examples can establish that there are no considerations with unvarying valence. If one holds to the distinction

---

<sup>24</sup> Urmson ('A Defence of Intuitionism', 113).

<sup>25</sup> Ross (*The Right and the Good*, 20-22).

<sup>26</sup> Dancy (*Moral Reasons*, 95).

<sup>27</sup> E.g. Dancy (*Moral Reasons*, 60-62); D. McNaughton *Moral Vision: An Introduction to Ethics* (Oxford: Blackwell, 1998), 192-4

between primary and secondary reasons then any convincing example of a switch in valence will merely be taken to show that the consideration in question is not a primary reason. The strategy of supplying examples is thus powerless in itself to show that there are no primary reasons. It can only show that the considerations in the examples are not primary reasons. To present the choice as being between holding that *every* consideration that can function as a reason is univalent and holding that *none* are would be to leave out of account an obvious and attractive range of middle positions.

One of us has argued elsewhere<sup>28</sup> that Ross is best read as offering a theory of primary and secondary moral reasons, in which secondary reasons have force only in so far as it is their presence in the particular context that ensures the presence of the primary reason. Take the case of lying. That an act is a lie is, for Ross, only a secondary reason against it. Lying is normally *prima facie* wrong. But, in Ross' view, this is due to the fact that in standard cases lying will be both harmful and in breach of an implicit undertaking between members of a society to tell the truth. But we can imagine cases where this isn't so, such as the one Dancy offers us where I play Contraband with my children. Since the whole point and fun of the game is to lie, no harm is done and the implicit contract is suspended. In such cases, that a statement would be a lie does not tell at all against making it; indeed, given the purpose of the game, it may well be a reason for making it. There is nothing in Ross' account that rules out the possibility that any property of an act may be secondarily morally relevant on some occasion. But many of these features do not, in themselves, seem to have any moral significance. The beauty of Ross' account is that it explains how some features, such as lying, are ones that do normally have moral significance -- because in all standard contexts they matter morally -- while allowing that there can be special circumstances in which their normal relevance is cancelled.

We develop below an account of thick intuitionism according to which non-moral features can enter weak moral principles, but only with evaluative riders attached. The evaluative riders lend moral shape to more commonly important non-moral features.<sup>29</sup> Thin intuitionism, by contrast, seems unable to offer any satisfying explanation for

---

<sup>28</sup> D. McNaughton, 'An Unconnected Heap of Duties?' *Philosophical Quarterly*, **46**, (1996), 433-447

<sup>29</sup> The strategy we employ is, of course, not the same as the one attributed to Ross in the previous paragraph, but develops Ross' approach in a significantly different way.

the fact that there are some non-moral properties that are more 'central' than others, and yet can still switch valence on occasion.<sup>30</sup> It explains neither why they do commonly count, nor why they don't count when they don't.

How could a Ross-style thick intuitionism be shown to be unsustainable? It is not sufficient to show that Ross' particular list is in error, for he may have made a mistake of mere detail. What would have to be shown is that no sensible amendments to Ross could yield a plausible thick intuitionism that fitted well with our moral intuitions. Since a thick intuitionism occupies a midpoint between the thin and the fat, Ross' theory is vulnerable on both sides. On the one, there might be arguments to show that there are plausible fat principles. That is, it might turn out that there is some feature, on Ross' list or some successor, that is a plausible candidate for a primary reason and that can be wholly spelled out in non-evaluative terms. On the other, it might turn out that no plausible examples of primary reasons can be found, even when couched in evaluative terms. It might be that there just are no features that only count one way.

Clearly, fully assessing both Ross' list and possible amendments to it is beyond the scope of this paper. But by discussing a few examples we intend to bring out certain structural features of any plausible thick intuitionism. Some of the features on Ross' list appear to fit the thick model very well. It can plausibly be held that the fact that an act is just, for instance, always counts in its favour. And justice is a thick moral concept whose content, it seems reasonable to assume, cannot be spelled out in descriptive terms because it is naturally shapeless. Certainly Ross' account of justice makes unabashed use of the ineliminably evaluative notion of merit, for he takes justice to require distribution according to desert.

Non-maleficence is less straightforward. First, can one be injured as a result of *fair* competition? If no, then the concept of injury involves that of fairness, and is thick. If yes, then an evaluative rider is required: it is *prima facie* wrong to injure others unless in the context of fair competition. In any case, there are other grounds for thinking injury a morally thick notion. It is linked to (among other things) the idea of thwarting another's needs, where needs include morally thick aspects of a life.<sup>31</sup>

---

<sup>30</sup> Dancy (*Moral Reasons*, 67).

<sup>31</sup> Thanks to Berys Gaut for this suggestion.



Other members of the list are trickier still. Take the duty to keep promises. It is hard to see how promise-keeping could fail to be on any deontologist's list of prima facie duties. Yet it might well be claimed that, while promise-keeping is a notion with moral significance, it is not itself an evaluative notion, but one that can be spelled out in non-moral vocabulary.<sup>32</sup> So it looks as if what we have here is a fat rule (a univalent non-moral feature). However, there are occasions on which the fact that I have promised carries no moral weight, gives me no reason to keep it. Suppose that I have promised to do something deeply immoral, such as a contract killing. Do I have *some* moral reason to carry out the killing in virtue of the promise, a reason that is outweighed by the duty of non-maleficence? It seems more plausible to think that here the promise counts for nothing, morally speaking. Or take the case of a promise extracted under severe duress, such as torture or threats. The duress does not just weaken my obligation; it cancels it.

We suggest that a moral principle may draw attention to a feature that is always relevant, and relevant in the same way, under certain implicit conditions.<sup>33</sup> In our example, provided that one's promise is not given under duress and is not an undertaking to do something immoral, then it supplies one with a moral reason to act so as to keep it. We don't normally spell out these conditions when articulating the principle, both to keep it simple and because the principle is, as it were, written with the standard case in view. But the list of conditions is not open-ended, and it is knowable in advance. It has a rationale. The conditions are inherent in a proper understanding of the nature and role of promising in the fabric of our moral life. Since promising is an institution for placing oneself under a moral obligation to perform an act, there would clearly be something self-defeating in allowing that one could use the institution to place oneself under a moral obligation to do an immoral act -- an act that one had a moral obligation not to do. Similarly, if one could be coerced into putting oneself under an obligation, then those who acted immorally in acting coercively could put their victims in the position of being under an obligation to them, which is clearly morally objectionable.

---

<sup>32</sup> J. Searle, 'How to derive "ought" from "is"', *Philosophical Review*, **73** (1964), 43-58. Repr. in P. Foot (ed.) *Theories of Ethics* (Oxford: Oxford University Press, 1967), 101-114.

<sup>33</sup> We are grateful to Tom Baldwin for this suggestion.

Note that we are back with thick intuitionism, because the conditions cannot be spelled out in purely non-moral terms. The notion of coercion here, for example, is not the mere use of force, for promises that are forced from one can still have moral weight. Once Milosevic had agreed to withdraw from Kosovo it would have been a breach of faith if he had attacked NATO troops, even though his promise was extracted by the use of force and the threat of further force. And we condemn those Argentinian soldiers who, during the Falklands war, are alleged to have indicated that they were surrendering and subsequently opened fire, even though force was used to make them surrender. It is, rather, *unjust* coercion (the use of force where the person using it has no *right* to do so) that invalidates the normal force of a promise, as does the use of trickery -- another counter-instance that can be treated along similar lines.

In the case of promising, the primary moral reason is that one made a promise meeting the relevant conditions: it is not a promise to do something immoral and it is not extracted under unjust duress. That an action would be the keeping of such a promise always counts in its favour. We have a univalent feature. But it is not purely non-moral -- it relies upon the notions of immorality and injustice. Hence it does not fit the fat mould. And the list of conditions is not open-ended in the way thin intuitionism suggests. Even if our account of the conditions is not complete and someone were to come up with an ingenious counter-example requiring supplementation or amendment to the principle, we can be confident that the needed amendments would be variations on the kind of theme we have already seen.

It is plausible to suppose that the concept of fidelity, which Ross uses as a blanket term to cover obligations of this type, is governed by the evaluative conditions we have spelled out. Someone who ignores a promise extracted from her under unjust duress is not committing a breach of faith. The notion of fidelity is thus, as one would have expected, a thick moral concept that cannot be fully explicated in purely non-moral terms. Indeed, there is a further respect in which fine evaluative judgement is required in order to decide whether there are reasons of fidelity for acting in a certain way. Ross rightly stresses that not all agreements or commitments are made by explicit promise. In many cases one comes to have a commitment in virtue of undertaking a series of actions; it is often a matter of moral judgement whether the way one has so far acted has placed one under a binding commitment that it would be a breach of faith not to honour.

Other prima facie duties on Ross' list may need similar qualification if we are to discover a feature that always counts one way. For example, it may be that one need owe no duty of gratitude to someone who has helped you, but who has acted immorally in procuring you the benefit that he is now conferring. We suspect that whether you are beholden to your benefactor in such a case will depend on the precise circumstances -- in particular, on his intentions. But it looks as if we need a similar exclusion clause to one of those we had in the case of promising, and for roughly the same reason.

It is worth mentioning one complication. We have suggested that there are *weak moral principles* with implicit conditions. These principles spell out complex evaluative features that always (morally) count one way. But our original exposition (following Urmson) was in terms of primary *reasons*. The original idea was that there were features of actions that were univalent reasons for acting. But the univalent feature that constitutes the reason here turns out not to be 'having made a promise' but 'having made a promise that was neither a promise to do something immoral nor extracted under unjust duress'. Dancy has urged that this is not the right way to think of reasons.<sup>34</sup> I have reason to return your book today, because I promised to do so. It is true that, had the promise been extracted under duress, or had my returning the book been an immoral thing to do, then I would not have had such a reason. But that does not mean that the *absence* of such defeaters is itself part of the reason why I should return it. Similar remarks apply if we construe prima facie duties in terms of right-making properties. What makes it (prima facie) right to return the book is that I promised. It is not part of what makes it right that my promise was not extracted under duress, though if it had been, the returning of it would not have been (prima facie) required of me.

There are two ways of responding to this objection. The first is to tough it out. To say that I promised is to give the most important part, but only a part, of the full reason, which should include the absence of the exclusionary circumstances. Though it sounds odd to say that the fact that the promise was not coerced is part of the reason why it is right to return the book, that oddity is to be explained in terms of the normal conventions of conversation and thought. It is not metaphysically odd or redundant. The other response is to admit the force of Dancy's objection, and to reformulate thick intuitionism accordingly. We might say, for example, that the fact that I have promised does not always provide a reason for acting, but always does so whenever certain specifiable conditions are met. This

---

<sup>34</sup> Dancy (*Moral Reasons*, 77).

would be to give up on the idea that there are univalent reasons here, but this only commits us to the letter of thin intuitionism in this instance, not its spirit, which holds that there is no way of specifying in advance how the valence of promising will alter from one new context to another.

#### IV IN DEFENCE OF THE THICK

So far we have tried to show that there are some reasons for thinking that there might be plausible weak moral principles involving thick concepts, and none involving non-moral concepts. Without a good deal more work we could not show that a fully articulated and plausible thick intuitionist moral theory can be constructed. What we hope to do in this final section is to give further reasons for preferring a thick to a thin theory.

A thin intuitionist might claim that the flexibility of her view accords better with our moral intuitions. But it is not as though thick intuitionism is inflexible: whilst the sign of the valence of the thick concepts cannot reverse or fall to zero on this view, the magnitude of the valence can vary widely. In the case of promise-breaking (where the promise meets the relevant conditions) the valence is always negative, but the magnitude depends on such factors as the content of the promise, to whom it was made, and the circumstances under which it was broken. Suppose someone goes too far in the keeping of a promise, and does the wrong thing thereby -- she should have broken the promise in order to save a life. Does the valence of promise-breaking reverse in such a case? No, it is simply that its magnitude is outweighed by another consideration. Hare points out that the term 'industrious' is 'normally used to commend; but we can say, without any hint of irony ... "too industrious".'<sup>35</sup> Is this a case of the valence of 'industrious' reversing, or is it a case in which industriousness has been carried too far in relation to other considerations? We suspect the latter. And we are even more convinced in the central cases, such as that of someone being 'too just'. Justice always counts in favour, but it must be weighed against other considerations. It might be possible for an act to be the worse for being just. But this is not a case of valence reversal; it simply means that the act was wrong, and the right act was unjust or less just. (It might also be a remark about the motivations of the agent, but that discussion would take us too far afield.)

---

<sup>35</sup> Hare (*The Language of Morals*, 121).

Thin intuitionism does not appear to do justice to the role of thick moral concepts in our moral thinking. Recall Dancy's challenge to the fat or thick intuitionist: what reason have we to suppose that the valence of any consideration (non-moral or thick) is invariant across all possible cases? Why cannot the valence of a consideration in a particular case be reversed (or equal zero) in a new case? Our view is that, when it comes to certain thick concepts, the onus is on the thin intuitionist to justify her claim that the valence *can* vary -- it is a weakness of the position that its radical holism marks no distinction between the thick and the non-moral. There are a number of thick concepts, roughly those corresponding to the traditional virtues and vices, that seem to play a crucial role in moral justification. These moral concepts mediate between the non-moral and the thinly moral. Part of moral wisdom is the ability to locate an act in terms of all the thick properties under which it falls. Moral judgement is then aided by the fact that we know the valences of the thick properties -- we know on which side of the scale to weigh them. We have already mentioned justice, fidelity and gratitude, but there are many other examples. That an act is cruel, mean, or dishonest counts against it; that it is kind, generous or honest counts in its favour. On thin intuitionism, thick moral properties have no more intrinsic moral significance than non-moral properties. It will, presumably, turn out that these properties are 'commonly more important' than some others (although thin intuitionism owes us an account of why), but that not only understates their force, it seems to mislocate their centrality. It is not just that it is helpful to look at them first because they often count; their counting is central to their being thick *moral* concepts.

The valences of thick properties are discovered ontogenetically as the concepts are acquired: learning about rightness and acquiring an understanding of the thick moral concepts occur in tandem, and are mutually reinforcing. Learning about justice, for example, deepens the understanding of rightness. And learning about rightness deepens the understanding of justice: we see how justice and its lack feed into rightness and wrongness. To fail to see that justice has invariant positive valence is to fail to understand fully the relevant concepts, where invariance is a modal notion because, as we have seen, it applies counterfactually -- of any act, if it were just, that would count in its favour.

Davidson, initially in response to Goodman, utilizes the notion of 'fit' between predicates.<sup>36</sup> For example, he sees a lack of fit between the mental and the physical. In somewhat similar fashion, we might re-cast McDowell's shapelessness claim as the claim that there is a lack of fit between the moral and the non-moral predicates we use. And the thick intuitionist claims that there *is* a degree of fit between thick and thin moral properties -- the fit here being responsible for the modal fact of invariant valence. The thick intuitionist gives us the best account of moral modality.<sup>37</sup>

Dancy wants to make room for an element of modality in the moral domain.<sup>38</sup> He attempts to account for this as follows.<sup>39</sup> What we learn in a particular case, when we notice a feature making a moral difference, is that it is possible for the feature in question to make a difference. And in certain modal logics (e.g., S5) the following holds: if possibly P then necessarily possibly P. So we have necessity in the picture. However, we are not convinced that this is a satisfactory account of moral modality -- it is too easily paralleled in low-level non-moral cases. My car does not start, and I learn that this is due to water in the distributor. So I learn that water in the distributor can make a difference to whether my car starts. Hence I learn that necessarily it is possible for water in the distributor to make a difference to whether my car starts.

The fat intuitionist denies shapelessness and claims that there is a degree of fit between some non-moral features and the thin moral properties. And she also owes us an explanation.

---

<sup>36</sup> D. Davidson, 'Emeroses by other names', *Journal of Philosophy*, 63, (1996), 778-780. Repr. in *Essays on Actions and Events* (Oxford: Clarendon Press, 1980), 225-8. Davidson, 'Mental Events', in L. Foster and J. W. Swanson (eds.), *Experience and Theory* (University of Massachusetts Press and Duckworth, 1970). Repr. in *Essays on Actions and Events*, 207-224.

<sup>37</sup> A lot more needs to be said here of course, concerning such tricky matters as scientific 'laws' and their support of counterfactuals but space does not permit.

<sup>38</sup> Dancy (*Moral Reasons*, pp.69-70).

<sup>39</sup> Dancy (*Moral Reasons*, p.70, and n. 11, 72).

What other work do weak principles do for the thick intuitionist? Do they, for instance, help us see what is relevant in a new case? On Ross' view, I come to apprehend the prima facie moral principles by a process of intuitive induction; that is, I notice that one or more particular acts are prima facie right in being (say) just, and then come to realize that there is a self-evident necessary general truth here. Since, however, I could only come to recognize the general principle because I am capable of recognizing straight off that an action is prima facie right in virtue of being just, I do not use this general moral principle to *infer* that another just act is prima facie right -- I can see this straight off, on normal occasions, just as I could in the first cases.<sup>40</sup> Dancy concludes from this that 'moral principles play no epistemological role'.<sup>41</sup>

However, first, if one admits, as Ross does, that there are exceptional circumstances in which one 'apprehends individual facts by deduction from general principles'<sup>42</sup> then it is an exaggeration to say that principles play no epistemological role. And, second, Dancy's argument seems to assume that the only epistemological role a principle could play is by acting as a major premise in a deductive inference. But it can inform what we are able to see 'directly' (this is part of the point of the discussion of shape). Cases vary in complexity and obscurity. It may be that it is only when we have grasped the general principle that we can see in less straightforward cases both that an action is just (among other things, we know that this is something to look for) and that it is prima facie right in virtue of being so. Weak principles help us apprehend what to do; they help us capture relevant features and help us see how they are relevant.

Weak principles also help us to justify our actions, even in tricky cases. It may be unclear whether the fact that I have promised does make a difference. Where there is genuine uncertainty, it is helpful to be able to appeal to the principle of fidelity and to reflect on the kinds of case that constitute the exceptions. Only conditions of a certain kind are allowed to disqualify the making of a promise from having its normal force. While we still have to judge

---

<sup>40</sup> W. D. Ross, *Foundations of Ethics* (Oxford: Clarendon Press, 1939) 168-71.

<sup>41</sup> Dancy (*Moral Reasons*, 95).

<sup>42</sup> Ross (*Foundations of Ethics*, 171).

whether this case falls under the general rubric, having a general rubric helps us to justify our action in addition to alerting us to look out for the relevant conditions.

Whilst we favour thick intuitionism over its fat and thin rivals, we hope to have done justice also to their common underlying conception of morality. Moral judgements cannot be 'read off' strong principles. Rather, factors must be assessed *in situ* to determine their role, if any, as reasons. And these must then be weighed.