Penultimate draft of Rowland Stout, "Internalising Practical Reasons", *Proceedings of the Aristotelian Society* 44:3, 2004, 229-43.

INTERNALISING PRACTICAL REASONS

by Rowland Stout

ABSTRACT Practical reasons figure in both the justification and the causal explanation of action. It is usually assumed that the agent's state of believing rather than what they believe must figure in the causal explanation of action. But, that the agent believes something is not a reason in the sense of being part of the justification of what they do. So it is often concluded that the justifying reason is a different sort of thing from the causally motivating reason. But this means that in a causal process of acting the justifying reasons have done their work by the time the agent has the appropriate beliefs and desires. Transforming these into behaviour is not guided by reason. This conception of action in which there is no role for reason in the part of the process where anything actually gets done is not acceptable. So the original assumption that beliefs rather than the believed facts figure in the casual explanation of action should be challenged.

It is a quite standard assumption in the philosophy of action, and one that I aim to challenge, that a proper account of agency must appeal to the causal role of psychological states – in particular beliefs and intentions or beliefs and desires. Donald Davidson's causal theory of action in 'Actions, Reasons and Causes' exemplifies this assumption.[1] And for many people working in this area the central

[1] Davidson, 1980, essay 1.

task is, as Al Mele describes it in *Springs of Action*, the 'exploration of the roles played by a collection of psychological states in the etiology of intentional behaviour'[2]

But Davidson's causal theory was the development of a deeper idea in the philosophy of action - the idea that a proper account of agency must appeal to the causal role of *reason* or *reasons*. We find this in Aristotle's conception of teleological causation. '"Why is he walking about?" We say: "To be healthy", and having said that, we think we have assigned the cause.'[3] Explaining something in terms of what it is for the sake of is to explain it in terms of a reason for it. And we find the idea also in Kant's conception of the rational will. 'Will is a kind of causality of living beings in so far as they are rational'.[4]

Now suppose the standard assumption that I am challenging were right and action had to be understood in terms of the causal roles of psychological states. And also suppose that action must be understood in terms of the causal roles of reasons for action. Then the neatest way to make sense of this would be to say that reasons for action are themselves facts about psychological states. My reason for going to the shop is not that I can buy milk in the shop; it is that I believe I can buy milk in the shop. This belief, combined with my desire for milk, both justifies and causally explains my going to the shop.

I will argue that this neat move fails to do justice to our notion of practical rationality. What really justifies my going to the shop is the fact that I can buy milk in

---

[2] Mele, 1992, 3. This assumption is also the guiding principle of cognitivist psychology, the result of putting aside behaviourist misgivings about theories with hidden variables to be filled by inner states.

[3] *Physics*, Book 2, 194^{b}34-5.

[4] *Groundwork*, 97 (2nd edition). Davidson's approach to action may be seen as an attempt to accommodate this Kantian insight in a materialist ontology.

the shop: the believed fact rather than the fact that I believe it. To treat practical rationality as simply concerned with the rational responses to different psychological states would be a crude form of psychologism every bit as bad as the psychologism that Frege objected to concerning the principles of logic.

Many philosophers concerned with practical rationality do not accept the neat solution of identifying practical reasons with facts about psychological states. But they still try to capture the idea that the rational explanation of action is psychological by arguing that the psychological states that figure in the causal explanation of action somehow internalise the reasons for the action. In this paper I will argue against these less crude ways to psychologise practical reasons too. In general I want to reject any attempt to account for the causal dimension of rationality by internalising practical reasons *as* or *in* psychological states. In particular I reject the idea that practical reasons must be internalised as or in *beliefs*.

I

Causal role of practical reasons. But I do want to retain the Kantian insight that rationality is involved in the causal process of acting. This insight seems to be rejected by anti-rationalists, like Brian O'Shaughnessy, who argue that there is a perfectly good notion of action with no constitutive role for rationality.[5] But even if there were a perfectly good notion of a-rational or non-rational acion it would not undermine the insight that a central notion in understanding agency is that of action for a reason. What counts as a reason is an open question at this stage, but it makes no sense to deny that reason in some sense has a constitutive role in intentional action.

[5] O'Shaughnessy, 1980, vol 1, 61.

Another threat to this Kantian insight comes from anti-causalists, who argue that explanations of what people do in terms of their reasons are not causal explanations. Anscombe (1957), while arguing that actions give application to a certain sense of the question 'Why?', one which demands an explanation which is at the same time a justification, denies that such explanation is causal. And this denial is part of a very powerful tradition in the philosophy of action including Ryle (1949), Melden (1961), von Wright (1971), as well as plenty of current philosophers, including in particular Jonathan Dancy (2000, ch. 8) whose recent book will be a focus for much of my argument in this paper.

In fact these anti-causalists do not always deny that in action there is some involvement of reasons in the causal process. They often have a particular model of what this causal involvement might be and reject that model. So Ryle claims that motives 'are not happenings and are not therefore of the right type to be causes. The expansion of a motive-expression is a law-like sentence and not a report of an event.' This amounts to no more than denying that explaining an action in terms of its motives is a form of *event* explanation. Equally Melden rejects only what he calls *Humean* causal explanation of action in terms of motives and intentions.[6]

But other anti-causalists leave less room for the possibility of reasons having any sort of causal role in action. For example Dancy rejects any sense in which normative explanation is causal.[7] This is partly because he takes normative explanation not to be factive, whereas any kind of causal explanation must be factive. And even if I am right in charitably (as I see it) attributing to some anti-causalists the possibility of accepting a sense in which reasons have a causal role in action, they certainly do not exploit this possibility. Although by and large these philosophers are

[6] Ryle, 1949, 113 and Melden, 1961, 102.

[7] Dancy, 2000, 159 ff.

very good at avoiding the kind of psychologism about reason that I am arguing against here, I think that the price of avoiding psychologism would be too high if it meant ignoring the causal role of reasons in action.

Now the anti-causalists would be right to point out that there is a big step from Kant's insight that action involves the causality of living beings in so far as they are rational to the claim that the reasons for action are themselves the causes. When you act for a reason you transform the world in the light of that reason. But it is you who transforms the world not your reason.

But think of this analogy. Gravity exerts a certain force of attraction on two bodies, where the force is determined partly by the distance between the two bodies. It is gravity and not the distance between the two bodies that is pulling them together. Still we do talk of the distance between the two bodies being a cause of there being a certain force acting on them. It is a cause in the sense of being something that figures in the causal explanation, not in the sense of being something with causal power. It is not an *efficient* cause in Aristotle's sense of that which makes what is made and that which changes what is changed.[8]

In the same way we can say that in acting, one's transformation of the world comes *in response to* and *as a result of* a reason. So the reason figures as a causal factor in an explanation of why the world has changed in that way. This means that a reason, while not necessarily being a cause in the sense of being something with causal power, is a cause in another sense. The point shows what is wrong with one of Dancy's attempts to cast doubt on the idea that normative explanation is causal. He asks rhetorically: 'Can a body be caused to move by the fact that one person owes

---

[8] Aristotle, *Physics*, 194b31-2.

another a favour?'[9] The answer of course is no; but it does not follow that the fact that one person owes another a favour cannot figure in the causal explanation of a body moving.

II

Crude psychologism. On the face of it, the reason I go to a shop might be that I can buy milk there. But it is often argued that the real reason is not that I *can* buy milk there, but that I *believe* that I can buy milk there. According to this argument, the rationality that figures in action is the rationality of making certain responses to certain psychological states. This is what I shall call crude psychologism.

Jonathan Dancy retains the term 'psychologism' for an even cruder view, but one that I think is genuinely incoherent and should not really concern us here. Dancy distinguishes between a conception of practical rationality according to which the agent's reasons are that they have certain psychological states and a conception of practical rationality according to which the agent's reasons are these psychological states themselves. Only the latter conception gets to be called 'psychologism' by Dancy.

But psychological states are never even candidates for being reasons in the sense of providing justifications. To the extent that a reason justifies an action there should be some inference available from the reason to the conclusion that that is the thing to be done. But you cannot have an inference from a psychological state like a state of believing, any more than you can have a justification from a mountain or a

[9] Dancy, 2000, 161.

pint of milk. You can only have an inference from something that is propositionally structured.

The matter may be different for causes. Psychological states, like someone's state of believing in ghosts for instance, may be said to cause things. And since reasons for action are supposed both to justify and to cause, this would seem to leave us in a quandary, since psychological states are not the right sorts of things to justify. But this quandary is very easily resolved. As I pointed out earlier it makes better sense to talk of reasons causally explaining actions than of them causing actions. The fact that someone believes in ghosts may causally explain their strange behaviour, and it is at least the sort of thing that might justify what they do.

Dancy does have an argument for the claim that there are two distinct types of reason that may be at play here – the psychological state and the fact that the agent is in such a state.

> 'Consider the difference between my nervousness, which is a mental state, and that I am nervous, which is not. My nervousness may explain my jumping whenever there is a loud noise; that I am nervous explains why I take beta-blockers when I have to sing in public, since my reason for doing so is that I am nervous, while my nervousness is not my reason for jumping at loud noises. (In fact I don't do this for a reason at all.) So explanations that appeal to mental states as *explanantia* are not equivalent to explanations that appeal to such things as "that the agent is in such-and-such a mental state".'[10]

[10] Dancy, 2000, 123.

But this argument does not show that my nervousness may be a reason in the sense of justifying what I do. On the contrary, as Dancy acknowledges, the thing that my nervousness is supposed to be a reason for – my jumping – is not something that has a justification. What Dancy does not provide is an example where a psychological state provides one kind of justification and the fact that the subject is in that state provides another. At best he has raised the possibility that psychological states may explain what people do in a way that does not justify what people do. This possibility is not sufficient to force us to take seriously the idea that psychological states themselves (rather than facts about them) figure in practical rationality.

Just to compound the possibility for confusion there is a perfectly good sense in which someone's belief justifies what they do. But this is not when their state of believing justifies what they do, but when *what* they believe justifies what they do. I may be able to infer that a dossier should be sent to the examinations committee from my belief that a student has plagiarised an assessed essay. But this is not a case of inferring something from my state of believing something, but a case of inferring something from *what* it is I believe – namely the *fact* that the student has plagiarised that essay. If we mean this when we talk of a belief justifying then we are not guilty of a crude psychologism after all.

So, what can justify an action is either what is believed or the fact that the agent is in that state of believing. The latter is the crude psychologistic view, and I think it presents a badly distorted picture of practical rationality. Although it is sometimes the case that we have to take our own psychological attitudes into account when deciding how we should behave, these are rather special cases. Generally we have to take into account conditions in the world that apply quite independently of our attitudes to them. The rational sensitivity that constitutes agency is not in the first instance rational sensitivity to our own psychology.

Dancy presents an example of someone whose reason for action is that they have a certain belief:

> 'For instance, that I believe that the cliff is crumbling is my reason for my avoiding climbing it, because having that belief I am more likely to fall off (I will get nervous). This is a case where that I believe what I do is genuinely my reason for action, in a way that is independent of whether my belief is actually true. … I recognize that if the cliff were not crumbling, I would still have just the same reason not to climb it as if it were, so long as I continue to believe it to be crumbling. But this is quite an unusual situation, not at all the normal case.'[11]

Seeing how unusual this sort of case is leads us to see that in normal cases this cannot be what is going on. There is a genuine contrast between the case where the reason for action is that one believes that the cliff is crumbling and the case where the reason is just that the cliff is crumbling. And the latter is the standard sort of case. The fact that we have certain beliefs (and likewise desires) does not normally constitute reason for our action.

Consider the rule that justifies a football referee blowing the final whistle. The rule is not the following:

> Rule 1. The referee should blow the final whistle if and only if they *believe* that 90 minutes of non-injury time have been played.

[11] Dancy, 2000, 124.

It is the following:

> Rule 2. The referee should blow the final whistle if and only if 90 minutes of non-injury time *have* been played.

If you blow the whistle after 85 minutes have been played, mistakenly thinking that 90 minutes have been played, then you have made a mistake – you have failed to apply the rule that should have been governing your behaviour. But if Rule 1 were the rule that should have been governing your behaviour, then this would not have been a mistake.

If you act on the false belief that 90 minutes of time have been played you are still acting rationally in a sense. But the rule that explains why you are acting rationally is not Rule 1 but Rule 2. You are doing what this rule tells you to do on the assumption (false assumption) that 90 minutes of time have been played. The fact that you are working on this assumption does not contribute to the rationality of your behaviour.

So, what are we to make of your saying: 'The reason I blew the whistle was that I believed 90 minutes had been played'? This would be said as part of a justification or at any rate an excuse for what you did. But I have just argued that the fact that you believed that 90 minutes had been played does not itself contribute to a justification of your behaviour. This seems to be a straightforward contradiction.

Dancy considers two options here. The first is that the fact that the agent believed something is not a proper part of the explanation of what they did, but is an 'enabling condition for an explanation which explains the action in terms of the

reasons for … doing it.'[12] I shall put forward some considerations in favour of something like this option shortly.

While Dancy has no argument against this idea he prefers to follow up another option. This is what he calls the appositional account. 'This hears "He is doing it because he believes that *p*" as "He is doing it because *p*, as he believes".'[13] The referee blows the final whistle because 90 minutes have been played, as he believes. And this is supposed to be the case even when 90 minutes have not been played.

So, for Dancy, rational explanation of action is not *factive*: i.e. we can explain an action by something which is not true. This is an extraordinary view. It is important to remember that Dancy does not regard such explanation as causal, and this may help make the position more palatable. But I do not think it really does. A true explanation must deal in truths. Equally a justification must deal in truths. I can offer a *possible* justification of what the referee did by saying that 90 minutes have been played. But this is not an actual justification since 90 minutes have not been played. Rational explanation of an action must reveal that action's sensitivity to reasons in the world. According to Dancy's suggestion the correctness of an explanation of action is independent of what is actually the case. And this marks a failure to see practical rationality as properly world-involving, every bit as bad as psychologism.

Dancy attempts to provide intuitive support for his non-factive conception of action explanation by switching quite illegitimately from talking of *the* reason he did it to talking of *his* reason for doing it.[14] It sounds much better to say that your reason for blowing the whistle was that 90 minutes had been played, although in fact 90

[12] Dancy, 2000, 127.

[13] Dancy 2000, 128-9.

[14] Dancy, 2000, 132.

minutes had not been played, than it does to say that the reason you blew the whistle was that 90 minutes had been played, although 90 minutes had not been played. This is because 'your reason' can be taken to mean something like 'what you took the reason to be'.

This suggests a possible line of thought which is that you actually have no justification at all for blowing the whistle; it is just that you are in a situation in which you think you have a justification for blowing the whistle. Assuming 90 minutes had been played you would have had a justification. But since 90 minutes have not been played you only have a possible – a *might have been* – justification. When you say that the reason you blew the whistle was that you believed 90 minutes had been played what you should have said instead is that you believed the reason you were blowing the whistle was that 90 minutes had been played.

But this line of thought may be too uncompromising. I agree with Dancy, and indeed with the psychologistic account, that there is at least some way of justifying your early whistle blowing. But I do not think it is by appealing to the non-fact that 90 minutes have been played. Instead the non-fact that 90 minutes have been played is an *assumption* structuring a way of justifying your behaviour. The way of justifying your behaviour *works on that assumption*. And according to this way of justifying things, blowing the whistle was the thing to do. Since the assumption that this way of justifying things works on is false, it is a bad way of justifying things, but a way nonetheless.

So, starting with a way of justifying things that would justify your behaviour assuming 90 minutes had been played, we can derive a more limited way of justifying things that embeds the assumption that 90 minutes have been played. This will justify your whistle blowing whether or not 90 minutes have been played.

Saying that you believed that 90 minutes had been played is the same as saying that you were working on the assumption that 90 minutes had been played. And this is the same as saying that the way of justifying behaviour that was causally operative with you was one that worked on that assumption. So saying that you believed that 90 minutes had been played is to say something about a causally operative way of justifying your behaviour. That you had the belief you did is not some kind of input into a way of justifying what you did; it is a fact about the structure of the way of justifying behaviour that was causally operative with you.

When you say that the reason you blew the whistle was that you believed 90 minutes had been played, this is what I think you should more properly have said:

> 'I believed that 90 minutes had been played and that meant that I had reason to blow the whistle, and that was why I blew the whistle.'

Treating the fact that the agent has some relevant belief in this way as a sort of enabling condition for a rational explanation of what they did is to allow some sort of role for that psychological fact in the causal explanation of their behaviour. But it is to see the psychological fact as a fact about how someone's process of acting works. It is not to see that fact as an input into such a process.

III

Sophisticated psychologism. The view that your real reasons for action are facts about your beliefs and desires (or intentions or inclinations or whatever) is often attributed

to David Hume.[15] It is also attributed to Donald Davidson, who claimed that a primary reason for an action is a belief and a pro-attitude paired together.

Davidson argues for this claim from his principle that 'a reason rationalises an action only if it leads us to see something the agent saw, or thought he saw, in his action.'[16] But this principle by itself should not force us to psychologize practical reasons. The reason for an action might be some fact about how things are outside the agent. By seeing that such a thing is a reason for the agent we are thereby led to see something the agent saw or thought they saw in their action, just as Davidson's principle requires. The reason does not *itself* have to be the agent's seeing or thinking they are seeing something in their action.

There is a more sophisticated way to make the internalising move in the philosophy of action. It is to concede that reasons for action need not be merely psychological facts, but may be conditions obtaining in the world outside, but at the same time to require that these reasons and the laws that demand sensitivity to them must be internalised in the agent's attitudes as part of the *causal* process that results in that agent's behaviour. On this view, the structure of rationality is, as we normally take it be, a properly world-involving thing. But it must be internalised in the agent's attitudes before it can have a causal role for that agent. An attitude then provides a reason for an action only in the sense that for a person with that attitude that action would be rational. It would be rational in virtue of the world-involving structure of rationality internalised by that attitude.[17]

---

[15] See for example Smith, 1994, chapter 5.

[16] Davidson, 1980, 3.

[17] This view has such widespread acceptance that it is really quite arbitrary to pick this or that philosopher who does accept it. For some arbitrarily chosen examples consider John Broome, 2001,

Michael Smith, in an attempt to endorse this sort of position, has distinguished between what he calls motivational reasons and normative reasons.[18] Motivational reasons, according to Smith are explanatory and not justificatory. They do not contribute to the *rationality* of the action just to the *intelligibility* of the action. He argues that normative reasons are truths and not psychological states, but that only psychological states provide motivational reasons.

Smith's terminology here is certainly a bit misleading. Given that motivational reasons are merely explanatory and not justificatory it is wrong to say that they are reasons for the agent *to* act in certain ways. This is wrong in the same way as it would be wrong to say that a large mass being close to a stone is a reason for the stone to accelerate towards the mass, or to say that my nervousness is a reason for me to jump.

Dancy provides an alternative and more satisfactory way to distinguish motivating and normative reasons.[19] According to Dancy, normative reasons are reasons that *favour* the action, and motivational reasons are reasons *in the light of which* an agent acts. But if motivational reasons are understood in this way then they cannot be radically different kinds of things from normative reasons. The reasons in the light of which one acts must, at least potentially, be reasons that favour one's action. So if motivational reasons are facts about psychological attitudes, then so are normative reasons, and we are back with the first way of making the internalising move – crude psychologism - that presents such a distorted conception of practical rationality.

175, Robert Brandom, 1994, 29 ff., Alan Gibbard, 1990, chapter 3, Cullity and Gaut, 1997, 2, Bernard Williams, 1981.

[18] Smith, 1994, chapter 4.

[19] Dancy, 2000, 98 ff.

But in response to this someone like Smith might insist that what he means by a motivating reasons is not a reason in the light of which one acts, but is a reason that just causally explains one's acting. So he might concede (although this would contradict what he actually says)[20] that the motivating reason is not a reason for the agent to act. At the same time he would have to accept that the normative reason and the motivating reason are related. Somehow we can see the normative reason in the motivating reason.

The obvious way to make this move is to claim that the motivating reason is the fact that the agent has a certain belief and the normative reason is the fact that is represented by that belief. Dancy considers something like this line of thought (though at this stage in his argument he is assuming that the motivating reason is to be the belief itself rather than the fact that the agent has it). His objection is this:

> 'My response to this version of the content-based strategy is that it simply awards itself the concept of a good motivating reason, without really doing anything to show that it makes sense. It awards itself this prize because, given the terms of the debate, it needs to do so if it is to show a good sense in which a motivating reason can enjoy a normative status.'[21]

But this response really fails to engage with Smith's much more extreme view that a motivating reason does not enjoy a normative status. Its status is entirely that of a causal explanation. Dancy does not accept that an explanation of action that reveals reasons is causal at all. 'If causal explanations of action are to be possible, I must say,

---

[20] Smith, 1994, 95.

[21] Dancy 2000, 119.

at least they cannot include psychologistic ones except in other than a reason-specifying style.'[22] Although Smith's motivating reasons are not normative, they are at least in a 'reason-specifying style'. Given this Dancy can find no space for Smith's conception of a motivating reason. But as far as I can see he has not argued against it.

The point would be much clearer if we dropped the misleading terminology of motivating *reasons*. According to this internalising strategy, in order for an agent to be motivated by reasons these reasons must be internalised in psychological attitudes – beliefs, desires, intentions, etc. The psychological attitudes are not themselves taken to be reasons that *contribute* to the rationality of the action, although their existence *reveals* the rationality of the agent. These internalisations of reasons are supposed to be what cause the agent to act. And the fact that the agent has such attitudes causally explains their acting that way.

This internalising strategy does not give us such a distorted conception of practical rationality as the simpler Humean strategy. However I think the real objection to it is that, by separating reasons and causes for action, it gives us a distorted conception of the causal process that constitutes action. For although an action may be justified by facts about the world on this view, such a justification does not figure in the causal explanation of what the agent does. An action, then is not caused in line with a justification by facts about the world, but is caused by something that does not justify it at all – namely a set of psychological attitudes.

Strangely, these psychological attitudes are themselves usually taken to be both caused and justified by facts about the world. So the process of reasoning may involve a genuine transformation of things in line with reason. But this view fails to take seriously the idea of *action* as a process of rationally transforming the world.

---

[22] Dancy, 2000, 167.

Instead the only transformations that characterize the rationality of agency are taken to consist in the production of psychological attitudes. The move from these attitudes to the transformation of the world outside is not governed by rationality. Reason has done its work before action gets started.

So first you have reasons doing their stuff; then you have a causal response to these reasons. The transformation of the world is not taken to be a manifestation of the reasons in action, but rather a response to such a manifestation. Action is then regarded as a two-stage process. It involves the rational production of intentions and beliefs; then it involves the merely causal response to these attitudes.

According to this internalising picture, that aspect of action that involves transformation of the world is just subjection to a psychological law: do what you believe you should do; or do what you intend to do. But there is no other aspect of action, since the first stage - the production of the attitudes - is not strictly speaking part of action at all, as it does not involve making anything actually happen.[23] According to this picture, rational agency consists of two stages - a rational non-active stage and then a non-rational active stage bolted together. But this does not look like a picture of agency at all.

For neither the sophisticated psychologistic philosopher like Smith nor the anti-psychologistic philosopher like Dancy do reasons (understood normatively) have any causal role in action. What neither even considers is the possibility of causal processes that are genuinely sensitive to reasons – and that a person's acting is such a

[23] Some philosophers have argued that the work of agency is over before the body even starts to move. For example, Pietroski (2000), developing Hornsby's (1980) identification of acting with trying, argues that actions are *inner causes* of bodily movement. But I take it that this is a position only to be adopted if one has painted oneself into a corner, and my task here is to see if we can step out of that corner.

process. Of course, making sense of such a possibility faces all sorts of challenges, but it is at any rate not just obvious that such challenges cannot be met.[24]

Department of Philosophy

University College Dublin

Dublin 4

Eire

rowland.stout@ucd.ie

REFERENCES

Anscombe, G. E. M. 1957. *Intention*. Oxford: Blackwell.

Brandom, R. 1994. *Making It Explicit*. Cambridge Mass.: Harvard University Press.

Broome, J. 2001. 'Normative Practical Reasoning', *Proceedings of the Aristotelians Society supp. Vol*. 75, 175-193.

Cullity, G. and Gaut, B. (eds.) 1997. *Ethics and Practical Reason*, Oxford: Oxford University Press.

Dancy, J. 2000. *Practical Reality*. Oxford: Oxford University Press.

Davidson, D. 1980. *Essays on Actions and Events*. Oxford: Oxford University Press.

Gibbard, A. 1990. *Wise Choices, Apt Feelings*. Oxford: Oxford University Press.

Hornsby, J. 1980. *Actions*. London: Routledge.

Mele, A. 1992. *Springs of Action*. Oxford: Oxford Univeristy Press.

O'Shaughnessy, B. 1980. *The Will*. Cambridge: Cambridge University Press.

Smith, M. 1994. *The Moral Problem*. Oxford: Blackwell.

[24] I try to meet some of these challenges in Stout (1996).

Williams, B. 1981. ‘Internal and external reasons’, in his *Moral Luck*. Cambridge: Cambridge University Press.