

Agency and Action

ROYAL INSTITUTE OF PHILOSOPHY SUPPLEMENT: 55

EDITED BY

John Hyman and Helen Steward



CAMBRIDGE
UNIVERSITY PRESS

correctness does not mean we are less completely guided by reason than we should be. In such an area, then, we can be autonomous by being the authors of our own law. But since morality is an area in which legitimate questions arise about the correctness of imposing this law rather than another on ourselves, autonomy does not consist in being the authors of our own law. It consists in being guided by our rational grasp of the independently correct law.

Such an account of moral autonomy leads us to a further question. We might agree with Kant's judgment that morality is an area in which questions about the correctness of a specific act of legislation seem to be legitimate, and that this is why it seems to be an area in which we cannot be the authors of our own law. But if we become convinced that the questions that seem to be legitimate are really illegitimate, because we cannot give any reasonable answers to questions about the correctness of this or that legislation, we will conclude that our initial conception of moral autonomy rested on an illusion.

In saying this I am repeating the earlier point that our conditions for autonomy in a given area must reflect our views about the sort of law that is available in that area. I have avoided discussion of Kant's account of the correct law in the area of morality; but we cannot avoid discussion of it if we want to decide about his account of moral autonomy.

The Structure of Orthonomy*

MICHAEL SMITH

According to the standard story of action, a story that can be traced back at least to David Hume (1740), actions are those bodily movements that are caused and rationalized by a pair of mental states: a desire for some end, where ends can be thought of as ways the world could be, and a belief that something the agent can just do, namely, move her body in the way to be explained, has some suitable chance of making the world the relevant way. Bodily movements that occur otherwise aren't actions, they are mere happenings (Davidson 1963, Davidson 1971).


The feature that bodily movements have that makes them especially suitable to count as actions, as distinct from (say) the ends that agents desire, is that they are the events in the world over which agents are supposed to have direct rational control: agents' abilities bottom out with movements of their bodies (Danto 1963). This is why the occurrence of an action, as distinct from a mere happening, does not depend on whether the agent's bodily movement causes and rationalizes the end she desires. That merely affects whether that which she directly controls, the movement of her body, can be redescribed as the action of intentionally bringing about the desired end (Davidson 1971).

We can therefore sum up the roles of the different elements in this Humean story of action in the following terms. When desires for ends and means-end beliefs combine to cause and rationalize bodily movements in the way required for direct control by an agent—this is what the 'DC→' signifies in figure 1—then those bodily movements count as actions of that agent.

Much has been said both to clarify and to defend this standard Humean story (Hornsby 1980, Peacocke 1979, Smith 1987, Wilson 1989, Mele 1992, Smith 1999). Though more doubtless needs to be

* An earlier version of this paper was presented at the Royal Institute of Philosophy conference on *Action and Agency*, St John's College, Oxford, September 2002, and as one in a series of three Erskine Lectures delivered at the University of Canterbury, Christchurch, March 2003. I would like to thank Jonathan Dancy, Jennifer Hornsby, John Hyman, Cynthia Macdonald, Graham Macdonald, Philip Pettit, John Tasioulas, Jay Wallace, Ralph Wedgwood, and Bernard Williams for their many useful comments.

Figure 1. The standard Humean story of action



The crucial elements are in bold.

said, my aim here is not to add anything further by way of clarification or defence. Rather, taking the story for granted, I wish to raise a question concerning the status of actions themselves. If the rational etiology of a bodily movement—its genesis in desires for ends and means-end beliefs—is what turns a mere happening into an action, an event under the control of an agent, then does the rational etiology of those desires for ends and means-end beliefs turn an action into something which is under the control of an agent in some more significant way still?

More specifically, does the rational etiology of the desires for ends and means-end beliefs that produce an action turn a *mere action* into an *orthonomous action* (Pettit and Smith 1990, Pettit and Smith 1993, Pettit and Smith 1996, Smith 2003)? In other words, does the rational etiology of these desires and beliefs show the extent to which the action was performed by an agent who exercises his *orthonomy*, that is to say, his distinctive capacity to get things *right* or *correct* (this is the 'orthos')? As we will see, philosophers of quite different persuasions can and should agree that the rational etiology of the desires for ends and means-end beliefs that produce action in the standard Humean story does indeed have this kind of import. Moreover, as will emerge, this fact about the rational etiology of desires for ends and means-end beliefs provides us with an illuminating picture of both responsibility and control. Interpreted ecumenically, it provides a framework in which we can recognize, and in terms of which we can make sense of, a whole range of cases in which people act correctly, thereby displaying their responsibility.

The alternative and more standard account of responsibility and control asks whether an agent manifests his capacity for *autonomy* (Christman 1989). But, as a casual glance at the literature on autonomy reveals, such value as we place on autonomy looks to be entirely derived from the value we place on orthonomy (Watson

1975, Wolf 1990). Autonomy is the mere capacity an agent has to rule himself (this is the 'autos'). But we value agents ruling themselves to just the extent that, in so doing, they thereby manifest their capacity to get things right (Pettit and Smith 1993). The main aim of the present paper is thus to lay out the structure of orthonomy. My hope is to encourage philosophers to think in terms of this structure—to think in terms of orthonomy, not autonomy—when they engage in further theorising about the nature of responsibility and control.

The paper is in four main sections. In the first I describe a minimal conception of orthonomous action. In the second section I explain why, at least according to common sense, we must go beyond this minimal conception. In the third section I argue in favour of the more expansive, common sense, conception of orthonomous action by showing how it can be underwritten in both realist and irrealist terms. And then, in the fourth and final section, I show how the common sense conception of orthonomous action provides us with framework in which we can recognize, and in terms of which we can make sense of, a puzzling case of responsibility and control that has been much discussed in the recent literature.

1. Radical Humeanism

As I said, the standard Humean story tells us that actions are the product of a desire for some end and a belief that something the agent can just do has some suitable chance of bringing that end about. So far, however, this story tells us nothing about the rational etiology of the desires and means-end beliefs themselves. That requires a supplement to the standard Humean story.

According to the most minimal supplement, the one that Hume himself would presumably have accepted, the desires and means-end beliefs that produce actions differ significantly in their rational etiology (Milgram 1995). Because means-end beliefs are representations of how things are, the radical Humean view is that they are evaluable in terms of truth and falsehood, and, as a consequence, that they are also evaluable in terms of how well or ill justified they are, given the evidence available to the subject. This means that, assuming that rational subjects have the capacity to grasp the evidence available to them and to form beliefs that accord with that evidence, the justification of subjects' beliefs turns on whether or not they have exercised such evidence-grasping and belief-forming capacities as they have. The exercise of these capacities, in turn, is

what underwrites the possibility of rational criticism of the actions they perform. Actions are rationally criticizable to the extent that the beliefs that produce them are rationally criticizable.

Of course, since agents differ in their capacities—since some are more capable than others of evaluating evidence, for example—it follows that there is another dimension to the evaluation of subjects as well, namely, an evaluation in terms of how capable they are. But an evaluation of subjects of this kind does not constitute criticism of them because such an evaluation does not suggest that the subjects fail in respects in which they could have succeeded. Though 'ought' implies 'can', 'acted badly' does not imply 'could have done better'. What makes criticism apt is thus a subject's failure to exercise such capacities for grasping evidence and forming beliefs as she has.

By contrast, however, at least according to the most radical version of the Humean story, the desires for ends that produce actions—that is to say, those desires that are not themselves derived from further desires for ends and beliefs about means (Smith 2004: more on this later)—are, in Hume's memorable phrase, 'original existences' (Hume 1740, 415). This means not just that they elude evaluation in terms of truth and falsehood, but also, and supposedly as a consequence, that they are not evaluable in terms of how well or ill justified they are either. Agents may have conflicting desires for ends, of course, but not even these conflicts are supposed to be resolved via recourse to rational principles. Rather they are resolved more or less mechanically by the causal force of the desires involved. Unlike beliefs about means-ends, then, desires for ends are supposed to elude rational criticism altogether (Smith forthcoming).

Spare though it is, this account of the rational etiology of desires for ends and beliefs about means to ends provides us with all we need in order to provide an account of the structure of fully orthonomous action, at least by the lights of radical Humeanism. It provides us with an interpretation of the claim that an agent manifests his capacity to get things correct. When an agent acts in a fully orthonomous way—when he manifests his capacity to get things correct—the following conditions must all obtain: facts about which of the means available to the agent are means to which of the ends he desires must impact, in the right kind of non-accidental way, on the evidence available to him; this available evidence must in turn fix, in the right kind of non-accidental way, his beliefs about which of the means available to him are means to which of his desired ends; these means-end beliefs must then join together, in the right

kind of non-accidental way, with a desire for an end to rationalize the appropriate bodily movement, that is, a bodily movement picked out as an available means to that desired end; and this bodily movement must produce the desired end, again in the right sort of non-accidental way.


What exactly does it mean to say that all of these things must come about in the 'right kind of non-accidental way'? Facts produce evidence in the right kind of non-accidental way when they produce evidence in the way required for beliefs formed on the basis of that evidence to count as knowledge (this is what the 'K→' signifies in figure 2), and bodily movements cause desired ends in the right kind of non-accidental way when they produce desired ends in the way required for the bringing about of those desired ends to count as intentional conduct: that is, roughly speaking, when they come about in the way specified in the agent's pattern of practical reasoning (this is what the 'PPR→' signifies in figure 2). An agent's possession of the capacity to get things correct may therefore require that he be situated in a suitably obliging world, where a world is suitably obliging if it is one in which agents can have knowledge and act intentionally. Plainly not all worlds are suitably obliging in this sense: think of the myriad examples of wayward causal chains that undermine both agent's having knowledge and their acting intentionally.

But for other events to happen in the right kind of non-accidental way what is required is not so much that the agent be situated in a suitably obliging world, but rather that the agent himself meets a certain condition of ideal rationality. Available evidence produces belief in the right kind of non-accidental way, for example, when the formation of the belief, in the light of that evidence, counts as an exercise of the ideal capacity to rationally believe in accordance with the evidence available to him (this is what the 'IR→' signifies in figure 2), and the agent's desires for ends and means-end beliefs cause his body to move in the right kind of non-accidental way when the movement of the body counts as an exercise of the capacity to directly control his body (this is what the 'DC→' signifies in figure 2). The question we must ask is thus whether the agent has and exercises these rational capacities.

We can picture the radical Humean conception of fully orthonomous action in the following way (see figure 2).

The radical Humean conception of fully orthonomous action thus plainly allows that various things can go wrong when an agent acts. For example, some of the connections required for an action's being the upshot of the agent's distinctive capacity to get things

Figure 2. The radical Humean conception of fully orthonomous action



The crucial elements are again in bold.

right might not be in place. The available evidence about means-ends might be misleading, say, because not appropriately related to the facts about means-ends. The agent might fail to form beliefs about means-ends in the light of the evidence available to him. The agent's desire for an end and means-end belief might fail to cause a bodily movement. Or the events caused by the agent's bodily movement might not be the agent's desired ends. Alternatively, though all of these connections are in place, some of them might be in place by fluke or accident. For example, though the agent's beliefs about means-ends might accord with the available evidence, perhaps he doesn't have that belief as a result of having exercised his belief-forming capacities. Or though the desired end is caused by the agent's bodily movement, it might happen via a wayward causal chain. In that case, though the agent gets things right, his action would not be a manifestation of his distinctive capacity to get things right.

With this catalogue of things that can go wrong before us, it becomes clear that there is a natural division within this class. For whereas some of the things that can go wrong are things for which it makes sense to hold the agent himself responsible, at least assuming that he has the requisite rational capacities, others plainly aren't. Whether an agent forms appropriate means-end beliefs in the light of the available evidence, for example, is plausibly a failure for which he could, in principle, be held responsible. For all that is required for the formation of appropriate beliefs, given that the evidence for forming those beliefs is available, is that the agent fully exercises his rational belief-forming capacities, again, assuming he has such capacities.

But other things that can go wrong are plainly not things for which it would make any sense at all to hold the agent responsible. Whether the facts about means to ends manifest themselves in the evidence available to an agent is plausibly an example of a failure of this kind. The mere exercise of an agent's rational capacities may be insufficient to ensure that such facts manifest themselves to him. Good luck is required: location in a suitably obliging world. Likewise, whether or not the bodily movement that is caused by an agent's desires for ends and means-end beliefs produces the desired end may not be something that the agent can control once that bodily movement has been performed. The mere exercise of his rational capacities will have no effect on interference that occurs beyond his body when such interference can't be anticipated. Again, good luck is required: location in a suitably obliging world. (The reason for the qualifications—the 'may's—is that (say) the failure of the facts to manifest themselves in evidence could be the result of some prior action of the agent's of hiding the facts from himself, in which case he would be responsible. Likewise for the failure of his bodily movement to cause and rationalize the desired end.)

It thus follows that fully orthonomous action itself turns out to be, in part, a matter of good luck. But it also follows that we can define a more partial kind of orthonomy, a kind that will always be a legitimate aspiration for an agent given that agents are, by definition, capable of the exercise of such rational capacities as they have. This more partial kind of orthonomy—let's call it 'narrow orthonomy'—is a matter of agents' getting things right to extent that it is up to them whether or not they get things right: that is, to the extent that they fully exercise such rational capacities as they have. This is a capacity an agent has whether or not he is situated in a suitably obliging world and no matter how impoverished his rational capacities are. All that matters is that he fully exercises such capacities as he has to rationally manage his own psychology.

Narrow orthonomy thus looks like it will be of the greatest significance in philosophy, as it promises to provide a relatively straightforward and uncontroversial account of the nature and scope of responsible conduct. For, according to this account, the concept of responsible conduct simply piggy-backs on the concept of responsible believing and desiring. Responsible conduct is a matter of an agents' controlling their behaviour in the sense of acting after having exercised such rational capacities as they have, capacities they have to get things correct rather than incorrect independently of whether they are in a suitably obliging world. In order to produce such an account of responsible conduct all we need to do

is to provide a fully spelled out story of the nature and scope of an agent's rational capacities.

Importantly, however, and especially in the light of the potential significance of narrow orthonomy for the understanding of responsible conduct, radical Humeanism suggests that there is a further natural division within the class of things that can go wrong. For whereas there are various ways in which things can go wrong in the way in which means-end beliefs and desires for means are produced—the beliefs formed may not be true, or they might not be formed in the light of the available evidence, or the agent might irrationally fail to form a desire for the believed means to his desired end—and whereas some of these are the agent's own responsibility—it is up to the agent to exercise such capacities as he has to believe what is supported by the available evidence, and it also up to him to exercise such capacities as he has to desire the believed means to his desired ends—there is simply no way in which things can go wrong as regards the production of an agent's desires for ends. This is not to deny that desires for ends are caused in various ways, for of course they are. It is simply to say that, at least according to radical Humeanism, it does not matter to the narrow orthonomy of the resulting action how an agent's desires for ends are caused. This is the cash value of Hume's claim that such desires are original existences. Because desires for ends are not the product of a rational capacity, they are not subject to a correctness condition. Desires for ends are not subject to rational control.

When the point is put as bluntly as this it will doubtless raise eyebrows. 'Surely', it will be objected, 'the fact that an agent's desires for ends are not the product of a rational capacity doesn't suffice all by itself to show that they are not subject to a correctness condition. To think that it does is to be taken in by the ambiguity of the phrase "correctness condition". The existence of a correctness condition for desires for ends turns on whether or not there is some standard or other by which we can assess such desires. But while it might be agreed that *rationality* does not provide such a standard, that is consistent with there being some other system of norms that does provide such a standard: norms of *morality*, say.' The objection is a good one, as it forces us to sharpen the point.


Though it is true that we could relativize orthonomy to a system of norms, we mustn't lose sight of the fact that what makes orthonomy of such special interest is its potential to provide an account of responsible conduct. This constrains our choice of a system of norms relative to which orthonomy can be judged. Being responsible is, after all, an inescapable feature of the circumstances we face

as rational actors. But it follows from this that we must therefore evaluate orthonomy relative to a system of norms to which we are subject, as rational actors, in a similarly inescapable way. Given that being subject to norms of reason is constitutive of being a believer and desirer, and given that such norms provides a natural interpretation of what it means to get things right, this makes the norms of reason the obvious choice. Conversely—at least if the norms of morality do not reduce to the norms of reason (see Smith 1994 for the alternative view)—given that being subject to norms of morality is not similarly constitutive of being a believer and desirer, it follows that the norms of morality would be a very poor choice of norms relative to which we should judge orthonomy. For in that case there would be no grounds on which to base the inescapability of responsibility: no basis for the universal expectation that people do in fact get things right, at least to the extent that they can.

We can represent this radical Humean conception of narrow orthonomy in the way suggested in figure 3. According to this conception, it is irrelevant to the narrow orthonomy of an agent's action whether the evidence available to him is misleading, and whether the end the agent desires fails to result from his bodily movement due to some unforeseeable interference. All that matters is that the agent exercises such rational capacities as he has: his actual rational capacities (this is what the 'AR→' signifies in figure 3).

However, and to repeat, the distinctive feature of the radical Humean conception of narrow orthonomy is that though it insists that the exercise of such rational capacities as an agent has will have all sorts of effects on which means-end beliefs he has, it tells us that it will have no similar effect on which desires for ends he has. This is because desires for ends are not the product of an agent's rational capacities. In assessing the radical Humean conception of narrowly orthonomous action, the question we must ask is thus whether this is really credible.

Figure 3. The radical Humean conception of narrow orthonomy



2. Common sense

As we will see, the problem with the radical Humean conception of narrow orthonomy is that it is inconsistent with a widely accepted and common sense view about the relationship between two kinds of reasons for action: motivating reasons and normative reasons (Woods 1972, Smith 1987, Smith 1994). The point will take a little explaining, however, so let me begin by making the distinction itself.

It is now familiar that we can distinguish between two kinds of reasons for action: normative reasons and motivating reasons. Motivating reasons are psychological states that teleologically explain what an agent does (Wilson 1989). If, as many think, such teleological explanations reduce to a kind of causal explanation, then this amounts to the claim that motivating reasons are psychological states that cause and rationalize an agent's doing what he does (Davidson 1963). Unsurprisingly, these psychological states are thus the desires for ends and means-end beliefs that we have been talking about so far in spelling out the standard Humean story. Normative reasons, by contrast, are the considerations to which we appeal when we construct a justification of an agent's conduct (Nagel 1970, Dancy 2000). The difference is that captured by the following pair of claims, both of which are true at the moment of my writing:

- (i) My reason for tapping on the keys to my laptop is that I want to write a paper about philosophy of action and believe that I can do so by tapping on the keys to my laptop (this is the motivating reason)

and

- (ii) My reason for tapping on the keys to my laptop is that people are counting on me to write a paper about philosophy of action (this is the normative reason).

To repeat, the distinction is that between psychological states that teleologically explain and considerations that justify.

The importance of making this distinction in this way becomes clear when we ask whether all actions must be done for reasons. For though this question gets answered resoundingly in the affirmative when reasons are understood to be motivating reasons—as we saw at the outset, what makes an action an action is the fact that there is a desire-belief pair that causes and rationalizes it—the question gets answered just as resoundingly in the negative when reasons are

understood to be normative reasons. An agent can act without there being any considerations at all that justify his doing what he does (and, just to forestall any concern about equivocation, let me clarify that from now on all talk of justification will be talk of rational justification). He can, after all, be mistaken about which considerations provide a rational justification of his conduct, and he can be mistaken about whether or not such considerations obtain in the circumstances in which he acts.

Moreover, and perhaps even more strikingly, an agent can act without believing that there are any considerations that rationally justify his doing what he does, and even while believing that the considerations that bear on what he is doing all dysjustify (to borrow an excellent term of Michael Stocker's (2004)), rather than justify. For example, though I am tapping on the keys to my laptop because I want to write a paper about philosophy of action and believe that I can do so by tapping on the keys to my laptop, I may yet believe, perhaps correctly, that I am required *not* to write such a paper: that everything is to be said against, and nothing in favour, of my doing so. Of course, to continue writing a paper about philosophy of action in such circumstances I would have to be rather perverse and irrational. But some people clearly are perverse and irrational in exactly this way (Stocker 1979, Smith 1999).

The distinction between motivating and normative reasons needs to be handled with some care, however. For, as Bernard Williams has pointed out, facts about considerations that rationally justify cannot be divorced entirely from considerations of explanation (Williams 1980). When someone has a normative reason to act in a certain way this must be the sort of consideration that could figure in an explanation of her conduct, if not on that occasion, then at least on others. In the light of the huge gulf between the two sorts of reason, however—motivating reasons are psychological states that explain, normative reasons are considerations that justify—we must ask how that connection gets forged.

The obvious answer to give is that the connection gets forged by a rational requirement of response upon recognition of a normative reason (Pettit and Smith forthcoming). What makes a normative reason a normative reason is, *inter alia*, the fact that, when agents recognize that they have such a reason, they are rationally required to respond appropriately by acquiring corresponding motivations. In other words, reason requires those who judge that they have a normative reason to act in a certain way to be motivated to act in that way. A failure to be so motivated indicates that the agent suffers from weakness of will, or compulsion, or some other such

form of practical unreason. This explains the possibility of perverse agents who act on considerations that dyjustify.

This common sense view of the distinction between normative reasons and motivating reasons and their relations is widely accepted in the philosophical literature. Indeed, it is common ground among both those who are sympathetic to a Humean conception of human psychology, and those who are hostile such a conception. Here, for example, is Williams, a Humean.

Does believing that a particular consideration is a reason to act in a particular way provide, or indeed constitute, a motivation to act? ... Let us grant that it does—this claim indeed seems plausible, so long at least as the connexion between such beliefs and the disposition to act is not tightened to that unnecessary degree which excludes *akrasia*. (Williams 1980, 107)

On the plausible assumption that Williams takes *akrasia* to be a form of practical unreason, this amounts to the claim that the norms governing practical reasoning require that those who believe that such-and-such provides a reason to act in a certain way are motivated to act in that way.

Christine Korsgaard, a Kantian, puts forward a similar view.

Thus, it seems to be a requirement on practical reason claims that they be capable of motivating us. ... Practical-reason claims, if they are really to present us with reasons for action, must be capable of motivating rational persons. (Korsgaard 1986, 11)

and she elaborates on what she means by this in the following terms:

...if a person did know [about the existence of a reason] and *if nothing were interfering with her rationality*, she would respond accordingly. (Korsgaard 1986, 14, footnote 9)

Her idea is thus the same as Williams's. In so far as people are rational they are motivated to do what they believe that there is a normative reason for them to do.

Thomas Scanlon, who has misgivings about both Humean and Kantian conceptions of human psychology, makes much the same point.

Rationality involves systematic connections between a person's judgments and his or her subsequent attitudes. ... [A] rational person who judges there to be compelling reason to do A normally forms the intention to do A, and this judgment is sufficient explanation of that intention and of the agent's acting on it (since

this action is part of what such an intention involves). (Scanlon 1998, 33–34)

In other words, according to Scanlon, people who judge themselves to have a normative reason to act in a certain way are rationally required to be motivated to act in that way. But if these theorists are right then it follows that, contrary to radical Humeanism, desires for ends are subject to a correctness condition. For reason requires that agents' desires for ends covary with their judgments about what they have normative reason to do.


This, in turn, tells us something important about fully orthonomous action. If desires for ends are indeed subject to such a norm of reason then it follows that, when agents act in a fully orthonomous way, their desires for ends must be the product of a capacity to have motivating reasons that accord with their judgments about what they have normative reason to do. Agents whose desires for ends do not covary with their judgments about their normative reasons in this way violate a norm of practical reason. Moreover it also follows that when agents act in a fully orthonomous way their judgments about normative reasons must meet their own correctness condition.

Agents' judgments about their normative reasons look, after all, like they can be true or false, and more or less sensitive to the available evidence. I might think that I am required to write a paper about philosophy of action, but I might be mistaken, as I am not in fact required to write such a paper. Or, abstracting away from whether or not I am required to write such a paper, my judgment that I am subject to such a requirement might bear no relation to such evidence as regards what I am and am not required to do as is available to me. To be completely correct, then, agents' desires for ends must not only covary with their judgments about what they have normative reason to do, but their judgments about their normative reasons must in turn be the product of the evidence available to them, evidence that is itself a manifestation of the facts about their normative reasons.

It thus follows that, at least according to common sense, radical Humeanism quite dramatically misrepresents the structure of fully orthonomous action. Common sense tells us that the desires for ends and means-end beliefs that produce fully orthonomous action must be understood in a far more symmetrical fashion (see figure 4).

According to common sense, fully orthonomous action requires not just that agents be in a position to have and act on knowledge of means to ends, but also that they be in a position to have and act on

Figure 4. The common sense conception of fully orthonomous action




knowledge of what there is normative reason to do. In this respect, common sense tells us that fully orthonomous action requires that we live in a doubly obliging world—facts of both kinds must make themselves manifest—and it also requires that we have and exercise ideal belief-forming capacities about means to ends in the light of our evidence, ideal judgment-forming capacities about the normative reasons we have in the light of our evidence, and ideal desire-forming capacities in the light of the judgments we make about the normative reasons that we have. Fully orthonomous action is paradigmatically virtuous action (Pettit and Smith 1993).

Correspondingly, common sense also entails that narrowly orthonomous action is a much more substantial accomplishment than radical Humeanism suggests. For narrowly orthonomous action commonsensically requires that an agent's beliefs about means-ends are the product of such belief-forming capacities as she has—her actual rational capacities—given the evidence available to her, that her judgments about what she has normative reason to do are the product of such judgment-forming capacities she has given the evidence available to her, and that her desires for ends are the product of such capacities as she has to form desires in the light of the judgments she makes about what she has normative reason to do (see figure 5).

The glaring differences between this common sense conception of narrow orthonomy and the radical Humean conception are

Figure 5. The common sense conception of narrowly orthonomous action



striking. Moreover the common sense nature of the common sense conception reveals the truly revisionary nature of radical Humeanism.

3. Realist and irrealist interpretations of common sense

The mere fact that radical Humeanism is a revisionary doctrine is no objection to it. For if there are compelling reasons to revise common sense then we have no choice but to embrace those revisions. But are there such reasons?

The obvious grounds for revising common sense in the direction of radical Humeanism would be that we can make no sense of the key elements in the common sense conception of orthonomous action. To repeat, these key elements are:

- (i) the idea that there are facts about what we have normative reason to do;
- (ii) the idea that these facts manifest themselves in the form of evidence about what we have normative reason to do;
- (iii) the idea that, via the exercise of our rational capacities, we can arrive at judgments about what we have normative reason to do;


and

(iv) the idea that, via the exercise of our rational capacities, we can acquire corresponding desires.

The question we must ask next is whether we can give a plausible interpretation of these four key ideas.

There is, of course, one completely straightforward interpretation of these key ideas. Realists hold that there are indeed facts about what we have normative reason to do, and hence that judgments about what we have normative reason to do are none other than expressions of our beliefs about this domain of facts (see figure 6).

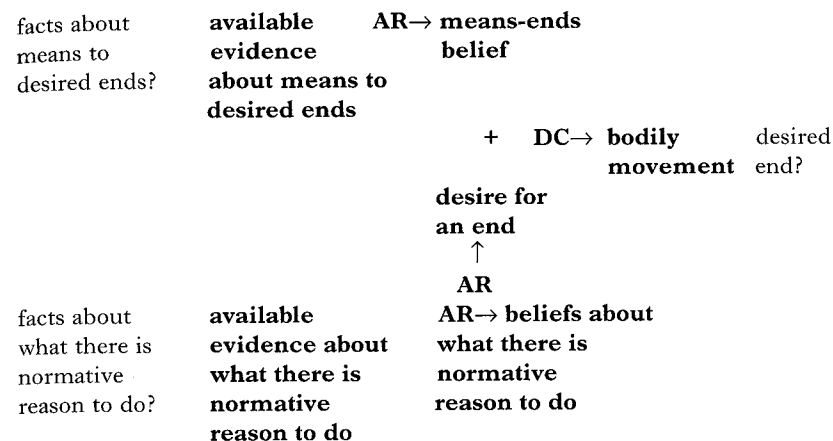
Figure 6. A realist interpretation of the common sense conception of fully orthonomous action



Fully orthonomous action is a matter of an agent's being suitably placed to have knowledge of both means-ends and normative reasons, of her having and exercising idealized capacities for the formation of rational beliefs about both means-ends and normative reasons in the light of the evidence available to her and the formation of desires in the light of her beliefs about her normative reasons, of her having and exercising the capacity to move her body in the way required to realize her desired ends, and of the world obliging her by allowing that desired end to come about in a manner that corresponds to her pattern of practical reasoning.

There is also a corresponding realist picture of the structure of narrowly orthonomous action (see figure 7).

Figure 7. A realist interpretation of the common sense conception of narrowly orthonomous action



As before, narrowly orthonomous action abstracts away from the assumption that the agent is located in a suitably obliging world, and of her having idealized rational capacities. According to the realist, it simply requires that she exercises such belief-forming capacities as she has about her normative reasons, and such capacities as she has to form desires and act in the light of these beliefs. Such capacities may, of course, be very limited. But the crucial point is that, since this is all we could possibly expect of an agent, it follows that this is all that there could be to holding her responsible.

Now it might be thought that this realist alternative is a non-starter. After all, how could we take seriously the idea of there being facts about normative reasons for action, facts belief in which are subject to a requirement of reason to pair up with a desire to act accordingly? Isn't the radical Humean right that, since desires cannot be true or false, it follows that there is no way to make sense of such a rational requirement? And doesn't this entail that desires are not subject to such requirements of reason? There are, however, two quite different ways in which this realist option can be developed.

On the one hand, there is non-reductive realism of the kind that is argued for by Jonathan Dancy (1993, 2000), Derek Parfit (1997), and Thomas Scanlon (1998). In defending realism these theorists take the high road. They offer no analysis of facts about normative reasons (i), and hence no account of either how such facts manifest themselves in evidence (ii) or of how, on the basis of such evidence,