

Ross and the Foundations of Morality

Garrett Cullity

Ross's substantive moral theory—his theory of the content of morality—exhibits a distinctive type of pluralism. It is a pluralism about what makes right actions right: a view of the structure of morality according to which there is more than one fundamental right-maker.

This is attractive in various ways. Its picture of the structure of morality is faithful to an important part of moral experience—the experience of conflicting priorities that must be balanced against each other. It allows that often such conflicts can be resolved, since one of the conflicting considerations is weightier than the other, but that the consideration outweighed can still count in favour of the action that ought not, all things considered, to be done. It eschews the search for verdictive principles of the form “All actions of type T are morally wrong”; but it does so without giving up the ambition of providing an explanation of morality's underlying structure. Ross offers us a model for answering his question, “What makes right actions right?” that is not subsumptive but instead combinatorial. We explain what makes right actions right not by subsuming them as instances of general verdictive principles, but rather by identifying the most basic elements of morality and then seeking to explain its complexity by reference to the ways in which those elements can interact. This picture allows the determination of right action to be nuanced and context-specific; but it does so without trying to write all of the nuance into a set of verdictive principles and thereby making moral justification esoteric. It allows that the justifications for right action can still be as simple as “Because I promised”—and therefore that they are available to be *given* by ordinary reflective agents to each other.

Rossian pluralism, as I shall define it in Section I, is a genus of substantive moral theories that share these attractions. Ross's own theory is one member of that genus. This specific view treats as foundational a set of “general principles of prima facie duty”:¹ on the now-standard interpretation, which I follow, these principles specify pro tanto reasons (that is, reasons which have some weight, but need not be decisive), which serve both as reasons for or against acts of certain kinds, and reasons for the rightness or wrongness of those acts.² He has a specific view of the content of these duties, and a specific picture of how they interact to determine overall rightness.

¹ (Ross 1930), p. 30.

² (Urmson 1974–5), p. 113; (Stratton-Lake 2011), p. 149.

In this paper, I want to examine some of the further possibilities: possibilities for developing members of the same genus that depart from the specific features of Ross's own view. I shall do this by focusing on seven of those features:

Self-Containment. His foundational moral norms are the sole determinants of rightness and wrongness.

The Right and the Good. He has a hybrid view of the priority-relation between the right and the good.

Action-Focus. His foundational moral norms all concern action. (One of them is a norm for gratitude, but he treats that as an action too: "I use 'gratitude' for the returning of services, irrespective of motive.")³

Reason-Foundations. The relation involved in his foundational norms is the normative reason relation.

Subsumption. There is one way in which derivative moral considerations derive from the fundamental ones: by being instances of them.

Balancing. There is one way in which the considerations derived from different foundations interact with each other to determine the overall rightness of an action: by bearing a comparative weight.

Invariance. The foundational reasons he identifies always contribute in the same direction. So the principles describing his foundational norms are exceptionless.

Working through these features one by one in Sections II-VIII, my focus will be on identifying some significant limitations that they impose on Ross's substantive theory. In each case, I shall ask how his theory could be supplemented or modified to avoid the limitations, without departing from the moral general structure of Rossian pluralism.

In some ways, then, this is a critical project. However, I think it is the sort of work that Ross would have welcomed. In presenting his list of prima facie duties as provisional, he implicitly invites us to consider revisions and extensions.⁴ Here, I take up this invitation with respect to the seven structural features just listed. What is the case for structuring a Rossian pluralism that is not a pluralism of prima facie duties, and where does it lead us instead?

I: Rossian Pluralism

³ (Ross 1930), pp. 22-3.

⁴ (Ross 1930), p. 23.

Rossian pluralism, as I shall define it, is the conjunction of two claims. The first is that there is more than one derivationally fundamental right-maker.

The relation of right-making, I am (like Ross) going to assume, is a familiar part of ordinary moral thought and discourse. If you are told that it is right to conceal the truth from Fred, you can ask what makes that right. An answer might be: it will protect him from distress. Pressing further, you can then ask: why does *that* make it right to conceal the truth from him? The most obvious way to answer that second question is to appeal to some further part of the content of morality from which this derives. We might say: because being protected from distress will be better for Fred, overall. Then, if this answer is successful, it also identifies a right-maker. If the protection of Fred from distress makes it right to conceal the truth from him because it will be better for him, then it is also correct to say: its being better for Fred makes it right to conceal the truth from him. The first right-maker derives from the second.

It is an interesting question what further analysis (if any) there is of the right-making relation. I won't offer one here: I just introduce it as the relation we assert whenever we say "X is right because...". When we assert that something is a right-maker, we are making a claim about the content of morality. And in ordinary moral thought and discourse, when we try to explain *why* something is a right-maker, we typically do so by appealing to some other part of the content of morality—for example, another right-maker. This is what I am calling "derivation". One part of morality's content derives from another when it has the first part because it has the second.

Philosophy sometimes attempts something more ambitious. One way of answering the question, "What makes it the case that there is a reason to promote others' welfare?" is via the Kantian (or Nagelian) answer: "There must be a reason to do that if there are to be reasons for anything at all."⁵ This is not deriving one part of the content of morality from another; it is purporting to explain why morality *must* have the content it has.

Since derivation is a one-way relationship, there are then three possibilities. Either derivations never get started, or they go on for ever, or they stop somewhere. Ordinary moral thought tells us to dismiss the first possibility. (If there are any right-makers at all, the first possibility is ruled out, since any relationship of right-making consists in the derivation of one part of morality's content—an action's being right—from another—whatever makes it right.) No one has ever seriously proposed the second.⁶ All actual theories of morality adopt the third alternative. They all reach a point at which, when we ask "From what further part of the content of morality does this

⁵ (Kant 1996), Sect. I; (O'Neill 1992); (Nagel 1970).

⁶ However, for a defence of the "infinetist" analogue in epistemology, see (Aikin 2011).

part derive?" the answer is: "Nothing further". Monistic moral theories are those which recognize only one such stopping-point: one derivationally fundamental principle from which the rest of the content of morality derives. The first respect in which Ross offers us a pluralist theory is that he says there is more than one such stopping-point.

That gives us the first component claim of Rossian pluralism: there is more than one derivationally fundamental right-maker. The second claim is that, given this plurality, there is no further principle governing their contributions to determining the overall rightness or wrongness of an action.⁷ We can appreciate what is distinctive about that further claim by noticing that some prominent theories deny it while still embracing pluralism about morality's derivational foundations. Scanlon's contractualism is an example of this.⁸ This does not trace all morally relevant reasons to a single source. However, given the various morally relevant reasons for or against actions, it offers a single principle for determining their wrongness: the principle that wrong actions are the ones ruled out by rules that no one could reasonably reject. That principle does not propose a derivation of morally relevant reasons from some other part of the content of morality. Rather, it is a principle governing the contribution that the reasons generated by a plurality of foundational norms make to determining overall rightness or wrongness.

Any view that makes those two claims I call a version of Rossian pluralism. But if so, Ross's own view is only one possible species of this genus. Amongst the further claims that distinguish his specific version are the seven I began with. By working through these in turn, examining their limitations, and asking what supplementations we might make to avoid those limitations, we will be helped to appreciate the further range of possibilities that is open to a Rossian pluralist.

II: Self-Containment

Here is a first, simple objection to Ross. Suppose some action of mine would confer a great benefit on you: it would reunite you with your long-lost son, let's say. Would it be wrong for me not to do it? That depends on the impact on me. If it just requires passing on a phone number, then failing to do that would be wrong (all else equal); but not if it means giving up years of my life to help you. This suggests that Ross's Self-Containment claim is incorrect. His foundational *prima facie* duties

⁷ (Ross 1930), pp. 31, 41-42.

⁸ (Scanlon 1998), Chs 4 and 5.

are not the only determinants of rightness and wrongness. Whether my action is wrong can also depend on its impact on my own interests.⁹

It is hard to believe that Ross was unaware of this point, but if a Rossian view is to be carefully formulated, it will need to be explicitly accommodated. There is a straightforward way to do so, which appeals to a certain broad way of thinking about wrongness: an inadequate reasons conception, we can call it. On this conception, an action is wrong when, among the reasons that are available to you, there are serious other-regarding reasons against it, and no adequate countervailing reason in its favour. “Other-regarding” reasons we can gloss as those non-instrumental reasons that consist in facts about the relationships in which I stand to others.¹⁰ And we can add to this a variety of different possible proposals about what it is for a countervailing reason to be “adequate”.¹¹ On this conception, the force of complaining that your action is wrong is that you have no adequate justification for having treated others as you did.

The objection to Self-Containment is easily handled, then. It is true that Ross’s foundational moral norms are not the only determinants of rightness and wrongness. They should be understood instead as the putative determinants of our other-regarding reasons. The wrongness of my action is determined by the bearing upon it of *all* the available reasons.¹² These include its impact on my own interests. When this impact is great, that can make a difference to whether my response to your welfare was inadequate. So although Self-Containment should be rejected, it is easy to see how a Rossian pluralist can do so.

III: The Right and the Good

⁹ Compare (Scheffler 1992), Ch. 7. It is true that the reasons covered by Ross’s norm of promoting the good are not just reasons to act beneficently towards other people. They are reasons to make the world better. And harms to me will bear on that as well as benefits to you. So there is a stronger reason of this kind to perform the first of the actions in the reunion example than the second. But that is not very convincing as an explanation of the difference in wrongness between the two actions. The cost to me has a relevance to whether my action is wrong that is not exhausted by its contribution to how much worse it makes the world as a whole.

¹⁰ On the connection of moral wrongness to other-regarding reasons, see (Mill 1991), Ch. 4; (Frankena 1970), p. 158; (Scanlon 1998), pp. 6–7; (Nagel 1986), Ch. 10; (Scheffler 1992), Ch. 7; (Thomson 1997).

¹¹ The possibilities include making the adequacy of countervailing reasons depend simply on the relative strengths of the reasons for and against the action, as Ross does, (Ross 1930), pp. 28-9; or alternatively on whether the action calls for reactive attitudes such as resentment, indignation and blame (Watson 2004), Ch. 8; on which reasons must be recognized as prevailing if our interaction is to be governed by the exchange of reasons rather than coercion (Scanlon 1998), Ch. 5; or on whether the action meets the demands of second-person accountability we are entitled to address to each other (Darwall 2006), Ch. 5.

¹² This point is sometimes expressed by saying that not all morally relevant reasons are moral reasons—e.g. (Harman 2016), p. 368.

Ross is usually characterized as a non-absolutist deontologist. If by a deontologist we mean a non-consequentialist, and if by a consequentialist we mean someone who holds that the rightness of an action is fully determined by its relationship to an impersonal evaluative ranking of states of affairs, then this characterization of Ross is quite right. Ross's substantive ethical theory is developed in opposition to Moore's consequentialism, but his way of opposing Moore is by annexation rather than outright rejection. Acting to promote good states of affairs is treated by Ross as one source of prima facie duties alongside others. Since there is a prima facie duty to promote the good, there is no type of action that is wrong whatever the consequences, so he is not an absolutist; but since the rightness of an action is not fully determined by the goodness of the associated consequences, he is not a consequentialist.

However, notice that the case of Ross displays one of the problems with the attempt to classify deontologists and consequentialists in terms of whether they make the right prior to the good, or vice versa. Ross has a hybrid view on which in one part of morality (the prima facie duty of beneficence) the good is prior to the right, while in the other parts it is not. This is an interesting feature of his view, but it points to a limitation.

Ross presents a substantive theory of the right, structured as we have seen upon a plurality of foundational right-makers. He also has a substantive theory of the good, which is similarly structured: intrinsic goods are traced to the four basic classes of "virtue, pleasure, the allocation of pleasure to the virtuous, and knowledge".¹³ But when it comes to explaining the *relationship between* the right and the good, he offers us the same picture as Moore. The relevance of the good to the right comes from the prima facie duty we have to promote it. Indeed, that is the sole connection that he recognizes between the good and normative reasons. However, this is another respect in which his view calls for supplementation. To see why, we can turn to the next distinctive feature of Ross's view: its Action-Focus.

IV: Action-Focus

Ross's discussion of goodness is not confined to presenting a substantive theory of intrinsic goods, and to connecting this to the deontic part of morality via the prima facie duty to promote the good. He also has a discussion of "moral goodness", which covers the moral evaluation of three kinds of thing—the motivation of actions (which he distinguishes from the deontically assessable acts that

¹³ (Ross 1930), p. 140.

one is or is not motivated to perform), persons, and feelings.¹⁴ His leading example of a morally good feeling is sympathy with someone else's misfortune.¹⁵

That example draws our attention to an important point. Our morally significant responses to other people's welfare are not restricted to action-responses: they include emotions too.¹⁶ Others' misfortune calls for my action to alleviate it, but it also calls for my sympathy, whether I can alleviate it or not. However, reflection on this point reveals some substantial limitations in Ross's approach. One is that it is not clear what source the attribution of goodness to sympathy is supposed to have in his theory. He identifies the norms that are foundational for the deontic part of morality; but what are the derivational sources of the goodness or badness of our responses?¹⁷ Another is that there seem to be reasons to feel sympathy, and he has not accounted for those: the reasons generated by his foundational norms are reasons to act. And a further limitation is that our morally evaluable responses to others' welfare extend beyond the kinds he mentions: actions of benefiting and harming, and feelings of sympathy. Good responses to the welfare of others also include attention to it and perceptiveness about it; speaking up for it by expressing one's solidarity with those whose welfare is threatened or damaged; and thinking that both others' welfare itself and these various responses to it are important. And to each of these further responses, the first two points apply: Ross's view apparently lacks the resources to account for both their goodness and the reasons we have to make them.

How can we supplement Ross's theory to develop a version of Rossian pluralism that fills these gaps? A first thought will be to expand the principles of the theory so that they cover the morally relevant reasons we have not just for actions, but for other responses of feeling, thought and speech as well. However, in order to do this, we will need to modify Ross's principles in more than one way. While it makes sense to think that there is a *pro tanto* reason for the action of promoting the good, it makes less sense to think that *the good* is something we have a reason to sympathize or express our solidarity with. The proper object of these responses is not *the good*, but one important good-maker: namely, others' welfare.

Accordingly, a first model for expanding Ross's view to cover the full range of responses for which we have morally relevant reasons will be one that replaces his *prima facie* duties with principles of this kind:

¹⁴ (Ross 1930), Ch. 7.

¹⁵ (Ross 1930), p. 155.

¹⁶ We should be more careful than Ross was about identifying emotions with feelings, though. It is better to think of emotions as complex states, of which dispositions to have certain feelings are only one component. See (Solomon 2004).

¹⁷ Perhaps these claims are supposed to be derived from the goodness of virtue: virtue is one of his categories of intrinsic goods. But what determines which actions, persons, and feelings are virtuous?

(W1) Facts about others' welfare give us reasons to promote it, to protect it, to be sensitive to it in thought, to sympathize with others and to express our solidarity with them.

That would be a way of extending the theory's coverage of morally relevant reasons to capture these other responses. But the next challenge is to develop this in a way that relates the normative and evaluative parts of morality. Ross connects them via his fundamental norm for beneficence, which tells us to promote the good. But we saw that that fails to account for reasons to make responses other than action. How can we fill that gap?

V: Reason-Foundations

The simplest proposal for addressing this would be to treat principles such as (W1) as a single set of foundations for morality, treating them not only as the foundations from which its assignments of rightness derive, but as foundational to morality's value-assignments as well.

To motivate that idea, we could draw on a distinctive tradition of thought about value: the tradition recognizing different varieties of goodness. According to this tradition, there are different kinds of goodness and badness, and corresponding differences in the responses they call for. A great work of philosophy calls to be studied and understood, a baby to be nurtured, a charitable institution to be joined, a forest to be protected, a talent to be developed, a good state of affairs to be brought into existence and sustained. These different objects are all good, but good in different ways. What good things have in common is that they all call for a response of orientation in their favour; but there are different kinds of favour-responses, and to them there correspond different kinds of goodness.¹⁸

Given this, the simplest suggestion about the relation between the reasons morality gives us and the values it assigns is the following. Good things are those that call for favour-responses, and the relation of "calling for" just is the normative reason relation.¹⁹ If so, then (W1) can serve as a foundation not just for the reason-giving part of morality, but for the value-assigning part too. (W1) tells us about the particular kind of goodness that welfare has: it is something that calls for these favour-responses: promotion, protection, sensitivity, sympathy, and solidarity.

¹⁸ See (Ewing 1947), Ch. 5; (Von Wright 1963), pp. 8–13; (Rescher 1969), Ch. 2; (Nozick 1981), pp. 429–30; (Gaus 1990), Sect. 13–14; (Thomson 1992); (Anderson 1993), Ch. 1; (Scanlon 1998), Ch. 2, Sect. 3; (Zimmerman 2001), pp. 85–6; (Rabinowicz and Ronnow-Rasmussen 2004), p. 393.

¹⁹ A similar suggestion prefers to say that good things are those for which there are sufficient reasons to make favour-responses: see, e.g., (Skorupski 2010), p. 83; and compare (Chappell 2012), Sect. 2.

The simplicity of this suggestion is attractive. Notice that it is not committed to the reductive claim that the kind of goodness something has is its calling for a particular kind of favour-response. It just claims that good things are those there are reasons to make favour-responses to, and treats truths about reasons as the derivational basis for truths about value. It thereby offers the prospect of a single set of foundations for morality's attributions of reasons and value.

However, this simple suggestion faces objections, from two directions. First, there is the “wrong kind of reasons problem”.²⁰ If you are going to torture my children unless I praise you for being so menacing, then I have a strong reason to praise you. Praise is a favour-response, but here it is not a response to something good. So, apparently, there can be sufficient reasons to make favour-responses to objects that are not good. Moreover, in the other direction, something can be good or bad without there being a normative reason for you to respond to it with either favour or disfavour. If you suffer some impairment that renders you unable to praise Julia, that will deprive you of reasons to praise her. However, that fact about you is irrelevant to Julia's praiseworthiness: that is determined by her qualities, not yours. Reasons are capacity-dependent; value is not.²¹ And, apparently, reasons must also satisfy a condition of personal relevance which value need not. Perhaps Julia is a corporate lawyer in Hong Kong, who handles contractual disputes with great sensitivity and skill. That could be enough to make her praiseworthy—but not yet enough to give me a reason to praise her, in the absence of any relevant connection between her actions and me.

If these objections succeeded, what would they show? Notice that they are not objections to the claim that good things are those that call for favour-responses. They are objections to identifying “calling for” with the normative reason relation. My reason to praise the torturer does not make him praiseworthy; and Julia's praiseworthy qualities can fail to provide me with a reason to praise her. We can instead use “calling for” as a name for the relation expressed by the “-worthy” suffix in such remarks: the relation that praise bears to the praiseworthy. Identifying a trait as praiseworthy tells us about the kind of goodness that it has by asserting a relation between the trait and praise: the “-worthiness” or “calling for” relation.²² That same relation is often expressed by a range of other suffixes which are attached to the names of responses in evaluative terms such as “desirable”, “shameful”, “awesome”, “boring”, “outrageous”, and “repulsive”. Another name used by some writers for this relation is “fitness”. On this usage, to say that an action is praiseworthy is equivalent

²⁰ (Rabinowicz and Ronnow-Rasmussen 2004).

²¹ (Streumer 2018).

²² For Ross on “worthiness”, see (Ross 1939), pp. 275-6.

to saying that it is fit for praise.²³ This way of talking is convenient, so I shall follow it; but it is also potentially misleading, so a caveat needs to be borne in mind. When “fitness” is used in this way, it is not being used (as Ross sometimes does) to refer to all-things-considered appropriateness.²⁴ An action can be praiseworthy without praise being all-things-considered appropriate, and vice versa. “Fitness”, as I am going to use it, just means “-worthiness”.

If we insist on using “fitness” in this way, then the following claim is true by stipulation:

Value and Fitness. Something is good or bad in some way just in case it is fit for some favour- or disfavour-response, respectively.

Fitness, so conceived, just is the relation that a good thing bears to the particular favour-responses that it calls for.

Thus, if you agree with the objections to identifying fitness with the normative reason relation, that is no obstacle to accepting Value and Fitness (as a stipulation about “fitness”). Notice next that, although those objections suggest that we should not identify fitness with the normative reason relation, they also point to the following relationship between them:

Fitness and Reasons. Reason relationships derive from fitness relationships.

If there is a reason to praise you, it will arise in one of two ways. You might be praiseworthy: then the reason for praising you will be some fact about you in virtue of which you are praiseworthy. Or you might not be praiseworthy, but there might be a reason of the “wrong kind” for praising you, as when you are menacing my children. In the latter case, my reason for praising you derives from the reason I have to prevent my children from being tortured, which derives in turn from the prevention-worthiness of torture. So, either way, the reason derives from a fitness relationship. There are two main sources of reasons for action: they can be either facts about an object that make it fit for an action in response to it, or facts about an action’s relation to some other fit response. But perhaps (if the second objection is correct) in order to have a reason to make a response, some further conditions must be met: I must be capable of making it, and it must have the right kind of relevant connection to me.

Given Fitness and Reasons, we can improve the simple suggestion from the previous section with a simple refinement. That suggestion was that we treat a claim about reasons, (W1), as one of morality’s derivational foundations. The improvement is simply to replace (W1) with the corresponding claim about fitness:

²³ Equivalent, and not an explanation of it. On this usage, the claim that a praiseworthy action is a fit object of praise is not a substantive claim. The literature attacking fitting attitude analyses of value can be traced back at least as far as (Broad 1985), p. 265.

²⁴ (Ross 1939), pp. 52-5.

(W2) Others' welfare calls for promotion, protection, sensitivity, sympathy, and solidarity.

(W2) is not itself a claim about reasons, but it is a derivational source of reasons: reasons to make those responses to others' welfare, and reasons to do the other things that conduce to our making those responses (up to and including welfare-protecting praise for torturers). So (W2) still offers the attractive prospect of providing a single set of derivational foundations for both the reasons-theoretic and evaluative parts of morality.

VI: Balancing

If we change the structure of a Rossian theory in this way, we can make room for some further possibilities that Ross's own view seems unable to recognize. As Ross says, the prima facie duties he recognizes as fundamental right-makers function in the same way as Newtonian forces.²⁵ When a prima facie duty is instantiated, it can be outweighed and thereby fail to generate a duty proper, but it retains its force as a pro tanto reason. On this picture, whenever a consideration that derives from one of the foundations of morality is instantiated, it carries a reason-giving weight, and the only way in which such considerations can interact with each other to determine the rightness of an action is by contributing that weight to the overall balance of reasons for and against.

However, there are contexts in which it seems wrong to treat considerations of welfare as reasons to be balanced in your deliberation about what to do. If you have a fiduciary responsibility as a trustee, you cannot properly balance the interests of your relatives in benefiting from the fund under your management against the interests of the legal beneficiary. And if you are a surgeon whose patient decides against having an operation, you should not simply treat that as one input into your assessment of the balance of reasons for and against operating, to be weighed against the benefit they will receive if you go ahead with the operation.²⁶ In these cases, something has gone wrong even if you end up deciding not to embezzle the money or perform the operation. The fiduciary obligation in the first case and the patient's autonomous decision in the second make the consideration of welfare you are wrongly thinking about unavailable to be balanced as a reason by you.

We need to be careful with these examples. Perhaps there can be cases in which a consideration *is* a normative reason for an action but nonetheless it shouldn't register as such in your deliberation. So that reading of the cases remains open to someone who wants to preserve the idea that

²⁵ (Ross 1930), pp. 28–9; (Ross 1939), p. 86. For discussion, see (Pietroski 1993).

²⁶ As Daniel Groll says, the patient's will is "structurally decisive" in this case: see (Groll 2012), pp.699-706.

considerations of welfare always provide pro tanto reasons for the actions that promote it. That reading does raise the question: *why* should we exclude these welfare-considerations from deliberation if they really are reasons? But we would need to consider possible answers to that question before ruling this out.

However, notice that the view on which our normative reasons have deeper derivational sources—in the relationships of *fitness* that our responses can bear to their objects—readily makes room for the possibility that reasons themselves are context-dependent. When a principle such as (W2) is treated as a foundation from which the reasons we have to promote others' welfare derive, this allows for the possibility that such derivations can be blocked, in various ways. Two were mentioned earlier: the derivation of a reason for you to make a response from the fitness of that response can be blocked by your incapacity, or by facts about personal irrelevance. But perhaps they can also be blocked, in cases such as those just described, by fiduciary obligations, or by considerations of respect for the autonomous decisions of those whose welfare is at stake.

Notice that this then allows that in these cases, a person's welfare can be *fit* to be promoted, although features of the context deprive an agent of a reason to promote it. In the examples just described, your relatives' interests in financial security and your patient's interests in getting the best medical treatment remain promotion-*worthy*: your obligations as a trustee or your patient's decision have no impact on *that*. Those interests do, after all, provide you with reasons for actions of various other kinds: reasons to give your relatives financial advice, help them out of your own pocket, direct them towards other benefactors, and so on; reasons to draw your patient's attention to facts they may have overlooked.

Thus, grounding our normative reasons in deeper principles of fitness from which they derive has several advantages. It offers the prospect of a theoretically unified treatment of the value- and normative-reasons-generating parts of morality, and it allows that the derivation of reasons from fitness-relationships can be blocked where, apparently, it is.

VII: Subsumption

So far, we have been exploring some of the ways in which Rossian pluralism can depart from Ross by modifying the elements that a Rossian theory treats as derivationally foundational. Now we can consider the other main possibility for supplementing the structure of his substantive moral theory. This concerns the ways in which one part of the content of morality can derive from another.

“Derivation”, I stipulated, is the relationship that one part of morality’s content bears to another when morality has the first part because it has the second. Ross recognizes one such derivation-relationship: namely, subsumption. Derivative prima facie duties are those that can be subsumed as instances of one or more of the fundamental ones. This is the form taken by Ross’s derivations of the pro tanto moral reason not to lie from the foundational prima facie duties of non-maleficence and fidelity, the pro tanto moral reasons to reward and punish from the foundational prima facie duties of promoting the good and fidelity, and the pro tanto moral reason to obey the law from the foundational prima facie duties of gratitude, fidelity, and promoting the good.²⁷ Here, the idea is that it is because there are always prima facie duties not to harm and not to break one’s promises, and because all instances of lying are instances of either harming or promise-breaking, that there is a prima facie duty not to lie. Since all instances of lying are instances of either harming or promise-breaking, the prima facie duty not to lie is contained by, or subsumed under, the prima facie duties not to harm or break one’s promises.

On closer inspection, it turns out that Ross only discusses one kind of subsumptive derivation, and there are others. For example, suppose I have promised to repay the \$100 I borrowed from you in ten weekly instalments. Then, since I have a prima facie duty to keep my promise, I have a derivative prima facie duty to pay you \$10 in the first week. This is another way in which one duty (the duty to pay you the first \$10 instalment) can be contained in, or subsumed under, another (the duty to repay the whole \$100). But here the relationship is not one of instantiation. Paying you \$10 is not an instance of paying you \$100; it is part of paying you \$100. So derivation-by-subsumption can take different forms. Sometimes, as in the cases Ross discusses, it is an instantiation-relationship; but sometimes it is the relationship of mereological composition—the relationship between part and whole.

Having noticed this, we can next observe that not all derivations are subsumptive. Here is an example. If you have young children, squandering your income in a way that prevents you from being able to provide for them properly can be morally bad. Your income calls not to be squandered because your children call to be looked after. That is why your improvidence is morally bad and not just imprudent. But here, the derivation relationship is not subsumptive. If you restrain yourself from handing your wages to the publican or bookmaker, that is neither a specific way or a part of looking after your children: it is an enabling condition for doing so.

The sense in which not squandering your money is an enabling condition for looking after your children is this. Not squandering your money is something you need to do in order to be able to perform an action (namely, spending it appropriately) that will cause your children’s welfare to be

²⁷ See (Ross 1930), pp. 21, 54–5, 56–64, 27–8; and, for secondary discussion, (McNaughton 1996), pp. 436–8.

promoted and protected. In this case, then, the enabling condition is causal. Causal enabling conditions are those that fit the schema:

Condition C enables event E to cause event F.

In the example just given, C is not squandering your money, E is spending it on your children, and F is promoting and protecting their welfare. To satisfy condition C is to follow the derived norm; event F is the event of following the grounding norm. Notice that “enables” cannot mean “is strictly required for”. If I squander all my money at the pub, it remains possible that a stranger will give me a large sum on the way home. It does not even mean “makes it likelier that E will cause F”: perhaps I am a nasty fellow who is very unlikely to spend money on his children, unless it is acquired through an unexpected gift. Intuitively, it means “removes an obstacle to E’s causing F”.

That example illustrates one way in which C can enable E to cause F—by making it easier for E to occur. Another is by making it easier, when E does occur, for it to have F as its effect. For an example of this, consider one explanation of why it is worth adopting a gentle rather than an abrasive manner of speaking to other people: it makes it easier for apologies to lead to reconciliation. If that is right, then there is an enabling derivation from the fitness of reconciling breaches to the fitness of adopting a gentle manner.

The full class of causal enabling derivations includes a wide range of different causal relationships. It includes cases where C will cause someone else to cause F, will cause the removal of an obstacle to F, will prevent an obstacle to F from being imposed, will be a causal precondition for the removal of an obstacle to F, and so on. All the ways in which C renders it easier for some potential cause of F to have that effect are included.

However, beyond all of these possibilities, there are non-causal enabling derivations too. Consider: in order for some gesture of mine to amount to my expression of solidarity, it must express the right thing, and in order to do that I must possess an appropriate understanding of what I am doing. If I observe a minute’s silence because my parents have told me to, without understanding the significance of this gesture, then what I am doing could perhaps qualify as part of their expression of solidarity for the victims of a tragedy, but it could hardly be mine. So this is a way in which cultural understanding can have moral significance. It can be worth attaining because it is a precondition for my gestures to be capable of bearing the “meanings” that our cultural conventions can endow them with, and thereby qualifying as my expression of the solidarity that others’ welfare calls for. That explanation of the moral importance of cultural understanding is an enabling derivation, but a non-causal one. Instead, it instantiates this schema:

Condition C enables E to mean M.

Cultural understanding is a non-causal enabling condition (C) for my observing a minute's silence (E) to express my solidarity with the victims (M). The fitness of seeking cultural understanding derives in this way from the fitness of solidarity as a response to others' welfare.

In drawing attention to these further derivation-relationships, I do not take myself to be making a deep objection to Ross. Nothing prevents his view from accommodating such supplementations. But it invites an interesting further question for the Rossian pluralist: what is the full range of ways in which the content of morality can derive from its foundations? What other derivation-relationships are there?

§VIII: Invariance

A final, more radical, departure from Ross is encouraged by this question: Can the foundations of morality be captured by a set of *exceptionless* principles?

This is the holist's worry about Ross.²⁸ Consider gratitude, understood in his way as “the returning of services, irrespective of motive”.²⁹ If you have done me a service, is that a reason for me to act in a way that returns it? Sometimes, but not always. I have no reason to return services that were imposed on me against my will in an attempt to manipulate me.

More tellingly, the same point applies to the reasons we have to promote others' welfare; and it applies to fitness relationships as well as reasons. Your welfare, after all, is what benefits you, and if you are a bad person, the things you are benefited by attaining may themselves be bad. It may be in a fraudster's interests to succeed in swindling people; in a sadist's, to get the thrill of seeing other people suffer. But welfare of those kinds is not fitly promoted, protected, or sympathized with.

You might want to resist those judgements. You might deny that doing vicious things can be part of a person's welfare.³⁰ Conversely, you might accept that success in a vicious project can contribute to a person's welfare, but insist that it does then call to be promoted, to some degree—it is just that the resulting reasons there are to benefit people in those ways are outweighed. On either of those views, we could rest content with (W2) as our foundation for the part of morality that concerns our responses to others' welfare.

²⁸ Holism in the theory of reasons is the thesis that “a feature that is a reason in one case may be no reason at all, or an opposite reason, in another”. (Dancy 2004), p. 73. See his Chs 5 and 6 for discussion.

²⁹ (Ross 1930), pp. 22-3.

³⁰ See, e.g., (Foot 2001), Ch. 6; (Badhwar 2014); (Bloomfield 2014); and (LeBar 2013), p. 29.

However, if we do not accept either of those views, we must refine (W2) further. One way of doing so would be to find the characteristics of a person's welfare which make a difference to whether it calls to be promoted (etc.) or not, and incorporate them into a more tightly specified principle of the form:

(W3) When others' welfare is of type T, it calls for promotion, protection, sensitivity, sympathy, and solidarity

—where T is replaced by a description of those characteristics. After all, if Ross had been more careful he would not have said what he did about gratitude. Presumably he did not really think that there is a prima facie duty to return services of every possible kind. A more careful statement of his view would specify the kind of services which there is a pro tanto reason to return. The same strategy could be followed in order to arrive at a more accurate statement of the foundational principle for responding to others' welfare.

However, I think that project would fail. To produce a version of (W3) that was both true and comprehensive, we would need to replace “of type T” with “not constituted by the attainment of immoral ends”. That would exclude the problematic cases—but it would leave us with a principle that could not serve as foundational in a substantive moral theory. Such a theory is trying to identify the derivational foundations from which determinations of what is right, overall, are made. But this principle would itself require prior determinations of what is right, overall, for its application, since the judgement that an end is immoral, all things considered, is itself a judgement that it is wrong, overall, to pursue it. So a moral theory with this version of (W3) at its foundations would be circular and therefore unexplanatory. And any other version of (W3) would be false or incomplete.

What is the alternative? In my view, the correct claim setting out the core foundations of the morality of concern has to be formulated like this:

(W) Others' welfare calls for promotion, protection, sensitivity, sympathy, and solidarity, unless the fitness of those responses is undermined.

This allows for the exceptions, but does not yet say what they are. It asserts a relationship of “presumptive fitness” between welfare and each of those five responses—the relationship of calling for a response unless its fitness is undermined. When there is nothing objectionable about your welfare, it does call for those responses; but there are also cases—cases like the fraudster or sadist—where the fitness of those responses is undermined. Your interest in being reunited with your long-lost son is promotion-worthy, as is your interest in listening to the music you enjoy; the sadist's interest in seeing other people suffer is not.

A view that proposes exception-hedged principles such as (W) as foundational to morality remains a version of Rossian pluralism, since it still satisfies the two conditions set out in Section I. However, it now makes a radical departure from Ross. As we saw, Ross's model for explaining the determination of moral quality is the determination of motion by the interaction of Newtonian forces. His substantive moral theory has a compositional structure that parallels Newton's physical theory. The principles of prima facie duty identify the most fundamental moral force-types. Instances of those types are vectors: they have a direction (for or against) and a magnitude. Their interaction is determined by vector-addition. As Ross sees it, a consideration's always counting in the same direction is essential to its being foundational.

Departing from this, a theory with exception-hedged principles like (W) at its foundations grounds the determination of moral quality on elements that do not always count the same way. In my view, this has one great virtue: it seems correct. Welfare, when it is important, is important fundamentally. The fact that visiting you will cheer you up counts in favour of doing so. What makes that so is that being cheered up is good for you, and your welfare calls to be promoted. Welfare does not derive its promotion-worthiness from some further, more fundamental part of morality. So when welfare is important it is important fundamentally—but not every welfare-state is important.

However, even if *that* seems correct, the claim that an exception-hedged principle could really be derivationally foundational—that it could really be what most fundamentally *makes* right actions right—may well seem hard to accept. Indeed, it might seem that the derivational claim is obviously false. An exception-hedged principle like (W) has the form:

Object *O* calls for response *R*, unless it does not.

But claims of that form are never falsifiable, for any object *O* and response *R* whatsoever—so they are trivial.³¹ Nothing is said in advancing a principle like (W) until we are told under what conditions its exception-clause is and is not triggered.

Moreover, suppose we could do that. Suppose, that is, we embedded (W) within a broader theory that told us when the fitness of the responses it mentions is undermined. Then would we not be abandoning the claim that (W) itself is morally foundational? Our theory would contain a determinative justification for the triggering of the exception-condition in (W), and that (whatever it was) would then have to be morally more fundamental than (W). So Ross was right: only an exception-free principle has the right form to be credibly treated as foundational. Furthermore, if our theory did contain the resources to explain when the exception-condition in (W) is triggered,

³¹ (McKeever and Ridge 2006), Ch. 3.

then that would have to be because it put us in a position to replace (W) with a properly specified principle of the earlier form:

(W3) When others' welfare is of type T, it calls for promotion, protection, sensitivity, sympathy, and solidarity.

It would give us a specification of T.³² If our theory contained an explanation of the exceptions in (W), then it would actually have to be formulable with an exceptionless principle, and not (W), at its foundations.

However, I think this idea should not be dismissed so quickly. It is possible, as others have shown, to give explanations that treat exception-hedged principles as explanatorily basic.³³ In an explanation of this form, some applications of these principles—those in which the exception-clauses are not satisfied—can be invoked in the course of explaining why other applications are subject to undermining.

Here is a sketch of how an explanation of this form could be structured in the case of (W). Suppose undermining occurs just when a welfare-state is a state of “loving the bad” – it is a state of positive orientation directed towards an object that calls for the opposite response. Then we can start by applying (W) to those constituents of welfare that are not states directed towards objects that call for the opposite response. The state of suffering from physical pain, for example, is a (negative) constituent of welfare that is never a state of loving the bad. So the exception-clause in (W) is never triggered for this constituent of welfare: it always calls for the response of protection against it, sympathy with those who suffer it, and so on. We can then invoke this value-assignment in applying (W) to welfare-constituents that are states of cruel or malicious enjoyment of others' suffering. These *are* states of loving the bad, since they are states of positive orientation towards an object (others' suffering) that calls for the opposite responses of protection and sympathy; so here the exception-clause in (W) *is* triggered and states of cruel or malicious enjoyment do not call for promotion. In this way, a theory structured on base-level hedged principles like (W) can determine the value of those welfare-constituents that are responses to good or bad objects, at progressively higher levels. This could be done, without circularity, even in a theory with no unhedged base principles.³⁴ If so, the exceptions allowed for by principles like (W) can be explained without presupposing any further principle of form (W3). Indeed, a theory of this form can be more,

³² Compare (Richardson 1990).

³³ See (Horty 2007); (Horty 2012); (Lance and Little 2004).

³⁴ Suppose all the principles in your value theory have the form:

If and because X belongs to type T, then X has value V, unless (a) X bears relation U to Y and (b) Y has value W. Then there are two ways to settle the value of an object O, prior to any other value-assignment: (i) if O does not belong to any type T, and (ii) if O does not bear any relation U to another object.

not less, explanatory than one that has principles of form (W3) at its foundations, since it can tell us not just that some welfare-states do not call for promotion, but *why* that is so.

IX: Conclusion

I have canvassed six main departures from Ross's own version of Rossian pluralism. Departing from the Self-Containment of his view, we can claim that the norms generating other-regarding reasons are not the only determinants of rightness and wrongness. Departing from its Action-Focus, we can include among the foundations of morality norms governing responses other than action. Departing from its Reason-Foundations, we can give a fuller account of the relationship between reasons and value by seeing reasons as deriving from more fundamental norms of fitness. This also allows that those derivations can be blocked in various ways, so that the interactions between the considerations that determine the content of morality are not limited to the Balancing of pro tanto reasons. We can recognize relations other than Subsumption through which the derivative parts of morality's content derive from its fundamental parts. And, finally, it is possible to develop a Rossian pluralism that departs from Ross's Invariance by positing foundational principles for morality that are exception-hedged.

I do not claim that these suggestions exhaust the possibilities for Rossian pluralism. Although they involve progressively greater departures from Ross's own view, they retain the generic attractions of Rossian pluralism that we began by noting: the attractions of approaching the explanation of what makes right actions right not by searching for general verdictive principles of rightness, but instead trying to account for the complexity of morality by identifying its basic elements and the principles through which those elements interact.

Of course, further questions can be raised about the ultimate viability of all of these suggested departures from Ross—the last of them especially. I actually think a promising theory of morality can be developed that has all of these features.³⁵ However, I cannot hope to convince you of that here. What I do think this discussion goes some way to showing is that Ross broke important new ground in the kind of substantive moral theory he pioneered. He opened up a theoretical landscape that remains largely unexplored and will reward further investigation.

³⁵ (Cullity 2018).

Word length

Text: 8,085

Total: 9,502

References:

Aikin, Scott F. 2011. *Epistemology and the Regress Problem* (Routledge: London).

Anderson, Elizabeth. 1993. *Value in Ethics and Economics* (Harvard University Press: Cambridge, MA:).

Badhwar, Neera K. 2014. *Well-Being: Happiness in a Worthwhile Life* (Oxford University Press: New York).

Bloomfield, Paul. 2014. *The Virtues of Happiness: A Theory of the Good Life* (Oxford University Press: New York).

Broad, C. D. 1985. *Ethics* (Nijhoff: Dordrecht).

Chappell, Richard. 2012. 'Fittingness: The Sole Normative Primitive', *The Philosophical Quarterly*, 62: 684–704.

Cullity, Garrett. 2018. *Concern, Respect, and Cooperation* (Oxford University Press: Oxford).

Dancy, Jonathan. 2004. *Ethics Without Principles* (Clarendon Press: Oxford).

Darwall, Stephen. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability* (Harvard University Press: Cambridge, MA).

Ewing, A. C. 1947. *The Definition of Good* (Macmillan: New York).

- Foot, Philippa. 2001. *Natural Goodness* (Clarendon Press: Oxford).
- Frankena, William K. 1970. 'The Concept of Morality.' in G. Wallace and A. D. M. Walker (eds.), *The Definition of Morality* (Methuen: London).
- Gaus, Gerald. 1990. *Value and Justification: The Foundations of Liberal Theory* (Cambridge University Press: Cambridge).
- Groll, Daniei. 2012. 'Paternalism, Respect, and the Will', *Ethics*, 122: 692–720.
- Harman, Elizabeth. 2016. 'Morally Permissible Moral Mistakes', *Ethics*, 126: 366–93.
- Horty, John F. 2007. 'Reasons as Defaults', *Philosophers' Imprint*, 7.
- . 2012. *Reasons as Defaults* (Oxford University Press: New York).
- Kant, Immanuel. 1996. *Groundwork of the Metaphysics of Morals (1785)* (Cambridge University Press: Cambridge).
- Lance, Mark, and Margaret Olivia Little. 2004. 'Defeasibility and the Normative Grasp of Context', *Erkenntnis*, 61: 435–55.
- LeBar, Mark T. 2013. *The Value of Living Well* (Oxford University Press: Oxford).
- McKeever, Sean, and Michael Ridge. 2006. *Principled Ethics: Generalism as a Regulative Ideal* (Clarendon Press: Oxford).
- McNaughton, David. 1996. 'An Unconnected Heap of Duties?', *The Philosophical Quarterly*, 46: 433–47.
- Mill, John Stuart. 1991. *On Liberty (1859)* (Oxford University Press: Oxford).
- Nagel, Thomas. 1970. *The Possibility of Altruism* (Clarendon Press: Oxford).
- . 1986. *The View from Nowhere* (Oxford University Press: New York).
- Nozick, Robert. 1981. *Philosophical Explanations* (Clarendon Press: Oxford).
- O'Neill, Onora. 1992. 'Vindicating Reason.' in Paul Guyer (ed.), *The Cambridge Companion to Kant* (Cambridge University Press: Cambridge).
- Pietroski, Paul M. 1993. 'Prima Facie Obligations, Ceteris Paribus Laws in Moral Theory', *Ethics*, 103: 489-515.
- Rabinowicz, Wlodek, and Toni Ronnow-Rasmussen. 2004. 'Strike of the Demon: On Fitting Pro-Attitudes and Value', *Ethics*, 114: 391–423.
- Rescher, Nicholas. 1969. *Introduction to Value Theory* (Prentice-Hall: Englewood Cliffs).

- Richardson, Henry S. 1990. 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', *Philosophy and Public Affairs*, 19: 270–310.
- Ross, W. D. 1930. *The Right and the Good* (Clarendon Press: Oxford).
- . 1939. *Foundations of Ethics* (Clarendon Press: Oxford).
- Scanlon, T. M. 1998. *What We Owe to Each Other* (Harvard University Press: Cambridge, MA).
- Scheffler, Samuel. 1992. *Human Morality* (Oxford University Press: New York).
- Skorupski, John. 2010. *The Domain of Reasons* (Oxford University Press: Oxford).
- Solomon, Robert C. 2004. *Thinking About Feeling: Contemporary Philosophers on Emotions* (Oxford University Press: Oxford).
- Stratton-Lake, Philip. 2011. 'Eliminativism About Derivative Prima Facie Duties.' in Thomas Hurka (ed.), *Underivative Duty: British Moral Philosophers from Sidgwick to Ewing* (Oxford University Press: Oxford).
- Streumer, Bart. 2018. 'Reasons and Ability.' in Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity* (Oxford University Press, : Oxford).
- Thomson, Judith Jarvis. 1992. 'On Some Ways in Which a Thing Can Be Good.' in Ellen Frankel Paul, Fred D. Miller, Jr. and Jeffrey Paul (eds.), *The Good Life and the Human Good* (Cambridge University Press: Cambridge).
- . 1997. 'The Right and the Good', *The Journal of Philosophy*, 94: 273–98.
- Urmson, J. O. 1974–5. 'A Defence of Intuitionism', *Proceedings of the Aristotelian Society (Supplementary Volume)*, 75: 111–19.
- Von Wright, G. H. 1963. *The Varieties of Goodness* (Routledge & Kegan Paul: London).
- Watson, Gary. 2004. *Agency and Answerability* (Clarendon Press: Oxford).
- Zimmerman, Michael J. 2001. *The Nature of Intrinsic Value* (Rowman & Littlefield: Lanham, MD).