

Recalcitrant pluralism

Article

Accepted Version

Stratton-Lake, P. (2011) Recalcitrant pluralism. *Ratio*, 24 (4). pp. 364-383. ISSN 1467-9329 doi: 10.1111/j.1467-9329.2011.00507.x Available at <https://centaur.reading.ac.uk/24384/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1111/j.1467-9329.2011.00507.x>

Publisher: Wiley

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

RECALCITRANT PLURALISM

Philip Stratton-Lake

Abstract

In this paper I argue that the best form of deontology is one understood in terms of *prima facie* duties. I outline how these duties are to be understood and show how they offer a plausible and elegant connection between the reason why we ought to do certain acts, the normative reasons we have to do these acts, the reason why moral agents will do them, and the reasons certain people have to resent someone does not do them. I then argue that this form of deontology makes it harder to unify a pluralistic ethics under a single consequentialist principle in a plausible way, and illustrate this with reference to Rob Shaver's consequentialist arguments.

Introduction: Moral foundationalism

The various forms of deontology and consequentialism are foundationalist theories. They are foundationalist in the sense that they both offer differing views about what the most fundamental moral principles are, and maintain that all of our obligations can be derived from these principles. Principles are *moral* in the relevant sense here in so far as they specify what our *prima facie*, or actual duties or obligations are.

Although one might make quite valid distinctions between thin deontic notions such as right, duty, obligation, ought, required, etc, in this paper I follow Ross in using these terms interchangeably. So unless stated otherwise, when I talk of a right act I mean an act that we ought to do, rather than a merely permissible act, and when I talk of what we ought to do, I mean what we morally ought to do, or what is obligatory.

By a *fundamental* moral principle I mean one that is not derived from some other moral principle, and from which other moral principles are derived. So, for instance, if, as Ross thought, we ought to be honest because by being honest we will be keeping an implicit promise, and we ought to keep our promises, then the principle of fidelity to promises will be fundamental, or at least more fundamental, than the principle of honesty. This is because the principle of honesty is derived from that of fidelity. It may be that the principle of fidelity to promises is in turn derived from the principle of good promotion – that we ought to keep our promises when and because we will be producing the best outcome by doing so, and we ought to produce the best outcome. If that is true then fidelity is derivative and good promotion is fundamental.

Derivation here is an explanatory notion. One principle is derived from another if the duty that figures in the derivative principle is explained with reference to the duty that figures in the principle from which it is derived. Note that the notion of derivation here is understood with reference to duties, or obligations, and explanations of these. It is not an epistemological notion of derivation. The fundamental claim here is not that we can *know* that we ought to do a certain sort of act because by doing it we will be doing something else that we ought to do. Rather the claim is that we ought to do the one act because we will be doing something else that we ought to do.

It may be that I can know that I ought to be honest because I know that by being honest I will keep an implicit promise to express my genuine opinions, and I know that I ought to keep my promises. But the order of knowledge need not track the order of explanation of the duties known. For instance a non-derivative moral principle may be knowable with reference to some other principle. That does not mean that the principle known in this way is not really fundamental, or that it is both derivative and non-

derivative at the same time. So the fact that some basic moral principle can be known on other grounds and on the basis of our knowledge of other principles, and that our knowledge of the basic principle can be justified with reference to those other grounds and principles in no way casts doubt on its fundamental nature.

Of course, we would expect that what is epistemically fundamental and derivative will track what is explanatorily fundamental and derivative, but it need not do so; and when the epistemic order does track the explanatory order in this respect this will typically be because the relations of grounding and grounded between moral principles are more fundamental than those relating our beliefs in those principles.

What is more fundamental still in this respect is the content of the moral principles rather than the principles themselves. This is because moral principles, at least as I understand them, state that certain acts are right or obligatory, and if one principle is derived from another, this will be because the obligations picked out by one principle are derived from the obligations picked out by the other principle. Relations amongst moral principles simply mirror the relations amongst the obligations they record.

I said at the start that moral principles are those that state *prima facie* or actual duties. My favoured form of deontology takes its inspiration from W. D. Ross's *The Right and the Good*, so I think the best form of deontology is one spelt out in terms of principles of *prima facie* duty. I favour this form of deontology not just because of doubts about absolutist deontological theories,¹ but more importantly, because principles of *prima facie* duty are better suited to the task of discovering what is morally fundamental than are principles of duty proper.

I would go so far as to say that basic moral principles *must* be principles of *prima facie* duty. Principles stated in terms of duty proper can *never* be basic. This is because they all work at the level of what is recommended (peremptorily) – that is, they all state what we ought to do, or are required to do. But in trying to get at what is morally basic we cannot rest content at the level of what is recommended, but have to move to the more fundamental level of what recommends. This level is more fundamental because what recommends explains the relevant recommendation.

If I say to you that you ought to do some act, say, give me £20, it always makes sense for you to ask *why* you ought to do that act.² In answering this question I move from the level of what is recommended – your giving me £20 – to the level of what recommends, the reason why you should give me £20. This level is more basic because it explains the recommendation.

It is tempting to say that another ought must figure in the explanation of an ought, and so duties proper, recommendations, must figure in fundamental moral principles. But this cannot be right. For if a duty proper figures in the explanation of some other duty, then we can always ask of the more basic duty *why* we ought to do what it recommends, and this will move us back to the fact that recommends rather than to the fact that some

¹ By an absolutist moral theory I mean a theory that holds that there are one or more absolute moral prohibitions – that is, at least one type of act that is always wrong no matter what. Kant is the canonical absolutist, but of course act consequentialists are also absolutists, as they maintain that it is always wrong to fail to promote the good.

² One of the things I find unsatisfactory about moral discourse conducted in terms of rights is that there is no pressure to explain why we have some right.

action is recommended. To insist that a duty proper must figure in the basic level will lead, therefore, to an infinite regress of explanation.

Since Hume, it is often held that one cannot derive an ought from an is, and so an ought must figure in the explanation of another ought. But we are here talking about explanation, not logical derivation, so even if Hume's principle is true, it is not relevant here. You cannot logically derive an effect from its cause, but that doesn't mean that all causal explanations are deficient.³ Similarly the fact (assuming it is a fact) that one cannot logically derive an 'ought' from an 'is' doesn't mean that deontic explanations that don't mention oughts are deficient.

As I understand *prima facie* duties they are recommenders not recommendations. They do not tell you what you should do, but pick out features that give you reason to do certain acts. This point is quite lost in the terminology Ross uses in *The Right and the Good*. But despite the misleading terminology, it is clear that by '*prima facie* duty' he does not mean a special sort of duty, or recommendation.⁴ He means, rather, a moral feature of acts related to duty proper. The contexts in which Ross introduces this term, as well as much of what he says about the notion of a *prima facie* duty, make it clear that the way in which he thinks that *prima facie* duties are related to duty proper is by explaining them. Principles of *prima facie* duty pick out features of acts or situations that tend to make certain acts right. So if they are not defeated in some way they will pick out the reason why some act is right.

This means that principles of *prima facie* duty pick out explanatory reasons – the *reason why* some act is our duty, or the *reason why* we should do that act, or *why* it is right.⁵ At this point it may be asked what sort of explanation these facts provide. They clearly do not offer causal explanations. The fact that you have promised to do some act does not cause that act to be obligatory. Neither is this a conceptual, or teleological explanation. But although I can list the sort of explanation that recommenders do *not* stand to the obligations they explain, I have to confess that I cannot offer a positive characterisation of what sort of explanation recommenders provide of deontic facts (though I take some comfort from the fact that no one else has any idea either). What I *can* do, however, is say something about how recommenders relate to normative reasons and motivation. Once we are clear about these connections the task of subsuming something like a Rossian pluralism under some single moral principle in a plausible manner becomes, I suggest, much harder. So my proposal here is that the move from moral principles understood as principles of duty proper, to principles of *prima facie* duty

3 One might insist that a complete explanation must entail the explanandum, and so some relevant principle must figure in the explanation. This is not the place to dispute this claim. All I'll say is that this is a substantive claim in the theory of explanation, and does not flow from the very idea of explanation. One might think that causal explanations are explanations with reference to causes. On this view a principle could not figure in a causal explanation, as principles do not cause anything, and it is difficult to see how they could be part of a cause of anything.

4 "The phrase '*prima facie*' duty must be apologized for, since... it suggests that what we are speaking of is a certain kind of duty, whereas it is in fact not a duty, but something related in a special way to duty." (W. D. Ross, *The Right and the Good*. Stratton-Lake, P., (ed) (Oxford: Clarendon Press, 2002) p.20)

5 Once again, I am not saying that if these features obtain then we ought to do the relevant act. Rather, if we ought to do the relevant act, then one of these features will be the ultimate explanation why.

– from recommendations to recommenders – makes this form of moral pluralism more resistant to subsumption under unifying principles.

I have said the principles of *prima facie* duty pick out facts that explain why certain acts ought to be done. By facts here I mean instances of properties in certain things or actions, so will sometimes talk of properties, or features, that explain rather than facts. But when I talk of properties in this context I should be taken to mean that some act (or situation or thing) *has* that property – that is, some fact. What gets explained is a distinctive sort of fact – a normative, or more precisely, a deontic fact – the fact that some agent ought (or is required, or is obligated) to do some act. We may call such explanations deontic explanations.

In *Kant, Duty and Moral Worth*⁶ I talked of the reasons that figured in deontic explanations as ‘the normative reason why I ought to do some act’ (see, eg, 18-20). But the phrase ‘normative reason why’ is an incoherent mixture of two sorts of reasons – a normative reason and an explanatory reason. For some fact to provide a normative reason, is for that fact to count in favour of some attitude, or action (or some disjunction or conjunction of acts or attitudes).⁷ I follow Scanlon in thinking of reasons as basic and believe that all we can say about the way in which these facts count in favour of attitudes and actions is that they do so by providing a reason.⁸

I have said that deontic explanations are facts that explain deontic facts. The explainers are the *reasons why* these deontic facts obtain. So the explainers are a certain sort of reason. But they are explanatory, not normative reasons – they are the *reasons why* we ought to do some act, and ‘the reason why’ is an explanatory reason.

Although I do not know how to argue for this view, I believe that every such explanatory reason is also a normative reason.⁹ The reason that explains why we ought to do some act is always a normative reason to do that act. Or to put this another way:

If F explains why you ought to Φ , then F gives you a reason to Φ

So, for instance, if you ought to Φ because you promised to, then the fact that you promised to Φ gives you a reason to Φ . Or if you ought to Φ because Φ ing will prevent a great deal of suffering, then the fact that Φ ing will prevent a great deal of suffering gives you a reason to Φ . And so on. This link between explanatory and normative reasons is, I think, why I conflated these two notions in *Kant, Duty, and Moral Worth*. But although the explanatory and the normative reason are the same fact, the normative reason does not count in favour of the same thing that the explanatory reason explains. The explanatory reason explains the fact that I ought to do some act. It is a relation between the fact F and this deontic fact. The normative reason is provided by the same (explanatory) fact F, but does not count in favour of the deontic fact. It could not, for

6 Philip Stratton-Lake, *Kant, Duty, and Moral Worth* (London: Routledge, 2000).

7 Scanlon insists that they count in favour only of attitudes (See T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Massachusetts: Harvard University Press, 1998) p.21), but that causes unnecessary complications.

8 *What We Owe to Each Other*, p.17.

9 John Broome, also endorses this principle in his ‘Reasons’ in *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith (Eds). (Oxford: Clarendon Press, 2006) pp. 28-55.

deontic facts are not the sort of thing of which normative reasons can count in favour. What F counts in favour of is not the fact that I ought to do a certain act, but my doing this act.¹⁰

So by replacing principles of duty proper with principles of *prima facie* duty we move away from principles telling us what we ought to do to principles that state which facts are the *reasons why* we ought to do certain acts, and which provide us with reasons to do those acts. Because of their essential connection with obligations we may call these reasons *moral* reasons, and it is less confusing if we call the principles that specify these reasons *principles of moral reasons* rather than principles of *prima facie* duty.

Deontic Reasons

There is a tendency amongst Kantians, as well as other philosophers, to regard the obligation itself as the distinctive moral reason, but this cannot be the only moral reason we have, as Kant himself seemed to think, and once we recognise that facts that provide deontic explanations are moral reasons, the temptation to think that such deontic facts are any sort of normative reason should diminish. We always have a moral reason to do what we ought to do, not because the fact that we ought to do that act is a reason, but because the fact that explains this deontic fact is a reason. Once we are aware of the reasons provided by the facts that explain our obligations, it will seem very odd to suppose that the obligations themselves are also reason providing.

Suppose I tell you that you ought to be grateful to Brad, and you ask ‘why?’. I reply by pointing to the various things Brad has done for you in the past, often at some cost to himself, and claim that you ought to be grateful to him for that reason. Given the principled link between deontic explanations and moral reasons, by telling you the *reason why* you ought to be grateful I have, at the same time, picked out facts that give you a *normative* reason to be grateful. Let’s assume you accept this. So you know that the various favours that Brad has done for you in the past not only explain why you ought to be grateful, but also give you a moral reason to be grateful. It is hard to imagine that you might ever be tempted to say that I have not listed all of the moral reasons you have to be grateful – that I have left out a, if not the most important moral reason, to be grateful – namely, the fact that I ought to be grateful. The only reason I can think why someone might insist that this deontic fact is reason providing is that they cannot think of where else the distinctively *moral* reason to act could be located. Once we see clearly that this may plausibly be located in the facts that figure in the deontic explanation, the temptation

10 At this point it is worth noting an ambiguity in the notion of a normative reason. It can mean either the fact that provides the normative reason, or the fact that that fact provides a normative reason. We may, following Parfit, call the first a *normatively significant* fact and the second a *normative* fact. Understood as *normatively significant* facts, normative reasons are the facts that stand in the reason providing relation to some attitude or action. Understood as a *normative* fact, a normative reason is the relation the first fact stands to the relevant attitude or action.

This ambiguity will be important in getting clear about whether what explains deontic facts is the same reason we have to do what we ought to do, as the answer is that in one way it is, and in another it is not. The explanatory reason *is* the normative reason, where the latter is understood as the *normatively significant* fact: But it is not the normative reason if it is understood as the *normative* fact. The fact that F is a reason for me to do some act is not the same as the fact that F is the reason why I ought to do that act. But the fact that explains what I ought to do and which is a reason for me to do what I ought to do is the same.

to think that distinctively *moral* reasons must be provided by the deontic fact should at least diminish. I think they disappear.

Moral reasons and moral motivation.

So far I have talked about a certain type of explanatory reason – the one that figures in deontic explanations – and the normative reason the explanatory reason provides to comply with the obligation it explains. But there is a third sort of reason that is closely linked to the first two, which is a motivational reason. So it is to this, and the relation it stands to the explanatory and normative reasons discussed so far, that I now turn.

In one of her papers Christine Korsgaard makes the following claim about Kant's argument for the Categorical Imperative in *Groundwork* I:

Kant is analysing the good will, characterised as one that does what is right because it is right, in order to discover the principle of unconditionally good action. The assumption behind such an analysis is that *the reason why a good-willed person does an action, and the reason why the action is right, are the same*.¹¹

Although I think that this assumption has to be hedged by various conditions,¹² I think it is true – that morally good people will tend to be motivated to do what they ought to do by the reasons why they ought to do those acts. Since this thesis asserts a symmetry between the facts that explain why we ought to act in certain ways and why good people tend act in those ways, I call this the symmetry thesis.

Gripped by the Kantian view that the thought of duty is the only thing that will motivate good people, Korsgaard also claims that '[t]he good-willed individual does the right thing because it is the right thing', but given the symmetry thesis this cannot be correct. For if they do the right thing for the reason why it is right, and they do the right thing because it was right, it would follow that the right act is right because it is right, and that, of course, is no explanation at all.

Given that the facts that explain why we ought to do certain acts give us a moral reason to do those acts, it is no surprise that the symmetry thesis is true, at least once it is freed from Kantian prejudices about moral motivation. For all it states is that virtuous individuals will be motivated to do what they morally ought to do by the moral reasons they have for doing those acts, which is just what we would expect from the practically wise.

Indirect consequentialists will object to the symmetry thesis on the ground that it rules out a view they endorse. Indirect consequentialists might hold that we are unlikely to perform the acts that will produce the best outcome if we are motivated only by consequentialist considerations. For various reasons the best outcome is much more likely to be produced if we do not always aim to bring about the best outcome. It may be produced by a motivational set that treated non-consequentialist considerations as moral reasons – not merely as derivative reasons, but as reasons on their own account. So such

11. 'Kant's Analysis of Obligation', in *Creating the Kingdom of Ends* (Cambridge: Cambridge University Press, 1996, pp. 43—76) p. 60.

12 See my *Kant, Duty, and Moral Worth*, pp. 16-20

consequentialists may have motivations which are indistinguishable from those of Rossian deontologists. If this form of consequentialism were correct, then the features that would motivate good people to do what they ought to do would *not* be the same as the reasons why they ought to do those acts, and the symmetry thesis would be false.

But that indirect consequentialists would reject the symmetry thesis does not seem to me to cause a serious problem for the thesis. This is in no small part because the symmetry thesis has a great deal of initial plausibility to it and a rationale that is independent of any particular normative theory. If you are motivated by some consideration it seems you must regard that consideration as a reason to act, and in so far as you are rational we would expect you to be motivated by what you regard as reasons to act. A normative theory that requires us to be motivated to do what we ought to do by anything other than the reasons why we ought to do those acts requires that we sever this link between our normative judgements and the explanation of our actions. That counts against the normative theory rather than against the symmetry thesis.

Furthermore, it is not even clear that such a severance is psychologically possible. Could we persuade ourselves that non-consequentialist reasons have no independent normative force, and gain any normative force they have only by association with consequentialist considerations, and, at the same time, get ourselves to be motivated *non-derivatively* by these non-consequentialist considerations? It would seem that the only way in which we could get ourselves to be non-derivatively motivated by non-consequentialist considerations, is to regard those considerations as providing non-derivative reasons. But we could not get ourselves to believe that while we believe consequentialism.

For both these reasons I think that the fact that some normative theory requires us to reject the symmetry thesis should cast doubt on that theory rather than the symmetry thesis with which it is incompatible.¹³

Furthermore, in so far as some theory requires us to reject the symmetry thesis, it will complicate the connections between the three moral reasons so far mentioned – namely, the reasons why we morally ought to do certain acts, the moral reasons we have to do them, and the reasons why good people will do them. The Rossian view simply identifies these reasons in an economical, elegant and plausible manner. By rejecting this economical and elegant symmetry a monistic theory like consequentialism undermines at least some of the supposed advantage it has over pluralistic theories by having a single grounding principle. A theory with a single grounding principle is, all other things being equal, better than one with more basic principles, because it is simpler.¹⁴ But if this gain in simplicity generates further complications elsewhere, it is not at all clear that there is an overall gain in theoretical simplicity, such that the monist theory is, in this respect, a better theory than the pluralist one.

Being wronged and reasons to resent.

¹³ There are many other criticisms of consequentialist attempts to separate motivation and justification. See esp “Indirect Consequentialism, Friendship, and the Problem of Alienation” by Dean Cocking and Justin Oakley. (*Ethics*, Vol. 106, No. 1 (Oct., 1995), pp. 86-111)

¹⁴ Brad Hooker makes much of this in his argument for rule-consequentialism in his *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Oxford: Clarendon Press, 2000), pp. 19-23.

I now turn to a fourth and final reason that the facts picked out by basic principles of prima facie duty provide. These facts are the ones that pick out who we have wronged when we wrong someone, and the feature in virtue of which we have wronged them. For someone to be wronged by my action is for that person to have reason to resent my action. The feature in virtue of which this person is wronged is the fact that provides this reason.

Arguably at least, not all wrong acts involve wronging someone. If, for instance, I drive home drunk without having an accident, then it seems I have acted wrongly without wronging anyone. For although everyone would have reason to disapprove of my reckless behaviour, I do not think that anyone would have reason to resent it. Nevertheless, most wrong acts involve wronging certain individuals or groups. If I fail to keep my promise to you, then I not only act wrongly, but wrong *you*. If I fail to be grateful to someone who did me a favour in the past, then I wrong that person, and if I fail to save someone whom I could easily help, without any significant cost to myself, then I wrong that person by ignoring his plight.¹⁵

One of the virtues of a Rossian theory of prima facie duty, or as I prefer to call them, principles of moral reasons, is that the facts picked out by these principles not only give us moral reason to act in certain ways, explain why we ought to act in those ways, and explain why we act in those ways (in so far as we were virtuous), but also picks out the individuals or groups who are wronged if we fail to act in those ways. They also pick out the features in virtue of which those individuals are wronged.

Ross maintains that the reason why I ought to keep my promise to A to do a certain act is simply that I have promised A that I would do that act; and the reason why I ought to be grateful to B is that B benefited me in the past. These reasons are captured in the principles of fidelity and gratitude. The principle of fidelity states that the fact that I have promised A to Φ gives me a moral reason to Φ , and the principle of gratitude states that the fact that B did me a favour in the past gives me a reason to be grateful to B. Assuming that these reasons are not defeated, it will turn out that my duty proper is to keep my promise, or be grateful, and the facts mentioned in these principles will explain these duties. Since it is the fact that I promised A that explains my obligation to do what I have promised, it is A who is wronged if I fail to keep my promise. Since it is the fact

¹⁵ It is often said that we only wrong the other person if they have right against us, and that this is not the case in relation to beneficence. Accident victims, or the starving or destitute, have no right to our help, so, it is claimed, we do not wrong them when we fail to help. We may question this claim. We do not need to attribute a right to be helped to ground the claim that in failing to help them we wrong them. All we need is the idea that we owe help to others in such cases, and we do not need to assume that the needy have a right to be helped to ground that idea. It may be sufficient for an act of help to be owed to another person or group for us to be unable to justify not helping on grounds the other person could not reasonably reject.¹⁵ And we could not justify our failure to help someone in dire need when helping would involve little effort on our part? Suppose, to use a variant of Scanlon's example, you are trapped under some electrical equipment, and you are receiving agonizing electrical shocks as a result. I am walking by and notice that I could end your suffering by flicking the fuse switch nearby. It is impossible to imagine that I could justify my failing to press the switch to you on grounds you could not reasonably reject. That would make my helping you in this way something I owe to you. This in turn would be sufficient warrant the claim that by failing to press the fuse switch I not only act wrongly, but wrong you.

that *B* benefited me in the past that explains my obligation to be grateful to *B*, it is *B* who is wronged if I miss an opportunity to express my gratitude to *B*, or reciprocate. It is *A* and *B* in each case who have reason to resent my not doing what I promised to do and my failure to be grateful respectively.

Being wronged is not the same as being harmed here. My promise to *A* may be to help *B*, and my failure to help *B* may be something that does not harm *A* in any way. It may, therefore, be the case that I harm *B* by breaking my promise. In such a case I wrong *A*, not *B*, even though it is *B* who is harmed and not *A*, for it is to *A* that I promised to help *B*. Similarly, when I fail to benefit *B* when I ought to, I will wrong him, but I need not harm him. Failing to benefit someone is not harming them. I could benefit all of my students now by giving each of them a five pound note, but I am not harming them by not doing this.

Principles of moral reasons also pick out the reason an individual has to resent my failure to do what I ought. These features are the same as the reason why I ought to do the acts I failed to do. In relation to *A* it is the fact that I promised him that I would do a certain act, and in relation to *B* it is the fact that he did me a favour in the past. In the first case, the reason *A* has to resent my failure to do a certain act is provided by the fact that I promised *A* that I would do that act. In the case of *B*'s resenting my not benefiting him, the reason to resent my failure is provided by the fact that he benefited me in the past. In each case the reason why I ought to do a certain act, the reason I have to do that act, and the reason another person has to resent my not doing that act are provided by the same facts – the facts picked out by the principles of moral reasons.

To summarise then: the best form of deontology is defined in terms of a plurality of principles of moral reasons. Each of these principles pick out four reasons – two normative and two explanatory. If *F* is the feature that is *prima facie* right, and succeeds in making Φ ing my duty proper, then:

1. *F* is the reason why I ought to Φ
2. *F* is a moral reason for me to Φ
3. *F* will be a reason why I Φ , in so far as I am virtuous.
4. *F* will be a reason for those wronged by my not Φ ing to resent my not Φ ing.

So although there is a plurality of basic moral principles this Rossian deontology, these principles do give us an elegant and very simple account of these four moral reasons. Given that any moral theory, in so far as it aims to provide the fundamental moral principles or principle, should provide principles of this form – that is, principles that pick out recommenders rather than simply provide recommendations – any attempt to reduce a Rossian pluralism to a single principle of moral reasons must pick out a single feature that provides a plausible account of all four reasons. At the very least a failure to do so will be a strong *pro tanto* consideration against the reductive theory, for in this respect it will have lost much of the simplicity the Rossian theory has, despite the gain in simplicity gained by the reduction itself. Rossian pluralism is resistant to such reductions because no single principle can pick out a single feature that provides a plausible account of all four reasons for every action that we ought to do.

Moral reasons and recalcitrant pluralism

It is much more plausible to subsume a plurality of moral principles under a single principle if these principles are principles of duty proper. This is because the recommendations contained in what will become the derivative principles will remain in place, and will simply be explained by some other recommendation, such as the recommendation to produce as much intrinsic value as possible. So the principle telling me to keep my promises is left untouched if grounded by a principle telling me to do this because by doing it I will be complying with a principle telling me to do something else, like maximise the good. But it is much harder to present a plausible monistic reduction if we are starting with principles of moral reasons. This is because the facts picked out by these principles must be able to be reasons in the four different ways I have outlined above.

What I want to do now is briefly illustrate how this form of deontological pluralism makes it harder to subsume the plurality of basic moral principles under a single principle in a plausible way by showing how it deals with a particular subsumptive strategy employed by consequentialists. The particular strategy I have in mind is that of 'expanding the good'. I'll focus on Robert Shaver's attempt to employ this strategy as he focuses on Prichard's and Ross's anti-consequentialist arguments.

Expanding the good

So what is the 'expanding the good' strategy? It's basically a set of simple instructions on how to absorb any apparent counter-example to consequentialism. The instructions are as follows:

- 1) For any supposed counter example to consequentialism, such as a case where some special personal relation seems to make a moral difference on its own account, take the act that constitutes honouring or respecting that relationship.
- 2) Treat the fact that this act occurred as a consequence of the act.
- 3) Maintain that that consequence has intrinsic value.
- 4) Maintain that it is the intrinsic value of this consequence that accounts for, and thus neutralises the apparent counter-example to consequentialism.

By following these instructions, it is argued, consequentialism can get the right answer to these difficult cases.

There are two standard objections to this strategy. The first is to maintain that this strategy achieves only a semantic victory for consequentialism, so is really just an empty move. The other is to claim that this strategy gives rise to the wrong sort of moral motivation. I have argued that normative moral theories are best understood as theories of moral reasons on the ground that such principles are best suited to the job of finding foundational moral principles. So understood the expanding the good strategy is not merely an empty move, but makes a substantive claim, and in my view is mistaken. The moral motivation point is correct as far as it goes, but picks up on only one of the four moral reasons picked out by such principles, and can be bolstered by the other ways in which the content of basic moral principles pick out moral reasons.

Family relations

In “What is the Basis of Moral Obligation?” Prichard gets us to consider the following scenario:

Suppose all the houses and provisions of a village are damaged by a landslide, and suppose every villager but one is hurt. Suppose the circumstances such that the one sound man could only devote himself to providing food and shelter for one person. The acquisition of food and shelter by each would be equally good. Hence on the [ideal utilitarian] view there would be an obligation and an equal obligation on the one sound man to find food and shelter for himself, for his parent, and for a stranger. But while it might be contended that he had the right to consider himself as much as any other...no one would suppose that he was under a positive obligation to do so. Moreover everyone would hold, contrary to the view, that the obligation to help his father was greater than that to help a stranger. (B 2).

In “The Birth of Deontology”¹⁶ Shaver argues that an ideal utilitarian may try to deal with this type of objection by expanding the good. Prichard’s objection seems to treat human well-being as the only good. It is on this ground that he thinks that the utilitarian would claim that the villager ought to help either himself, his father, or the stranger, as the benefit to each of them is the same. But, Shaver points out, ideal utilitarians need not be restricted in this way. They can allow that many other things are good. They could allow that a state of affairs in which a son helps his father is intrinsically good, and the intrinsic value of a state in which a son helps his father would enable the ideal utilitarian to claim that in the situation Prichard describes the villager ought to help his father rather than the stranger. This is because the good consequence ‘that a son helps his father’ breaks the tie between the other goods provided by the equal benefits he can confer either on his father, the stranger, or himself.

So the ideal utilitarian is not committed to saying that in such situations the villager ought to help either his father, the stranger, or himself. Rather utilitarianism can generate the more plausible answer – namely, that he ought to help his father, and would have acted wrongly if he benefited the stranger.

Shaver puts his point as one about what one ought to do. His concern was to show that the ideal utilitarian would recommend the right action in cases where personal relations seem relevant. But as a contribution to the debate about fundamental normative theory he would have to claim not only that the villager should help his father in preference to a stranger, but that the reason why he should do this is because by doing so he will bring about an intrinsically better state of affairs than any other action open to him. But if this is the reason why the villager ought to save his father, it would also have to be the reason for the villager to save his father, the reason why he would save his father (in so far as he is virtuous), and the reason his father would have to resent his son’s action if he helped the stranger, or did nothing.

The son’s motive

I think the fact that the act will produce the best outcome fails to provide a satisfactory account of any of these reasons. In defending this view I shall start with the villager’s

¹⁶ In *Underivative Duty: British Moral Philosophers from Sidgwick to Ewing*. Thomas Hurka (ed). (Oxford: Oxford University Press, forthcoming).

motivating reason. Let us suppose that the villager is a decent individual who aims not only to do the right thing, but for the right reason also. Let us also assume that his judgement about such matters is good and so he typically succeeds in doing the right thing for the right reason. So he rightly judges that he should help his father. Why would he help his father in preference to a stranger? Given the symmetry thesis this will be because of the reason why he ought to help his father. So if the consequentialist account of why he ought to help his father is correct, then the reason why he would help him in preference to a stranger would be because he will be producing a better state of affairs in the world than if he did any other act. The fact that he stands in any special relation to his father would not count at all; for according to consequentialism, the reason why he should help his father in preference to a stranger is not because he stands in this relation – that is the deontologist's answer – but because by doing this he will be bringing about an intrinsically better state of affairs.

The deciding factor in whether this state of affairs is best is the fact that his act will have as a consequence a state in which a father helps his son. But this should not lead us to think that the relation that figures in this state will figure in the villager's motivation after all. The fact that his act will have this consequence is the reason why the act will have the best outcome, since it breaks the tie between the value of the benefit conferred on either his father, or a stranger. But it is the fact that his act has produced the best outcome, not the fact that it has produced a state in which a son helps a father, that makes it the right act to do. Otherwise this really would look like a mere semantic victory for the consequentialist. What makes it stand out from the intuitionist account is that it is the production of impartial value that grounds this duty. This is not merely a rewording of the intuitionists view.

But it is very hard to accept that the son would be motivated by a concern to produce impartial value in so far as he is virtuous. If asked why he helped his father instead of the stranger he would, I would suggest, answer 'because he is my father'. His answer would be the same if we asked him why he thought he *should* help his father instead of the stranger. The reason why he would help, and his view about the reason why he should help, his father would be the same, and neither would mention the amount of intrinsic value that his act would produce.

Shaver is aware of this objection to the expanding the good strategy. He writes:

[someone] might reply that the problem is not that the ideal utilitarian cites consequences too far removed from those we really cite. It is rather that in deciding upon our duty, we do not think of these states of affairs as goods at all. But while this may be sometimes, or even often, true, the ideal utilitarian can note that if, on reflection, we do think these states of affairs are goods, it does not matter much that we do not in every instance of acting have this thought. ('Prichard's Arguments Against Ideal Utilitarianism').¹⁷

But it does matter. The motivating thoughts of good people typically reveal what they regard as the features that explain why they ought to act in certain ways. As far as that

17 In *Impartiality Vol. II*. Brian Feltham and Philip Stratton-Lake (eds) (Oxford: Oxford University Press, forthcoming). Shaver notes that Ross recognises this point in his *Foundations of Ethics* (Oxford: Clarendon Press, 1939) p. 69.

goes it is just irrelevant whether on reflection he comes to recognize that helping his father would bring about the best state of affairs. The villager could accept that in helping his father he would be bringing about a state of affairs in which a son helps his father (how could he deny that!), and may also accept that this state is intrinsically good. This in no way supports consequentialism over deontology. For a start he may think that this state of affairs is good because it is one in which the son has done what he ought to do. Then the duty would explain the goodness, rather than the goodness explain the duty; and the duty would have to be explained on other grounds.

Furthermore, although the villager might accept that his act will have a good consequence, it is difficult to imagine that this would get him to revise his initial account of why he ought to help his father – that is, of what really matters to him here. What matters to the villager, and what will motivate him to act, is that he would be helping his father, not that by doing so he would be making the world a better place. And if this fact about his father motivates him to do what he should, even after reflection, this shows that he regards this fact as a reason to do what he should.

Things are worse when we turn to the reasons for resentment. Suppose that the villager helped the stranger instead of his father, perhaps because he thought that he would produce the same amount of good either way and decided to settle the issue by tossing a coin. His father would then have reason to resent the fact that his son helped a stranger in preference to his father. What would provide this reason? If the consequentialist answer were correct then it would have to be the fact that his son failed to produce the intrinsically best state of affairs. But although that *could* be the reason for his father (and for anyone else) for that matter, to *disapprove* of his son's action, that could not be the reason for his father to resent his son's action. On the contrary, the fact that the villager has failed to bring about the best state of affairs leaves the fact that his father has been wronged completely unexplained. If anyone is wronged here, it would be the world which has not been made as good as it could have been. But that, of course, makes no sense.

Shaver also tries the same strategy in relation to cases where keeping a promise seems to mean that the right act is not the one with the best consequences. Exactly the same objections would apply to this instance of the expanding the good strategy as applied to the case of the villager helping his father, and hopefully that argument is clear enough that it does not need repeating.

Indeed, once it is appreciated that Ross's principles of prima facie duty are principles of moral reasons in the four different ways mentioned, even Ross's own attempt to reduce the principles of justice, self-improvement and beneficence to the principle of promoting the good, look unlikely to succeed.

Other, non-consequentialist monistic theories might do better. For instance an attempt to subsume all of Ross's basic principles under a Kantian principle of respecting others as ends in themselves might be able to provide a plausible account of the four moral reasons mentioned. But whether it does will depend on how the idea of treating others as ends in themselves is understood. If we take Kant seriously when he equates this principle with the principle of universalizability it looks hopeless. And Robert Audi's interpretation of treating others as ends in themselves looks too close to treating others beneficently to free itself from variants of Ross's criticisms of utilitarian reductions,

especially if the principle of beneficence is subsumed under the principle of promoting the good.¹⁸

But even if other monistic theories do better than consequentialism in providing a plausible account of the four moral reasons mentioned here, my claim is that a Rossian understanding of moral principles as principles of moral reasons makes it much harder for monistic theories to succeed in their reductions. So even if some monistic theories can overcome this extra difficulty, I hope at least to persuade you that there is this extra difficulty to overcome.

Department of Philosophy
University of Reading
Whiteknights,
Reading,
Berkshire,
RG6 6AA
UK

18 See Audi's *The Good in the Right, A Theory of Intuition and Intrinsic Value* (Princeton: Princeton University Press, 2004) pp. 91-2