

## REASONS AND RATIONALITY

Jonathan Dancy

One topic that exercises those who think about the interrelations between different normative concepts is the question whether one of these concepts is somehow basic, and if so which one. A concept might be normatively basic in either of two ways. We might be able to define all the other normative concepts in terms of it, but unable to define it in terms of others. In that case it would be definitionally basic. Alternatively, we might be able to understand the normativity of other normative concepts in terms of the normativity of the basic concept, but not vice versa. In this case it might be called normatively basic. The most plausible candidates for either of these roles are the concepts of 'ought' and of 'reason'. Concentrating on the definitionally basic, theorists have been attracted by three different positions with respect to these two concepts. The first is that 'ought' is basic. The second is that 'reason' is basic. The third is that neither is uniquely basic; but these two concepts together form a local holism. What this means is that one cannot have either concept without also having the other, but that the two concepts are not interdefinable.

Part of the motivation to find a single basic concept is just the philosopher's desire to find neatness and simplicity, no matter how unpromising the material, and where one cannot find them, to impose them. But against that, there is the awkward point that if one does find a single basic concept, say that of 'ought', and announces that all the others are clustered around it in one way or another, one seems to have deprived oneself of any ability to explicate the normativity of one's central concept. That can only be assumed as given; The normativity of other concepts can be explicated by appeal to the normativity of this one, but that is as far as one can go.

Be that as it may, those who take 'reason' to be definitionally basic tend to suggest that 'ought' can be understood simply as meaning 'have most reason'. On this account, the term 'ought' is technically redundant; it is a convenient one, but one can say without it everything that one can say with it.

Those who take 'ought' to be definitionally basic include W. D. Ross. His isolation test for reasons (which is really extrapolated from what he says about prima facie duties) has it that a feature is a reason if and only if, in the absence of further reasons, it would generate an 'ought'. And there are other attempts to turn the trick, all of which I have criticised in Ch. 2 of my recent [Ethics Without Principles](#). One of the most interesting is Broome's suggestion that a reason is a feature whose presence plays a certain role in a certain kind of explanation of an 'ought' – the sort of explanation at issue being what Broome calls a 'weighing explanation'. His idea is that the role of

a 'pro tanto' reason can be understood as that of making a certain sort of contribution to a weighing explanation.

It is consonant with this last approach to think that, though weighing explanations constitute one way in which we can explain the presence of an 'ought', there might be other ways. And if there were other ways, what we would get is an 'ought' for which there is no reason – or at least, no reason of the given sort. And that is just what Broome asserts. Some 'oughts' are held in place by, or grounded in, pro tanto reasons, and others are not. This is the suggestion that prompts the present paper. Broome gives a couple of examples, which I have criticised in my [Ethics Without Principles](#) (Ch. 2). But the general style of his approach can be grasped by considering this example: you ought not to believe both that p and that not-p. This sort of rational requirement is one for which no weighing explanation may be available. And if none can be found, it will be an ought that is not grounded in a pro tanto reason.

However, the very idea of an 'ought' that is not grounded in a reason is one that I find it very hard to come to terms with. This is not because such oughts would be inexplicable: Broome would be perfectly happy to allow that there might be an explanation of them, so long as that explanation is not a weighing explanation. What makes it difficult for me to allow ungrounded 'oughts' is the conception of reasons as favourers. I am one of the many who find it appropriate to think of the relation of 'being a reason for' in terms of the relationship of 'favouring'. Some find this confusing, because they immediately think of favouring as something that people do, and cannot quite see how a reason could do the same thing without an unsuitable and unhelpful anthropomorphism. But the expressions to bear in mind are those of 'speaking in favour of', or 'counting in favour of', or of 'being part of the case in favour of'. A consideration calls for a certain response – or, as I would say, favours that response. The feature which is the reason is what does the favouring (and we can call it the ground), and the action, or belief, or feeling, is what is so favoured; and these two things are linked by the asymmetric relation of favouring.

It is because I think in these terms that I find it hard to make sense of groundless, or reasonless 'ought's. For any such thing would be something we should do (or believe or intend or feel) though there is nothing to be said in favour of doing it. And this seems to me to be a most unsatisfactory situation. Sadly, however, it may be forced upon us. This paper considers whether that is so, what remedies there may be, and the extent to which they may be successful.

1

Consider the following three requirements:

1. Do what you have most reason to do.
2. Do what, if things were as you suppose them to be, you would have most reason to do.
3. Do what you believe yourself to have most reason to do.<sup>1</sup>

Some comments about these. First, I have chosen to express them as imperatives rather than as claims about what you ought (or would have conclusive reason) to do. This is because I don't want to prejudge any issues about the sort of force properly ascribed to these requirements (though even in calling them requirements I have to some extent committed myself.) Second, I need to say how I conceive of the difference between the second and the third requirement. The second requirement kicks in when you have some particular beliefs about the situation that confronts you, beliefs which are false, but which are such that what you believe, if it had been the case, would have given you a reason for, or favour, a certain course of action. You believe that the person next to you is in trouble, and it seems that you ought, so believing, to help her. Actually, she is not in trouble, but that doesn't make any difference. You should (at least try to) help her, because you believe her to be in trouble, even though she is not in trouble at all. The requirement that in this situation you help is not grounded in any views of yours about what you have most, or indeed any reason to do; you don't have to have any such beliefs for the second requirement to get a grip on you. It would be a different matter if, in addition to believing that she is in trouble, you also believe that this gives you most, or sufficient reason to help her. This is where the third requirement kicks in. But in fact one could be caught by the third requirement without being caught by the second. Take any case where the way one believes things to be is not one which, if it were the case, would give one most reason to act, but which one does in fact believe to give one most reason to act.

Third, we have to decide on some terminology. The term 'rational' can be used to characterise all three requirements. Some would say that the first requirement, which I would call a requirement of reason, is a requirement of 'substantive rationality'; they are thinking that to be substantively rational is just to do what you in fact have most reason to do. I, however, prefer just to talk of what you have most reason to do, and leave the notion of what is rationally required for other sorts of case (for the third sort of requirement, in fact). Others would say that the second requirement is a

---

<sup>1</sup> There are significant issues about how exactly to interpret these requirements. Most probably the versions I have given are inadequate – always assuming that there are adequate versions. One issue that my formulations do not address is whether one is required to comply, or merely to conform.

requirement of rationality. If you believe that the house is on fire around you, you are rationally required to jump, even though there is in fact no reason to jump. I think of this way of characterising the situation as less than satisfactory, partly because the notion of what there would, in certain circumstances, be most reason to do seems worth keeping apart from what it is, as things stand, rational to do. But my real reason is that I want to reserve the notion of rationality for cases like the third requirement. Those who act in breach of this requirement, one might say, are at odds with themselves, since they are failing to act in accordance with their own idea of what there is the most reason to do. And this idea of being at odds with oneself, hard though it is to get a good theoretical fix on it, is visibly applicable to other cases, e.g. when one believes that  $p$  and that not- $p$  – another breach of a requirement which we all want to call a requirement of rationality. So I reserve the title ‘rational requirement’ or ‘requirement of rationality’ initially for cases where one can be said to be somehow at odds with oneself,<sup>2</sup> and of my three requirements only the third is like that. Note also that the third requirement seems to apply independent of what reasons there may be for believing that there is most reason to  $F$ . Once one does believe this, no matter what one’s reasons for so believing, one is caught by the third requirement. So it is apparently quite different in nature from, and distinct from, any requirement of reason.

Return to the three requirements. The first question I want to ask is whether, if we do what these requirements require of us, we do something that we had reason, or most reason, to do. The answer to this question, with respect to the first requirement, is obviously yes: if we do what we have most reason to do, we have indeed done something that we had most reason to do. With respect to the second requirement, the answer is less obvious, and I will turn to consider it a bit later. It is with respect to the third requirement that the question has a bite. For if the answer were yes in this case, it would seem that by doing what we take ourselves to have most reason to do, we make it the case that we have some reason, and sometimes most reason to do it.<sup>3</sup> And this is boot-strapping; the belief that we have most reason to act will always give us some reason, and is even in danger of making itself true. We have to avoid this result somehow. But the only way to do this, it appears, is to deny that you have some reason to do what you take yourself to have most reason to do. Similarly, we need to deny that one always has some reason to do what one believes one has some reason to do. And this would mean that the third requirement tells us to

---

<sup>2</sup> I will consider later some cases which don’t fit this rubric but which we still might want to think of as requirements of rationality. But these won’t be anything to do with the second requirement.

<sup>3</sup> It will make it the case that we have most reason in any case where there is no contrary reason.

do something which we may have no reason to do. If we were to say that the requirement tells us something that we ought to do, this would be a case of an ought without a reason.

But let us leave the third requirement aside for a moment and concentrate on the relation between the first two. In speaking about reasons in the way I have, I have already expressed my adherence to a sort of objectivism about reasons. I think that our reasons are given us by, or grounded in, features of the situations that we face, and not by our beliefs about that situation – not even if we restrict ourselves to our reasonable or permissible beliefs, those that we are not at fault in forming.<sup>4</sup> There is an opposing view, which I reject, sometimes called subjectivism, which holds that our reasons are all given us by the world, not as it is, but as we take it (or defensibly take it) to be. And then there is a third view which maintains that in a way both views are right: there are objective reasons given us by the world as it is, and subjective ones given us by the world as we take it to be. The second view, subjectivism, would take it that my first two requirements are identical because the first one is just a misleading way of expressing the second. The first view, objectivism, insists that they are different, and that the second requirement is not to do with reasons at all. Acting in a way that you would have most reason to act if your false beliefs were true is not acting in accordance with reasons of any sort at all. The third view, by contrast, understands the first two requirements like this: first, do what you have most objective reason to do, and, second, do what you have most subjective reason to do. The fact that these requirements can be at odds with each other is understood here as an instance of a familiar situation where one has some reason to do one thing and some reason to do another. We should not expect all reasons to be on the same side, and so we should not be disconcerted by the fact that our objective reasons can require, or at least recommend, what our subjective reasons discourage, or even forbid.

My own view, however, is that things are not as cosy as this. We should not, I agree, be disconcerted by the fact – if it were a fact – that our objective reasons speak in favour of one thing and our subjective reasons speak in favour of another. This situation is no different in structure, one might say, from that in which our financial reasons speak in favour of investing for our retirement while other reasons speak loudly against. But it is not like this. Once I have considered my objective reasons and my subjective ones, it would seem, on the present picture, that I then have to decide what, overall, is the thing to do. But this is peculiar, for a start, because unlike everyone else, I cannot distinguish between my objective reasons and my subjective ones, for I cannot distinguish other than in entirely general terms between the world as it is and the

---

<sup>4</sup> I argue for this position in Ch. 2 of my Practical Reality (Oxford: Clarendon Press, 2000). What I go on to say about it here is merely a summary of what I said there.

world as I take it to be. Other people, however, can do this – not between the world and how they take it to be, of course, but between the world as it is and how I take it to be. So the purported distinction between objective and subjective reasons is really operating from a third person point of view; for the agent, the difference cuts no ice. We can say of you that though you have no objective reason to do it, you have sufficient subjective reason. But now what are we to say to you when you ask us what you should do, or when we ask ourselves what you should do? As far as your subjective reasons go, you certainly should do the deed; but the objective reasons are all on the other side. What recommendation are we to make overall? Or how are we to think of what you did once you have done it? There seems to be no answer to these questions, and this is not because the reasons in each side are equally strong, but because they are of the wrong sort, or sorts, to be stacked up against each other in the way in which ordinary reasons can be. It is not as if there are reasons on one side and reasons on the other, and the balance of reasons (and with it the 'ought') comes down on this side or on that. There are 'oughts' on both sides, and this is a quite different kettle of fish, and a very uncomfortable one.

So uncomfortable it is that it persuades me that there cannot be two complete packs of reasons, objective and subjective. Taking it that there certainly are objective reasons, reasons grounded in, or given us by, the nature of the situation we find ourselves in, I conclude that there are not also such things as subjective reasons, given us by how we suppose things to be. This rules out one way of understanding the relation between the first and the second requirement. But it does not mean that there is no other way of understanding that relation. After all, what the third view did was to convert what we would have (objective) reason to do if things were as we suppose into something that we do already have (subjective) reason to do in virtue of so supposing. If we refuse to make that move, there remains the perfectly good contrast between what we have reason to do as things are and what we do not have reason to do but would have reason to do if things were otherwise. The fact that, had things been as he (maybe even non-culpably) supposes, he would have had reason to do this remains relevant to our assessment of his behaviour, but its relevance can no longer be seen as that of constructing a second pack of reasons to which he is actually responding appropriately. The question then is how else we are to understand this relevance, if not in terms of actual but subjective reasons.

In the next section I investigate this question in some detail. This may seem out of place in a paper that is ostensibly about reasons and rationality, since I have announced that the second requirement does not concern reasons (that being the business of the first requirement) nor rationality (that being the business of the third). But my strategy is to try to use the results of my investigation into the second requirement in resolving the issues about the third. So we need to deal with the second one first.

2

How then are we to understand the relevance of the fact that, had things been as he supposes, he would have had reason to F? The answer that I gave to this question in Practical Reality was as follows. The first requirement is grounded in the reasons, in what the situation gives one reason to do. The second requirement is also grounded in reasons, but those reasons are of a special sort. As well as the reasons that favour actions, there are also reasons favouring combinations of belief and action. In terms of 'oughts': it is not true that if you believe her to be in trouble, you ought to help. That would be a belief making a reason out of nothing. But it is true that you ought to help someone whom you believe to be in trouble. This ought is different; it governs a combination of belief and action, which is perhaps best expressed disjunctively: you ought either to help her or not to believe that she is in trouble. The same distinction applies to reasons. It is not true that if you believe her to be in trouble, you have a reason to help her. But it is true that you have a reason either to help her or not to believe that she is in trouble. And often enough that reason will be conclusive, and generate an ought.

My best example of this idea of a reason on this sort of combination of belief and action is that of what I call 'hypocrisy'. You do have reason not to believe that other people ought not to do this while doing it yourself. But it is not true that if you believe that other people ought not to do this, you have a reason not to do it yourself. There might be no reason for you not to do it. One doesn't make reasons out of beliefs in this sort of way. But there is a reason against – some might say a ban on – a certain combination of belief and action.

The question then is whether one can really understand all 'subjective reasons' as objective reasons on combinations. Although I have claimed in print that this is possible,<sup>5</sup> I now have severe doubts. I do think that this manoeuvre will work sometimes, but I doubt that it can be got to cover the entire ground.

My first question is whether we can actually find a reason in each case. Take the reason not to believe she is in trouble without helping her. What is the consideration that disfavors this combination? If there is a reason not to believe she is in trouble without helping her, we should be able to say what it is. But I have a suggestion about that. The reason at issue might simply be the counterfactual fact that if she were in trouble, this would give you a reason to help her. So the fact that her being in trouble would be a reason for you to help her is a reason for you not to believe

---

<sup>5</sup> In Ch. 3 of Practical Reality.

she is in trouble without helping her. And presumably this manoeuvre, which is merely structural, can be repeated in all further cases.

More worrying is the relation between what we might call 'simple' reasons, by which I mean reasons simply to act, and the reasons on combinations. I am going to suggest that there is a sense in which the simple reasons will always win. If so, the battle between them and reasons on combinations seems not to be a fair one. We are used to the idea that there can be reasons on both sides of the question. But the clash between simple reasons and reasons on combinations does not look like that sort of clash. First, what we are to say to someone who has mistaken beliefs about the situation he is facing, but asks us what he should do? We probably say that he should change his beliefs and then do what in fact he already had reason to do. But, first, to say that he should change his beliefs is to say something that might not be true. Not all false beliefs are culpable; not all are such that one ought to abandon them. And if, as I suppose, the truth of a belief is no reason to believe it (we need other reasons than that), the mere falsehood of a belief cannot be a reason to give it up. Second, and more significantly, it seems that in any battle between a simple reason and a reason on a combination, the simple reason will win, in a way that threatens the idea that the complex reason is a reason at all. What I mean by this emerges when we think of a case where we have no reason to F, and most reason to do something else – walk on by, perhaps. But we believe her to be in trouble, and so believing we ought, or have reason, to stop and help. Now what, in this case, do we have most reason to do? A general maxim about reasons is that where one has potentially conflicting reasons, one ought, if possible to find some course of action that would accommodate all of them. And there will, in the sort of case we are thinking about, always be a way of accommodating both the simple and the complex reason, which involves doing what the simple reason is a reason to do (and changing one's mind about the situation). This is what I mean when I say that in such conflicts the simple reason will always win. Now one might reply to this that in acting in this way one has also done what the complex reason asked one to do, namely either to cease believing that she is in trouble or to help her. And this is true. But the point is that we have lost the very thing we were trying to preserve in all this, namely a sense in which someone who continues to believe that she is in trouble and helps her for that reason is somehow to be approved of. For we have said that there was only one appropriate resolution, and that it was not this one. And on this account, acting according to one's false beliefs is not going to be well thought of at all. One had no reason to do it, and in doing it one fails to take the only course that would accommodate the reasons one did have.

Finally, one might ask about the structure of the supposed reason on a combination. I have expressed this disjunctively: one has a reason either not to believe or to act accordingly. And one of the purposes of putting things this way – a purpose shared by John Broome in his discussion of similar matters – is the suggestion that there are always two ways to move in response to



reasons with this sort of structure. To take the sort of example that Broome considers: it is not true that if you have adopted the end, you have reason to take the means (there might be no reason to adopt the end, and if so the adoption of that end would give you no reason to pursue it), but it is true that you have reason either to abandon the end or to take the means. There are two ways out, but if you stick to the end there is only one way left, that of taking the means. Similarly with the reason disfavouring what I have called hypocrisy: you can respect this reason either by ceasing to believe that others should not act in this way or by ceasing to act in that way yourself. So this is the general idea; there are always two ways out. How is it supposed to work in the sort of case we are considering? The claim must be that you can either abandon the belief that she is in trouble, or help her. If you don't abandon the belief, only one way of respecting the relevant reason remains to you, and that is helping her. But this does not mean that you have a reason to help her.

Broome's original discussion of these issues did not contain any suggestion that the relevant reason is a reason on a combination. He said only (and rather opaquely) that the relevant disjunction gives the 'logical factor' of the reason. But he supposed that there was something going on here which is not captured when one thinks of the content of the reason as disjunctive. There is an asymmetry involved, inside the reason, which cannot be captured in disjunctive terms. The intuitive idea is that believing that she is in trouble 'normatively requires' (Broome's term) helping her, but not helping her does not 'normatively require' not believing that she is in trouble. There is, as it were, a sort of normative flow within that which the reason is a reason to do. And my suggestion about the ground of the reason fits this idea rather well, since there is a rather similar normative flow in the subjunctive fact that if she were in trouble this would give you a reason to help her. It is worth saying at this point, however, that Broome himself has abandoned this idea of a normative flow, or thrust, in favour of the merely disjunctive way of expressing the content of these reasons on combinations which I am at present questioning. His reason, as I understand it, is that there is no way of capturing the idea of normative flow in any logical terms. I don't myself think this is a good reason. Logic was not made for this sort of purpose, and its inability to capture something does little to show that that thing does not exist.

So now the question is whether the two disjuncts are on a par. The idea that they are not is the idea that there are not two equally good ways out, two equally good ways of respecting the reason at issue. And there is at least some reason to think that there are not. Think of the matter in terms of 'oughts' for a moment. Suppose we have a case where you ought, believing her to be in trouble, to help her. There does seem to be a sense in which believing her to be in trouble calls for helping her, but not helping her does not call for abandoning the belief that she is in trouble. I am groping for terminology here, not very successfully. Perhaps the real point is that if someone

were to say to himself 'I'm not going to help her, and so I'll give up believing that she needs help, and then I'll be fine', this would seem to be a very strange and inappropriate resolution.<sup>6</sup>

One should not entirely abandon the thought that there are two ways out. There are indeed two ways out. But if we are trying to capture the sort of normativity that is at issue when we say of someone that he should not, given what he believed, have done what he did, we will fail to do so if we think of the matter in terms of reasons on logical complexes. Even if having those beliefs gave him no reason not to act as he did, there is still a sort of normative flow from believing as he did to acting in some other way, and I do not think that we have managed to capture it.

3

According to my third requirement, you are rationally required to do what you take yourself to have most reason to do. But do you have most reason, or indeed any reason, to do this? The problem, as we saw, was that if the answer to this question is yes, your belief that you have most reason is in danger of making itself true, that is, of proving infallible. And it is plainly not infallible. This is the boot-strapping problem. But as I have expressed it so far, the charge of boot-strapping rests on a mistake. We need to start again.

Suppose that if we are rationally required to  $\Phi$  we have a conclusive reason to  $\Phi$ . (We haven't said what that reason is; we just know that there must be one.) Suppose next that there is a rational requirement that one does what one believes one has most reason to do, and that we ourselves believe that we have most reason to  $\Phi$ . Now this rational requirement is a wide-scope one: we are rationally required to [do what we believe that we have most reason to do]. It is not the narrow scope requirement that if we believe that we have most reason to  $\Phi$ , we are rationally required to  $\Phi$ . If it were, our belief that we have most reason to  $\Phi$  would seem, under these conditions, to make itself true, and one cannot give oneself conclusive reason to do something merely by believing that one has most reason to do it. This is the boot-strapping problem. But it only arises if we take the second requirement as a narrow-scope one. If we take it as having wide scope, we don't get the immediate boot-strapping result. But still, I want to say, something pretty

---

<sup>6</sup> I owe this way of thinking to Niko Kolodny, op.cit. The point is not, of course, that not helping her gives him no reason not to believe that she needs help. Though true, this is irrelevant, as is the equally true remark that believing her to be in trouble gives him no reason to help.

like that result does emerge even so. For even with the wide-scope reading, we get the conclusion that there is a conclusive reason for us to [do what we believe we have most reason to do]. So once we believe that we have most reason to  $\Phi$ , we have conclusive reason to [do what we believe what we have most reason to do, namely to  $\Phi$ ]. We are reading this sort of thing disjunctively, at the moment. So it reads:

If we believe we have most reason to  $\Phi$ , we have conclusive reason either not to believe this, or to  $\Phi$ .

But if we have conclusive reason, there must be some reason present, and what might that reason be? The only candidate, really, is that we believe we have most reason to  $\Phi$ . For it is our believing this that has put us in our normative fix, and demands of us that we act accordingly so long as we retain this belief.<sup>7</sup> But this gives us:

Believing that one has most reason to  $\Phi$  is a conclusive reason for either not so believing or  $\Phi$ -ing.

But what is the difference between a complex reason either to ensure that this reason no longer obtains or to  $\Phi$  and a simple reason to  $\Phi$ ? In either case, so long as the reason obtains, one should  $\Phi$ . So the belief that we have most reason to  $\Phi$  makes itself true again, being a

---

<sup>7</sup> Philip Stratton-Lake objected to this. He urged that we already have a reason to do whatever we believe we have most reason to do. Believing that one has most reason to F merely puts one within the grip of that reason; it does not create a new reason. And there is a distinct plausibility to this. One might say, analogously, that one has a general (and antecedent) reason to take efficient means to one's ends, not to believe a contradiction, and so on, and adopting an end, or a belief, does not add to the reasons already in play. However, adopting an end, or a belief, certainly does make a difference to the normative score; it constrains one, in a way additional to any constraints already in place, by limiting one's permissible options. What it does do, it seems, is to convert the antecedent reason to a specific one. So if one asks what reason one has to F, the answer has to be not only that one has reason to do whatever one believes one has most reason to do, but also that one believes one has most reason to F. So on any account it seems that the belief that one has most reason to F is at least part of a reason to F. And the rest of the reason is there already, as it were. So it is still true, even on this account that one can make a difference to one's own normative score merely by believing that there is such a difference.

conclusive reason to  $\Phi$ . So despite the temporary respite offered by reading the relevant requirement as wide-scope, eventually we come back to the same point: believing that one has most reason to  $\Phi$  is conclusive reason to  $\Phi$ . This is the boot-strapping point again. And similar reasoning will unsettle the weaker claim that we have some reason to do whatever we are rationally required to do. Believing that one has reason does not itself give one some reason, any more than believing that one is required creates a requirement.

Support for this general conclusion can be gained from the fact that you should not change your beliefs about what you have most reason to do in the light of your unwillingness to do it. This cannot be something that you have reason to do either. So the disjunctive conception of the content of the third requirement is not really appropriate. Even if there are in fact two ways in which one could come to accommodate this requirement, only one of them is rationally appropriate. (This is the Kolodny point again.)

This is enough to establish the general point. One might still think that other rational requirements are grounded in reasons, even if this one cannot be. What about the rational requirement either to take means to one's ends, or to abandon the end? Is there no reason to do this? If there were a reason, it would probably be an objective reason on a combination, of the sort that we saw earlier. I have to admit that if there is such a reason, I don't know what it is. But for my present purposes, I don't need to argue the case. The issue that interests me arises even if only some rational requirements are ones we have no reason to satisfy. And we have seen that this must be so. We are lumped with the view that there is a central rational requirement for which no reasons can be found, and this is uncomfortable at best. It is especially uncomfortable if one takes it that there should be reasons wherever there are oughts, and even worse if one has to admit that there is nothing to be said against doing something which is rationally banned.

4

I now return to the issue I discussed earlier, namely the relation between my first two requirements. I argued against the best account of the second requirement, which sees it as an objective but complex one, grounded in a reason, on a combination of belief and action; but I put nothing in its place. This is what I now want to try to do.

There are two questions to be answered when someone does what there is in fact every reason not to do, but what there would have been good reason to do, had things been as he supposed. The first is how he manages (to the extent that he does manage) to escape the usual discredit for doing what there is no reason to do. The second is how he manages somehow to get a sort of

credit for doing what there is no reason to do, credit that arises from the fact that he would have had a reason, had the situation been as he took it to be. I am supposing in this that the agent is not at fault for having the views he does. Not every mistake is a culpable mistake, after all. But in such a case, if the agent's perspective is non-culpable, and his behaviour from then on just what it should have been, given that perspective, it seems as if there is no culpability anywhere. So this seems to explain why the agent escapes the usual discredit.

This story is told in terms of culpability, not in terms of reasons or what the agent ought to do. The point is more on the evaluative side than on the deontic one. We have not said that the agent did what he ought to have done; in fact, we persist in saying that he did what he ought not to have done. But we do not award him the usual discredit for this. If we go on to ask why we think positively well of him, presumably the answer will be similar. Once he has got into his mistaken perspective, from that point on he acts as a good, or decent, or sensible person would have done. He acts, that is, in the way that a good person would have done in the situation that he (perhaps non-culpably) takes himself to be in. Even if his beliefs are beliefs that he ought not to have formed – by which I mean more than that they are false – still he acts well from then on. And, again, not all false beliefs are ones that one ought not to have formed.

In saying this we have left to one side the fact that what he did was wrong. He did the wrong thing, but still there is a respect in which he acted well. He did what a good person, even a person sensitive to reasons, would have done in the situation as he took it to be.

There is of course a question whether this evaluative story, which finds much in the agent to praise, can be held in place alongside the deontic one which condemns the action. But there is bound to be a tension in our characterisation of the situation viewed as a whole, and it seems to me that the tension we find when we combine evaluative praise with deontic condemnation is not so stark as to constitute an inconsistency. In fact, I would say that if there were no tension, our account would be wrong. One might worry that to say that what he did was wrong but that in many respects he acted well involves a sort of carving up of the territory in such a way that all the tension vanishes when it ought not to. It is as if the action was wrong but the agent good; so we have evaluation of agent on the good/bad scale, and assessment of act in terms of right and wrong. There cannot be wrong agents, and even though there can be bad acts it is often held that the badness of an act is a different feature from its wrongness. Since the act is one thing and the agent another, it will not be surprising if our assessments of them differ on occasion.

I don't, however, think that any of this is coherent, and even if it is coherent it is not what I was intending to say. In my view there are not two distinct property-bearers, agent and action, to be found here. If the act is wrong, this means that the agent acted wrongly. One can express the fact that the agent acted wrongly by saying that what he did was wrong, that he was wrong to do it, or

that he did a wrong thing; but in my view we should not take this to introduce a new, separate object in addition to the agent in action. The agent acted wrongly, but despite this, in many respects he behaved well. This is as good a tension as anyone could want.

So we approve of someone who acts in accordance with the second requirement, even if he has no reason so to act, and good reason not so to act. If there had been a reason for him to act, that would also have served as our reason for approving of him. Since there is no such reason for him, is there then no reason for us to approve of him? I have a suggestion about this. I said earlier that, if there were an objective requirement on a complex, we would be able to find a reason or ground for that requirement. The reason why he ought, so believing, to act accordingly is, in each case, that if things had been as he believed, that would have given him reason to act. This counterfactual fact about what would have been a reason in different circumstances itself favours his acting when he believes as he does. If this is defensible, it could also serve as an account of our reason for approving of his so acting, even if we give up the view that the second requirement is an objective, reason-grounded, constraint on a combination. Our reason for approving is just that, if things had been as he believed, this would have given him reason to act. As before, what I am suggesting here is a merely structural manoeuvre, and can be repeated in every case. Note, however, that the repeatability of the manoeuvre does not mean that it is the same reason in every case. In the example I have been using, the reason for our approval is that if she had indeed been in trouble, that would have given him a reason to help.

Further, if we are able to find a reason for approval in each case, we can stick to the idea that there is no normativity without a reason, at least as far as the second requirement goes. There is no reason for him to do what the second requirement requires, but there is a reason to approve of him if he does that, and this is enough.

5

I now turn, finally, to try to apply this general picture to the case of rational constraints. Even if there is a lack of reasons suited to generate rational requirements, it might be that rationality is a virtue. Irrationality could be a kind of defect, even if it is not wrong to be irrational. Niko Kolodny puts the point interestingly in the following way:

Rationality might be understood as a kind of executive virtue. Executive virtues, like courage and tenacity, are not dispositions to recognize and respond to a special kind of reason. They are, instead, dispositions that help one to execute one's beliefs about one's reasons, whatever they may be, or to execute one's intentions, whether or not one believes there are reasons for

them. Courage, for example, is a disposition not to be deterred by fear from doing what one believes one ought to do, or what one, perhaps akratically, intends. Suppose ... that rationality consists in having the attitudes that one believes one has reason to have. Then rationality seems a kind of executive virtue. It is a disposition to execute one's beliefs about one's reasons for and against one's attitudes.

I doubt that there is only one virtue of rationality, but that is not the point. Kolodny is suggesting, of the third requirement, that it demands that we behave virtuously. Those who do what they take themselves to have most reason to do are behaving virtuously in this respect, at least, even if they have no reason to do it and the features that they take to be reasons would not even be reasons if they were the case. As far as this goes, the idea seems to be analogous to what we said above about the second requirement.

Kolodny raises, but does not answer, the question why we approve of this virtuous disposition – and of the other virtuous dispositions involved in the idea of rationality. My problem, however, is that though I think it possible to understand the third requirement in terms of an executive virtue, in a way that is analogous to the way we understood the second requirement, I don't think it possible to give an analogous account of the reason we have to approve of those who do what they take themselves to have most reason to do. Such an account would go as follows. The consideration that favours our approving of your doing what you believe you have most reason to do is that if your belief were true, that would give you most reason to do it. But when we unpack this supposed reason, it becomes the tautology that if you were right in believing that you have most reason to do this, you would have most reason to do it. And a tautology can favour nothing. Further, in the second requirement version of this manoeuvre, the left hand side of the favouring relation was itself an instance of a favouring relation. But the truth that if you have most reason, you have most reason is not an instance of the favouring relation. Or, to put it another way, that one has most reason cannot favour one's having most reason. So we get no analogous account of the reason that we have to approve of those who do what they take themselves to have most reason to do.

This means that there is serious unfinished business. If we do approve of those who conform to the third requirement, there should be a reason for us to do so, and I have to confess at the moment that I don't know what it is. Nor, I confess, do I know why one should conform to other supposed rational requirements such as that one take means to one's ends, that one not believe a contradiction, that one believe what one takes to be entailed by things one believes, and so on. Of course there are advantages to be gained by behaving in these ways. But if we appeal to these, we find ourselves caught in the old debate between consequentialist and non-consequentialist understandings of the virtues. Consequentialists will point to the manifest spin-

off benefits of rationality, and say that it is because of them that these particular character traits count as virtues. Non-consequentialists don't need to dispute the spin-off benefits, but can say, first, that some such traits seem to be virtues even though they don't have any obvious benefits at all, such as proper pride and self-respect, and, second, that a virtue may have spin-off benefits without necessarily counting as a virtue only, or even mainly, because of them. Proper pride, for example, is not determined as the amount of pride that pays off in some way or other, and we seem to admire it for itself rather than for its consequences. Imaginativeness, and perhaps (the right amount of) inquisitiveness, are rather similar.

Consider consistency, which is no doubt a virtue in a believer. But why so? I don't think I know the answer to this question, because all the answers that occur to me seem to be wrong. Why one ought not to believe a contradiction? One common suggestion is that an inconsistency entails anything whatever. But even if it is true, it does not seem to me to focus on the right sort of point. Nor is the answer that if one does believe an inconsistency, one thereby guarantees that one believes something false. For there is not always something wrong with believing a falsehood, nor even with guaranteeing that one does so. It is no better appealing to such things as the (supposed) fact that belief aims at the truth. False beliefs fail, on this account, just as much as inconsistent ones; the only difference is that the inconsistent ones are guaranteed to fail, and that doesn't seem to be really the point. It is not that one is bound to be wrong about something (and don't forget the compensating advantage that one is bound to be right about something too), but rather that one's judgement is at odds with itself, and in a particularly glaring way. But this notion of 'at odds with' is no more than that of inconsistency. And at this stage I have run out of answers.

If there are virtues of rationality, there are reasons to approve of those who display those virtues, to try to display them ourselves, and to inculcate them in our children. One last possibility occurs to me, which is that in deciding that we don't know why these things are virtues, we are not deciding that we don't know the reasons for approving of those who display them. We think highly of someone who avoids inconsistency, whose beliefs are coherent, who has a good sense of what is relevant to what and so on – and these things are the reasons for approving of her. That we don't know – if we don't – why they are reasons, perhaps this does nothing to show that they are not.<sup>8</sup>

---

<sup>8</sup> Thanks to Niko Kolodny and Michael Smith, to audiences at Reading and at Aarhus, and to Brad Hooker and Philip Stratton-Lake, for helpful comments and responses to various versions of this paper, whose first outing was at a conference on 'Reasons and Rationality' organised by Michael Smith at the Australian National University in Canberra in June 2004.