

COPYRIGHT NOTICE:

Kieran Setiya: Reasons without Rationalism

is published by Princeton University Press and copyrighted, © 2007, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@pupress.princeton.edu

INTRODUCTION

THIS IS A BOOK about how one should live. And since I take it for granted that what one should do, all things considered, is what there is most reason to do, it is at the same time a book about practical reason. The view it defends is roughly this: that one should live and act as a person of *good character* would live and act, if she were in one's place; one should imitate the ethically virtuous person.

There are complications here. What am I to do when I have managed to end up in a ditch in which no virtuous person could be found? What about the fact that there are many different ways of being good? For the moment, let us set these issues aside.¹ I will argue that we cannot say what it is to have a reason to act, or understand the nature of practical reason, except in terms of ethical virtue. It follows, as I will try to show, that Aristotle was right: we cannot be fully good without the perfection of practical reason, or have that perfection without being good.²

When I describe my view to non-philosophers, it is sometimes met with blank incomprehension—not, I think, because its terms are specially obscure, but because it is hard to see why one would bother to defend it. It can seem too obvious a truth to count as a philosophical insight. Of course one *should* act as a virtuous person would. “[There] is no one who needs to be told that he ought to be just and brave and temperate. This is self-evident, and calls for no deliberation” (Pieper 1966: 33). But while I am sympathetic to the spirit of this remark, the claim I am defending is *not* self-evident. What I am arguing is not just that there is a sense of “should” connected with ethical virtue, but that this is the “should” of *practical reason*, of what there is most reason for anyone to do. It is a commonplace of modern moral philosophy that there is at least a nominal distinction here—even for those who hope that practical reason and ethical virtue will converge. My claim is that this distinction, the idea that there are standards of practical reason apart from or independent of good character, is a philosophical mirage.

In saying this, I reject the tradition that descends from Hobbes and Hume to economics and decision theory, on which practical reason is conceived as purely instrumental: it is a matter of means-end efficiency, not of

¹ I address the first question in this introduction, and the second in Part Two, section 1.

² This is a loose rendering of Aristotle's *Nicomachean Ethics*, 1144b31–2.

ethical virtue.³ And I reject the Kantian conception of practical reason, on which its standards derive from the nature of agency, as such.⁴ On this view, too, the condition of being properly responsive to reasons can be distinguished from the good condition of one's habituated character.

We can see more clearly what is distinctive of my view, and how it conflicts with these traditions, by examining the question "Why should I be moral?" If *ethics* is concerned with how one should live, or what one should do, all things considered, and *ethical virtues* are virtues of character, *morality* can be thought of as *part* of ethics, and the moral virtues as a *subset* of the ethical virtues. It is not, in the end, very easy to say what is distinctive about morality and the moral virtues, if anything. A first thought is that they are essentially *other-regarding*. Thus justice and benevolence are paradigms of moral virtue—by contrast with non-moral virtues like prudence, moderation, and means-end efficiency.⁵ Since nothing turns on the significance of grouping the moral virtues together, we need not look for a definition. Instead we can rely directly on our examples. To ask, "Why should I be moral?" is, in effect, to ask, "Why should I be benevolent, or just?"

This question can be heard in two ways. It might express doubts about the standing of justice and benevolence *as* virtues of character—as, for instance, in Nietzsche, on Foot's (2001, ch. 7) account of him, or Callicles in Plato's *Gorgias*. Or it might express doubts about the practical justification of morality that *concede* common assumptions about moral virtue. "Why should I be moral?" is meant to be a question of the second kind: it is asked by someone who is willing to agree that a *virtuous* person would be benevolent and just, in the ordinary sense, but who wants to know what reason he has to follow suit.⁶ He accepts that the so-called moral virtues are virtues of character, but wonders why he should not cast them aside.

³ For accounts of this kind, see, especially, Williams (1980; 1989) and Dreier (1997). Williams protests that his view is not instrumentalist, because it has a richer conception of practical thought. But he ties an agent's reasons to his "subjective motivational set," and that is enough to belong to the tradition I have in mind.

⁴ For Kantian accounts of practical reason, see Darwall (1983), Gewirth (1978), Korsgaard (1996, 1997), Nagel (1970), and Velleman (1989; 2000b).

⁵ A similar distinction is made by Slote (1992: xvi and *passim*). I think it is actually unclear whether means-end efficiency is an unconditional virtue. This issue is discussed in Setiya (2005); see also the conclusion of this book.

⁶ As an example of confusion in this area, consider the common mis-reading of Prichard (1912) as being concerned with the relation between practical reason and morality—the second question. In fact, Prichard takes it for granted that obligations provide reasons. His question is about the standing of what we *take* as our obligations *as* genuine obligations. It is thus analogous to the first of the two questions distinguished in the text.

On the view that I defend in this book, his question rests on a mistake. When I say that one should act as a good person acts, I am thinking of good character in general, not the moral virtues in particular. But I treat these virtues—ones like justice and benevolence—in the same way as any others. They are not subordinate to the non-moral virtues of prudence or efficiency, or of “consistency in action.”⁷ If a virtuous person would be moved by certain considerations, it *follows* that they count as reasons to act. So if justice and benevolence are really virtues, they correspond to reasons in their own right: it belongs to good practical thought to give weight to the kinds of considerations to which the just and benevolent person is sensitive. The answer to the question “Why should I be moral?” is not, on this account, supplied by further *reasons* to be moral, which are certified as reasons by a standard other than ethical virtue. It is supplied by the fact that having the moral virtues is a matter of being responsive to considerations that *therefore* count as reasons to act.

This view contrasts with most of those that figure in recent debate, where it is assumed that the standards of practical reason can be understood apart from ethical virtue, and that the question “Why should I be moral?” is about how the life of ethical virtue, and moral virtue in particular—the life of justice and benevolence—is to be justified by these standards.⁸

This is true most obviously of the instrumentalist approach, on which good practical thought is finding and taking the means to one’s ends, where the ends are set by brute desire. No doubt a virtuous person is good at doing these things. But, on the face of it, the converse implication does not hold. One need not have the virtues of character in order to be good at getting what one wants. And if one is selfish but efficient, the virtues of justice and benevolence may seem to get in the way.

As this suggests, the question of reasons to be moral will be pressing also for those who tie *self-interest* to practical reason, insisting that an agent should do only what will benefit her. Why should she keep a promise, or restrain her appetites, unless she stands to gain by doing so? Some have argued—in the spirit of Hobbes’ *Leviathan*—that the moral virtues *can* be justified in terms of self-interest or desire. These arguments are controversial, either because they seem to justify too little, or because they only show the benefits of being just or benevolent in general, not in every particular case.⁹ But these issues are beside the point. The argument of

⁷ For the idiom of “consistency in action,” see O’Neill (1985) and Gibbard (1999).

⁸ See Brink (1992) for an elegant account of this debate; an important precursor is Foot (1972).

⁹ For an argument that falls into the first trap, see Gauthier (1986); for the second, see Foot (1958–9).

this book is directed not only against those for whom the contrast between reason and virtue amounts to actual divergence, but also to many of those who hope to see them coincide. The question is whether the standards of practical reason can be so much as understood apart from ethical virtue, so that it is the task of a more or less elaborate argument to bring them back into line. In my view, the project of *Leviathan*, and projects like it, are misconceived right from the start. They wrongly assume that we can explain what practical reason is in terms of self-interest or the satisfaction of desire—and in isolation from the virtues of character.

A distinction of the same kind is implicit even in Kantian conceptions of practical reason, which aspire to demonstrate the rational authority of the moral law. In doing so, they begin with the nature of agency or practical thought, from which they hope to derive “internal” or “constitutive” standards of success. It follows that, even if the Kantian argument shows that we *should* be benevolent and just, the most it can be is a vindication of the virtues of character in terms of practical reason, *independently conceived*. It is the assumption of independence that I oppose.

It is already clear, in this sketch of my conclusion, that I am engaged in a kind of “virtue ethics.” I am happy to accept that description; but it could be misleading. Virtue ethics is many things to many people, and only some of them are at issue here. It will be helpful, then, to locate my project briefly within the space of ethical theories that appeal to the virtues of character.

I have three things, primarily, in mind: virtue ethics as concerned with moral perception, and hostile to moral principles; virtue ethics as competing with consequentialism and deontology; and virtue ethics as Aristotelian naturalism. This book does not fit squarely in any of these conceptions; its topic is virtue ethics as a theory of practical reason.

There is, nevertheless, some overlap, particularly with the first conception: virtue ethics as (what has come to be called) “particularism.”¹⁰ There are in fact two questions here, not always clearly distinguished. On the one hand, there is the question whether the content of morality, or of practical reason more generally, can be codified in non-ethical terms—for instance, whether we can express, with a finite non-moral description, the conditions of application of every moral concept. One kind of generalist says that we can. One kind of particularist denies it: he claims that the class of things that fall under a moral concept may be “shapeless” at the level of non-moral description. On the other hand, there is a question about the role of ethical principles in the practical thought of the ethically

¹⁰ On particularism in ethics, see, especially, McDowell (1979), Wallace (1991), and the papers in Hooker and Little (2000).

virtuous person: does she decide what to do by applying a set of principles to the situation at hand? These questions are obviously connected: if the content of morality or practical reason cannot be codified in non-ethical terms, there is a kind of principle on which the virtuous person cannot rely, simply because there is no such thing. But the questions are nonetheless distinct. It might be possible to capture the content of ethics in finite terms, without its being true that knowledge of this description figures in the psychology of ethical virtue. Nothing I say in this book will bear in a direct way on the first question, about the existence of finite principles. But I will argue for a sceptical position about the need for ethical principles, of almost any kind, in the practical thought involved in the virtues of character.¹¹

I will have less to say about the other conceptions of virtue ethics. According to one of them, virtue ethics is to be conceived as an alternative to consequentialism and deontology.¹² Consequentialists want to define right action in terms of “the good”—the idea of a good outcome or state of affairs. The right action is that which generates the most good, or a sufficient amount of good. Deontologists define the good in terms of right action, or at least reject the consequentialist definition of the latter. In contrast with both views, virtue theorists (of the relevant kind) hope to explain right action, and the good, in terms of ethical virtue. Each of these theories is characterized by a claim of “*explanatory primacy*” (Watson 1990: 451).

I want to distance myself from this, in two ways. First, to repeat a point I have already made, my topic is not morality in a narrow sense, but ethics in general. I am not interested in what makes an action *morally* right or wrong, in particular, but in what one should do, all things considered. And I leave room for non-moral virtues. Second, although I am arguing for a metaphysical connection between ethical virtue and practical reason, I do not claim that the connection is *asymmetric* in any interesting way. We can say what it is to be a reason for action in terms of ethical virtue, or so I will claim. But that is not to say that the virtues of character have explanatory primacy. The connection between reason and virtue runs in both directions: it is a matter of reciprocity, not priority.

This fact is worth stressing, and I return to it below, and in the conclusion of the book. For now, two further points. The first depends on being careful about the distinction between metaphysics and epistemology. Even if it were true, in some sense, that ethical virtue is more basic than practical reason in the metaphysical order of explanation, it would not follow—and I do not believe—that it is epistemically prior. It would not

¹¹ See Part Two, section 1.

¹² For a classic expression of this approach, see Watson (1990) building on Rawls (1970).

follow that claims about what there is reason to do must always be derived from claims about ethical virtue that are antecedently justified, or that the order of justification cannot go the other way.¹³ If reason and virtue are connected in the way that I propose, assumptions about practical reason might be used to prove conclusions about the virtues of character—though the converse holds as well.

The second point is also connected with issues of metaphysical explanation. On the third conception mentioned above, virtue ethics is identified with Aristotelian naturalism, according to which we can explain what a human virtue is in terms of human nature, human flourishing, or the human function.¹⁴ It is often taken for granted that an account of this kind is necessary. But that assumption is mistaken, in at least two ways. First, we should not forget that the Aristotelian tradition is only one possibility here; there is also the “sentimentalist” virtue ethics of Hutcheson and Hume, with its radically different view of the metaphysics and epistemology of virtue. Second, even those who appeal to Aristotle do not always agree that his naturalism has explanatory ambitions. In “The Role of *Eudaimonia* in Aristotle’s Ethics,” John McDowell argues that “disputes [about what one should do] could evidently be conducted as disputes about what it is the business of a human being to do” (1980: 13). But there is no suggestion that the justification for claims about “the business of a human being is to be found in an independent, ‘value-free’ investigation of human nature.” On the contrary, the reference to “human nature” here is “a sort of rhetorical flourish, added to a conclusion already complete without it” (1980: 19). In effect, McDowell reads Aristotle as a quietist or anti-foundationalist about human virtue.

It does not matter for my purposes whether McDowell is right about this, or exactly what his doctrine amounts to—though, like him, I want to resist the bad idea that the *reasons* for being virtuous turn on an appeal to self-interest or human nature.¹⁵ What matters is that our investigation of virtue and practical reason does not depend on any particular view about the metaphysics of virtue. It is, I think, compatible with the sentimentalism of Hume’s *Treatise* and with the naturalism or quietism—whichever it is—of the *Nicomachean Ethics*. On these metaphysical questions, my argument is basically silent.

¹³ On antecedent justification, see Pryor (2000).

¹⁴ For accounts of this kind, see Thompson (1995)—following Anscombe (1958: 38, 41)—MacIntyre (1999), Hursthouse (1999, Part III), and Foot (2001).

¹⁵ For a reading of Aristotle that may fall into this trap, see Williams (1985, ch. 3). McDowell (1980; 1995b, §§2–3) and Watson (1990) argue, in different ways, that virtue ethics should not be conceived on an egoistic model. See also Hurka (2001, ch. 8).

In the rest of this introduction, I attempt to do two things. In section 1, I clarify and make precise the claim about reason and virtue that the rest of the book defends. In doing so, I solve a problem for virtue ethics that has seemed decisive to some: an argument by Bernard Williams that it makes no sense for the less-than-virtuous to imitate the fully virtuous person.¹⁶ If this is right, it is not clear how the virtue theorist can state the connection between ethical virtue and reasons for us to act. Responding to Williams, I explain how this connection can be made. In section 2, I sketch the argument of the book as a whole. It begins with a theory of action developed in Part One, which figures as a premise in the argument about reason and virtue that occupies Part Two.

1. “SQUEEZING THE GOOD INTO THE RIGHT THROUGH THE TUBES OF IMPERFECTION”

According to Aristotle, we engage in deliberation “where the outcome is unclear and the right way to act is undefined” (*NE* 1112b9–11).¹⁷ Here “deliberation” seems to mean something *active* or *intentional*, the kind of deliberative thought in which we deliberately engage, surveying our circumstance and what we can do to change or affect it. As Aristotle says, we engage in active deliberation mostly when our decision is difficult—though it need not be important. Much of our time is spent acting intentionally but without deliberating, in this sense: we get out of bed, make breakfast, and go to work, often without so much as wondering why.

Aristotle also said that “what is decided is what has been previously deliberated” (*NE* 1112a15–16), and he apparently means that every decision (*prohairesis*) is a product of deliberation. It is obvious at once that there is a problem here. If deliberation is something we do intentionally, and if “decision” has its ordinary meaning, then it seems wrong to say that every decision depends on prior deliberation. Indeed, it *must* be wrong, since it would initiate a vicious regress in which even the decision to deliberate would have to be deliberated in turn.

What we need here is a concept of deliberation that applies even when the thinking behind an action is not itself intentional. In fact, it will be useful to cast our net more widely than that. I want to count as “practical thought” or “practical reasoning”—terms I use interchangeably—not only deliberation as an intentional action, but the motivation of action done for reasons (though perhaps without deliberation, in the strict

¹⁶ The argument appears in Williams (1995a: 189–90).

¹⁷ I use Irwin’s (1999) translation of the *Nicomachean Ethics*, with occasional transliteration of the Greek. The title of this section is taken from Williams (1995a: 190).

sense), the balancing of reasons (even when it is not conscious), and the forming and revising of intentions and desires. There is no perfectly natural phrase for the topic I have in mind. “Practical thought” may suggest something wholly cognitive, a matter of evaluative belief, or thoughts about what there is good reason to do. We should not assume that practical thought in the broad sense must always include such things, and it certainly is not exhausted by them. “Practical reasoning” may suggest some kind of *calculation*, perhaps about the means to one’s ends.¹⁸ But the practical thought involved in acting for a reason, for instance, need not involve calculation in any ordinary sense, and we should not assume at the outset that it must depend on our desires. Finally, both “practical thought” and “practical reasoning” may suggest a restriction to mental goings-on, things that terminate with desires or intentions to act, not with action itself.¹⁹ I want to count *acting* for a reason as a central instance of practical thought. After all, it is relevant to one’s degree of responsiveness to reason whether one merely forms the intention to act, about which one is then weak-willed, or whether one actually tries to do something about it.²⁰

The principal topic of this book is practical thought, understood in this broad and somewhat artificial sense. The virtue theory of practical reason is a claim about what constitutes *good* practical thought. But since virtues are traits of character, it is more readily framed in terms of our *dispositions* to engage in practical thought, or what I will simply call “dispositions of practical thought.” A disposition of practical thought is good, as such, just in case it is a disposition to engage in good practical thought. According to the virtue theory, the property of being good as a disposition of practical thought can also be explained in terms of virtue.

The Virtue Theory: Being good as a disposition of practical thought is being a disposition of practical thought that is good as a trait of character.

There are several things to clarify here. First, this is meant to be a claim of *identity*, on which the property of being good as a disposition of practical thought just *is* the property of being good as a trait of character, applied to dispositions of practical thought. A consequence is that one’s dispositions of practical thought are good, as such, just in case (and just to the extent that) they are good as traits of character. Second, we need to say

¹⁸ For the suggestion about calculation, see Anscombe (1963: 65), following Aristotle (*NE* 1139a12–15).

¹⁹ For this restriction, see Broome (2001: 175).

²⁰ On weakness of will, see, especially, Holton (1999). In stipulating that acting for a reason is a kind of practical thought, I take no stand on the substantive issues, if any, raised by Aristotle’s doctrine that action is the conclusion of practical syllogism.

more about the dispositions that constitute good character: what is the psychology of ethical virtue? That will have to wait until Part Two, in “Character and Practical Thought.” (More generally, that section examines the detailed application of the virtue theory.) Third, we want to know what the virtue theorist says, not only about our dispositions of practical thought, but about our reasons to act. It is sometimes said that virtue ethics will face a decisive problem here; in the rest of this section, I argue that it does not.

I began this book with the fundamental question of ethics: “How should one live?” And I took this to be a question about *reasons*: “What is there most reason to do?” “Reasons,” here, are considerations that count in favour. Reasons for doing something are considerations that count in favour of doing it; reasons for caring about someone are considerations that count in favour of caring. So, for instance, the fact that it is a beautiful day is a reason for me to quit work early; the fact that this is my child is a reason for me to care about him; and so on. In claims like these, reasons are facts that support or tend to justify an attitude or a course of action, for a particular person in a particular context. There is a risk of ambiguity here, since we talk about people’s reasons for doing things without assuming that their reasons are good. There is a distinction between the use of “reason” in *explaining* someone’s action, or attitude, and its use as a term for *justification*—although they are surely connected. We will examine the explanation of action by reasons in Part One. For the moment, we are interested in practical justification: we are interested in *ethics*, and the “should” that signifies what there is most reason to do.

The problem is that, while it is possible to state the virtue theory as a claim about dispositions of practical thought, as above, it is not clear how to relate the virtues of character, or the ethically virtuous person, to reasons for acting or the ethical sense of “should.” Bernard Williams (1995a: 189–90) has argued, influentially, that this puzzle amounts to a fatal flaw: one’s reasons are sensitive to one’s defects of character in a way that prevents them from being understood in virtue-theoretic terms. His target is the tempting formulation that I employed, with caveats, at the beginning of this book: that one should live and act as a person of good character would live and act in one’s place. To reason well in practical matters is to reason in ways that are characteristic of those who have the virtues of character, “someone like Aristotle’s *phronimos*, [. . .] someone who has been properly brought up” (Williams 1995a: 189). Taking the virtuous person as a standard, we are led to a proposal that looks like this:

(V) The fact that *p* is a reason for A to ϕ just in case A would be moved to ϕ by the belief that *p*, at least to some extent, if she had the virtues of character.

However tempting it may be, this formulation is deeply flawed. For one thing, it assumes that the virtues of character are mutually compatible, and that may not be so. But there is a further and more fundamental problem. As Williams has shown, it is a mistake to explain the idea of a reason for acting in *imitative* terms.

[According to (V)] what A has reason to do in certain circumstances is what the *phronimos* would have reason to do in those circumstances. But, in considering what he has reason to do, one thing that A should take into account, if he is grown up and has some sense, are the ways in which he relevantly fails to be a *phronimos*. Aristotle's *phronimos* (to stay with that model) was, for instance, supposed to display temperance, a moderate equilibrium of the passions which did not even require the emergency semi-virtue of self-control. But, if I know that I fall short of temperance and am unreliable with respect even to some kinds of self-control, I shall have good reason not to do some things that a temperate person could properly and safely do. (Williams 1995a: 190)

To take an example from Watson (1975: 210), which is used by Michael Smith (1995: 110–1) in a similar context, suppose that I have just lost a very hard game of squash to an infuriating opponent. In this circumstance, the virtuous person would calmly shake his opponent's hand. But if I were to attempt that, I would fly into a rage and beat him with my racquet. The fact that I have a short fuse is a reason for me to take a cold shower, but it is not a reason by which I would need to be motivated if I had the virtues of character. In a circumstance like this, (V) gets it wrong: it cannot accommodate the reasons I have that would not be reasons for the *phronimos*, in a similar place.

The example depends on my having a defect of character, and on this defect's being practically significant. So perhaps our mistake was not to build this defect into the specification of my circumstance, before relating it to (V). But this won't help:

It will be no good trying to accommodate this difficulty, of squeezing the good into the right through the tubes of imperfection, by putting all A's limitations in the account of the circumstances. If the circumstances are defined partly in terms of the agent's ethical imperfection, then the *phronimos* cannot be in *those* circumstances and [(V)] cannot apply at all. (Williams 1995a: 190)

Since defects of character can be relevant to my circumstance in ways that bear on what I have reason to do, there is simply no way to explain those reasons in terms of the practical reasoning of the virtuous person who finds himself in exactly the circumstance that I am in.

Williams' argument is, I think, a decisive objection to (V). What is striking about it, however, is that it does not depend essentially on appeal to the virtues of character. The point of the argument is really quite general:

it is about imitation, not ethical virtue. If I am bad at practical reasoning, whatever that amounts to, I am bound to have reasons that do not apply to those whose dispositions of practical thought are better than mine. So the problem for (V) is just as much a problem for (R):

(R) The fact that p is a reason for A to ϕ just in case A would be moved to ϕ by the belief that p , at least to some extent, if her dispositions of practical thought were perfectly good.

Suppose that I am prone to a failure of practical reason: I tend to give too much weight to sunk costs. That will affect what there is reason for me to do, providing me with reasons that I would not have if my dispositions of practical thought were flawlessly good. For instance, it may be my heart’s desire to run a successful restaurant, but my obsession with sunk costs will get in the way. If I open a restaurant, and try to make a go of it, I will be unable to back out: unable to quit when things are hopeless, even though it would be self-destructive to go on. The fact that opening a restaurant is such a risky prospect, that most restaurants fail, and that pouring in more money doesn’t help—these may be reasons for *me* not to do it, even though they would not apply to someone who did not share my defect of practical thought. The problem here is exactly parallel to the one that Williams describes.

What this shows is that we have misunderstood the relationship between practical reasons and practical reasoning, in general. The problem of imitation does not depend on an alleged connection between practical reason and virtue of character; it is not a problem about virtue at all. It is about the need for care in stating the formal connection between reasons for action and good dispositions of practical thought, whatever their content may be. Our mistake was to connect particular reasons for action with the case in which someone has the whole array of good dispositions. In doing so, we necessarily obscure the way in which reasons can depend on defects of practical thought. Still, it is clear that there must be some connection here. It would make no sense to conclude that the standards of practical thought are conceptually independent of the facts about reasons to act. The only question is how to state their relation to one another. And this is a question for everyone—not just for those who are sympathetic to the virtue theory of practical reason.

Roughly speaking, the connection seems to be this: a reason is a premise for an episode of good practical thought whose other conditions are already in place. If the fact that p is a reason for you to ϕ , then it is good practical thought to be moved to ϕ by a certain array of psychological states, and you have that array—except, perhaps, for the belief that p . This belief would supply the final material for a good disposition of practical thought. We can make this idea more precise as follows:

Reasons: The fact that p is a reason for A to ϕ just in case A has a collection of psychological states, C , such that the disposition to be moved to ϕ by C -and-the-belief-that- p is a good disposition of practical thought, and C contains no false beliefs.²¹

The last clause of this formula—“no false beliefs”—is required to deal with a final difficulty. In an earlier example, which I have slightly modified, Williams imagines a thirsty person, presented with what seems to be a glass of cool, refreshing water (1980: 102). In fact, the glass contains odourless petrol. If I am in this situation, is the fact that I am thirsty a reason for me to drink the contents of the glass? As Williams says, the answer would seem to be “no.” If the glass contains petrol, the fact that I am thirsty is no reason to drink from it, at all; there is no good reason to drink what is in the glass. The inclination to say otherwise turns on the fact that I have a collection of psychological states—including the belief that the glass contains water—such that the disposition to be moved to drink by them, together with the belief that I am thirsty, is a good disposition of practical thought. What the example shows is that good practical thought corresponds to reasons only when it involves no false beliefs.²² A consequence of this refined claim is that our reasons

²¹ It is worth comparing this claim briefly with a different solution to something like the generalized problem of imitation: the “advice model” proposed by Michael Smith. His central claim is that “[what] we have reason to do in the circumstances in which we find ourselves is fixed by the advice our fully rational selves would give us about what to do in these circumstances” (Smith 1995: 112). Advice is cashed out in terms of *desire*; so reasons correspond, not to the motivations we would *have* if we were “fully rational,” but to the motivations we would *want* ourselves to have, in our actual position. This model assumes—I think, quite dubiously—that those who are “fully rational” must be good at giving advice. That may be true, but it is far from trivial: it is like the assumption of those who treat Prince Myshkin as a mentor, in Dostoevsky’s *The Idiot* (1869). (For a similar objection, see Millgram 1996: 218, n. 14.) In any case, there is a fundamental difference between the advice model and the claim defended in the text. My account is concerned with the connection between reasons for action and good practical thought in the circumstance in which they apply. The advice model, by contrast, relates the reasons of an agent to the practical thought of his *advisor*. It does not say anything, for its own part, about the standards of practical thought that apply to the circumstance the agent is in. Thus, even if the advice model is true—which I doubt—it does not explain the connection we have been looking for, the connection I said there simply must be, between practical reasons and the quality of the practical thought that takes those reasons into account.

²² A natural question to ask at this point is: why not correct for *ignorance* as well as false belief? But that “correction” would be a mistake. Reasons for action *may* correspond to practical thought that depends on ignorance of fact. So, for instance, there may be a reason for the gambler to bet on the horse with the best odds, even though it will lose the race—since he does not *know* that the horse will lose. His reason corresponds to good practical thought that depends essentially on his being in the dark about that. (For this example, attributed to Frank Jackson, see Dancy 2000: 65–6.)

supervene on our psychological states, together with the truth of our beliefs. The rest of the world is irrelevant. If agents are psychological duplicates, and their beliefs are similarly true or false, then their reasons must be the same.²³

The doctrine of *Reasons* is immune to the problem of imitation because it does not connect the assessment of practical thought with the case in which *all* of one’s dispositions are good. If the risk of failure is a reason not to open a restaurant, since I give too much weight to sunk costs, the disposition to be deterred by that consideration in my imperfect circumstance is a good disposition of practical thought. (More carefully, I ought to be deterred to some extent; not all reasons are decisive ones.) This is true even though I would not be deterred if I had the whole array of good dispositions.

As with the virtue theory, *Reasons* should be thought of as basically symmetric. It says that we can move from claims about reasons to claims about good dispositions of practical thought, and back again. It does not imply that claims of the latter kind have metaphysical or epistemic priority.²⁴ And it can therefore be accepted by non-reductionists about reasons, and by those for whom good practical thought involves knowledge of what there is reason to do.

We can state a similar claim about the *relative weight* of reasons to act:

Reasons correspond to collections of psychological states that fuel good practical thought. One reason is *stronger* than another just in case it is a good disposition of practical thought to be *more strongly moved* by the collection of states that corresponds to it, than by the collection that corresponds to the other.

(Here, “good” means “better than the alternative,” not “fine” or “good enough.” We would otherwise permit the impossible case in which two reasons are stronger than each other, because the corresponding and conflicting dispositions are both “good enough.”) With this principle in hand, we can finally explain the relationship between what one should do, all things considered, and the assessment of practical thought. What one should do is determined by the balance of reasons, which is fixed in turn by the relative weight of all the reasons there are. The facts about reasons correspond, in the ways I have been describing, to facts about good dispositions of practical thought.

²³ This claim is stronger than Scanlon’s (1998: 73–4) “universality of reason judgments,” but it is in the same vein.

²⁴ There is the following asymmetry: all reasons correspond to good dispositions of practical thought, but good dispositions need not correspond to reasons, since they may not take the relevant form—that of a disposition to be moved to *act* by, among other things, a *belief*.

None of this depends on the virtue theory of practical reason. Although it began as a question about reason and virtue, we have seen that the problem of imitation is entirely general: it applies to (R) as much as (V). In solving it, we have to articulate more carefully the formal connection between reasons for action and good dispositions of practical thought. That there *is* such a connection is something everyone should accept. What the virtue theory adds is a distinctive conception of what is good in practical thought: not efficiency, or practical consistency, but virtue of character. Being good as a disposition of practical thought is being a disposition of practical thought that is good as a trait of character. Since the facts about reasons correspond to facts about good dispositions of practical thought, the virtue theory can serve as a theory of reasons, and of what one *should* do, all things considered. As I remarked before, this is not a claim of metaphysical or epistemic priority. I am not saying that facts about good dispositions are more basic than facts about reasons, in any significant way. If the virtue theory is right, we can move seamlessly from practical reason to ethical virtue, and the reverse.

2. THE RELEVANCE OF ACTION THEORY

This book presents a single sustained argument for the virtue theory of practical reason. It is inspired by an ominous pronouncement in Elizabeth Anscombe's "Modern Moral Philosophy":

[It] is not profitable for us at present to do moral philosophy; that should be laid aside at any rate until we have an adequate philosophy of psychology, in which we are conspicuously lacking. (Anscombe 1958: 26)

Times have changed—not, I will argue, because it is any more profitable to engage in philosophical ethics apart from the philosophy of mind, but because we have Anscombe's work on which to draw. In the year before she wrote the words I have quoted here, Anscombe published the first edition of *Intention*, that difficult and remarkable book. Although I will criticize her work, below, its influence on my approach is deep. Like Anscombe, I think we cannot understand the character of practical reason without an account of what it is to act. More carefully, we need an account of what it is to act *intentionally*, and what it is to act for a reason. These are the questions to be explored in Part One.

In turning from ethics to action theory, I follow not only Anscombe, but a family of views in moral philosophy that may be gathered together under the heading of "ethical rationalism." According to the rationalist, in this sense of the term, the standards of practical reason can be derived, at least in outline, from the nature of agency or practical thought. The

philosophy of action is thus the foundation of ethics. The most prominent rationalist views belong to the tradition that descends from Kant, and my use of the term derives from and extends its application there. We find versions of Kantian rationalism in recent work by David Velleman (*Practical Reflection, The Possibility of Practical Reason*) and Christine Korsgaard (*The Sources of Normativity*).²⁵ For instance, Velleman claims that we can say what it is to act intentionally in terms of a “constitutive aim,” something to which everything we do is tacitly directed. Right action is what best achieves this constitutive aim. For Korsgaard, the standards of practical reason follow from the fact that rational agency depends on reflective endorsement. Right action must be compatible with *unconditional* endorsement and thus, so the argument runs, with respect for humanity as an end in itself. Understood in a schematic way, however, the project of ethical rationalism is not compelled to take a Kantian form. For instance, it appears in the version of *instrumentalism* on which desire-satisfaction is the goal of practical reason because desire has a distinctive place in practical thought.²⁶ On this view, too, the standards of practical reason derive from the nature of agency, as such.

Although I share the rationalist concern with agency, I do not share its optimism. One of the central arguments of this book is that the standards of practical reason *cannot* be derived from the nature of agency or practical thought, not even in outline, or in general terms. As I will try to show, it follows from this failure that the virtue theory is true. The reasoning here is hard to summarize; it depends on a premise about the metaphysics of *being good* that I defend in the second section of Part Two, “An Argument for the Virtue Theory.” But the upshot is this: those who reject the virtue theory of practical reason *must* accept some form of ethical rationalism. They must derive the standards of practical reason, at least in outline, from the nature of agency or practical thought. (This is so even if, like the instrumentalist, they do not take a traditionally Kantian line.) If ethical rationalism is false, as I argue it is, then the virtue theory must be true. The standards of practical reason are standards of good character, applied to practical thought.

This sketch leaves much to be filled in. Why should the metaphysics of *being good* force us to choose between rationalism and the virtue theory? How do the standards of good character apply to practical thought? What is wrong with ethical rationalism? These questions will be answered in Part Two. In the rest of this introduction, I want to say more about Part One, and the role it plays in the arguments to come.

²⁵ Velleman (1989; 2000b); Korsgaard (1996); see also Nagel (1970) and Railton (1997).

²⁶ See Williams (1980); Dreier (1997).

The topic of Part One is the claim that we act intentionally *sub specie boni*: “under the guise of the good.” On this “normative” conception of agency, to do something intentionally is to see some good in doing it; and to act for a reason is to act on a belief one takes to justify what one does, at least to some extent. As I will argue later on, this picture of agency is presupposed by the Kantian rationalist. But it is more widely shared. The same conception is at work in some Platonic and Aristotelian views, and it is often unreflectively assumed. I will try to show that it is false. In fact, one way to understand the argument of the book as a whole is to think of it as working out the consequences for practical reason of rejecting the claim that we act under the guise of the good.

We need to make some distinctions here. The normative conception of agency may derive from a doctrine about desire: “[what] is required for our concept of ‘wanting’ is that a man should see what he wants under the aspect of some good. [. . .] The conceptual connexion between ‘wanting’ [. . .] and ‘good’ can be compared to the conceptual connexion between ‘judgement’ and ‘truth’” (Anscombe 1963: 75–6).²⁷ If an agent must *want* to do what he intentionally does, it will follow that he must see what he is doing “under the aspect of some good.” (A slightly different view is possible here, on which it is one’s *intention* to act that involves, or is identified with, an evaluative judgement.)²⁸ Although I will mention these claims, below, what interests me most of all is an alleged connection between good or normative reasons—“reasons” in the sense of section 1—and the reasons for which we act. Although it is obvious that we can act for bad reasons, it is commonly held that we must at least *believe* that the reasons for which we act are good, that they are considerations that count in favour of what we are doing: “to say that a factor is the agent’s reason is to claim that the agent decided that it was a sufficient reason or sufficient grounds for acting in the way he did” (Milligan 1974: 183); an agent who acts for reasons “is *guided* by reason, and in particular, guided by what reason presents as necessary” (Korsgaard 1997: 221). (The quoted claims are excessively strong. A more modest view is that one must regard one’s reason as doing *something* to justify one’s action, even if it is not sufficient.)²⁹ More generally, deciding what to do is often assumed to be a matter of deciding what

²⁷ For similar claims, see Stampe (1987) and Tenenbaum (2003). The evaluative conception of desire is criticized by Stocker (1979) and Velleman (1992), among others.

²⁸ See Davidson (1978: 98–99); this picture of intention lies behind the problem of *akrasia*, as it appears in Davidson (1970a).

²⁹ For the claim that acting for a reason is acting for what one takes to be a *good* reason, see also Darwall (1983: 205), Bond (1983: 30–1), Walker (1989: 670), Broome (1997), Raz (1997: 8, 1999), Wallace (1999), Velleman (1992a: 101, 1992b: 140–2, 2000b: 9), and many others. A similar view is held by Anscombe (1963: 21–5).

it would be good, or even best, to do.³⁰ And descriptions of practical thought quite commonly imply that, in weighing reasons, one “come[s] to see that [one] has reason to do something” (Williams 1980: 104), as though responding to reasons were always a matter of making judgements that deploy the concept of justification. In coming to act upon the reasons we have weighed, “[the] basic case [is] that in which $A \phi$'s [. . .] because he believes of some determinate consideration that it constitutes a reason for him to ϕ ” (Williams 1980: 107).

In my view, all of these claims are false, and false not merely in marginal cases, where an agent is being irrational or suffers from *akrasia*, but in quite ordinary ones, as well. There is a sense in which, in acting for a reason, one must see the consideration on which one acts *as* one's reason for acting, but one need not see it as a *good* reason for acting, a reason in the normative or justifying sense. To clarify: it is not just that one need not see the consideration as *sufficient* reason; one need not see it as doing anything at all to justify what one does. This claim about agency will be the primary conclusion of Part One.

It supports a further claim about practical thought, in general. As I will argue in Part Two, if we do not act under the guise of the good, there is no reason to suppose that practical thought depends on evaluative judgement. Acting as one should may be the formal object of practical reason, in the sense explained above, in section 1, but it is not the material object of practical thought. We can engage in practical reasoning, and act on it, without asking whether the reasons for which we are acting are any good. We may decide to live like Countess Gemini, in *The Portrait of a Lady*, who cheerfully insists, “I don't care anything about reasons, but I know what I like” (James 1881: 307). But even this gives the wrong impression, that those who do not *reflect* about reasons are wanton or reckless in some way. That need not be so: even the virtuous person may not be articulate; she may not be able to say, and may not even have considered, why what she is doing is right.³¹

The arguments described so far belong to action theory, or moral psychology. What do they have to do with ethics and practical reason? In effect, that is the topic of Part Two. Its principal *negative* claim is that we cannot derive the content of practical reason, not even in outline, from the nature of agency or practical thought—and for reasons that depend upon the arguments of Part One. The project of ethical rationalism turns out quite often to rely upon the guise of the good. Thus, advocates of

³⁰ Watson: “Practical deliberation, as I think of it, is *reasoning about what is best (or satisfactory) to do with a view to making up one's mind about what to do*” (2003: 175).

³¹ In thinking about this fact, I have been influenced by Iris Murdoch; see, especially, Murdoch (1970: 2, 51–2). I say more about this in Part Two, section 1.

Kantian rationalism have typically assumed that, in acting for reasons, we aim to *justify* what we do. That is their mistake. The same point holds for the “recognitional” conception of practical reason, on which practical thought is thought about what there is reason to do, and for what might be called “Aristotelian rationalism.” I explain and criticize these views in Part Two, section 3.

The arguments of Part One are relevant even to forms of ethical rationalism that do not depend upon the guise of the good. (Think, for instance, of the instrumentalist view that practical thought is always directed by desire, so that its aim is desire-satisfaction; or the attempt to derive the standards of practical reason from the fact that, in acting for reasons, we know what we are doing and why.)³² By defending a positive theory of action, I hope to show, in general, that there is no room for the rationalist project to succeed. If the arguments of Part One are right, there is *nothing* in the nature of action, or of practical thought, from which the standards of practical reason could derive. The only alternative—so I claim in Part Two—is that the virtue theory is true.

This is still schematic, and the details will have to wait. But it is already enough to provoke a response. To many, the strategy of this book will seem ironic. It is a defence of virtue ethics that depends on the claim that we do not act under the guise of the good. But the guise of the good is a doctrine that descends from, and is an intimate part of, the dominant tradition of virtue ethics, the one that begins with Socrates, and flows through Plato, Aristotle, and Aquinas, to Anscombe and to us. This is not an accident. While the influence of Aristotle should be apparent here, it is overshadowed by the influence of Hume.

This may seem surprising. While Hume’s interest in virtue is obvious, he has usually been read as an instrumentalist about practical reason: a passion may be “call’d unreasonable” just in case “[it is] founded on a false supposition, or when it chooses means insufficient for the design’d end” (*Treatise* 2.3.3.7).³³ “Reason is, and ought only to be, the slave of passions” (2.3.3.4). This take on Hume is so widespread that it is often made true by definition: in much of the literature, “Humean” just *means* instrumentalist. More recently, some critics have urged that Hume is a *sceptic* about practical reason: he rejects its would-be standards, including

³² It is natural to read Williams as attempting to derive a broadly instrumentalist conception of practical reason from the role of desire in practical thought (1980). See also Dreier (1997). Velleman derives a theory of reasons for action from the cognitive content of intending to act (1989, 2000b). I criticize these views towards the end of Part Two.

³³ Hume (1739–40). I cite the *Treatise* in the main text by giving book, part, section and paragraph numbers.

the standard of instrumental reason, outright.³⁴ After all, he argues that it is “*impossible* [. . .] that [a] passion can be oppos’d to, or be contradictory to truth and reason” (2.3.3.5, my emphasis). This is how he comes to write one of the most arresting passages in his work:

’Tis not contrary to reason to prefer the destruction of the world to the scratching of my finger. ’Tis not contrary to reason for me to choose my total ruin, to prevent the least uneasiness of an *Indian* or person wholly unknown to me. ’Tis as little contrary to reason to prefer even my own acknowledg’d lesser good to my greater, and have a more ardent affection for the former than the latter. (2.3.3.6)

Hume is not, however, an instrumentalist or a sceptic. This is not the place in which to argue this at length.³⁵ The views that occupy this book are not meant as an interpretation of Hume, and do not correspond to his, except in a general way. What I take from him is an idea that lies behind the treatment of reason and passion in the *Treatise*: we know what we mean by “theoretical reason” or “reason, in a strict and philosophical sense” (3.1.1.12); but we have no clear conception of a “practical reason” that corresponds to this. We do not know what practical reason *is*, if it is meant to provide a standard for thought and action, apart from the standards of virtue and vice.

This means that there is a sense in which Hume *is* a sceptic about practical reason, as the passage above suggests. But the sense is limited: he is a sceptic about practical reason, I think, only on a *rationalist* account of what it must be. For Hume, ethical rationalism is the attempt to model practical on theoretical reason, and it comes in the end to a confusion of the two. When this confusion is cleared away, Hume finds it misleading, rather than helpful, to apply the term “reason” in ethics at all; he wants to limit its use to “reason, in a strict and philosophical sense.” And so his language is that of a sceptic: “reason alone can never be a motive to any action of the will” (2.3.3.1); nor can our passions conflict with it. Ethical virtue is left as the only intelligible standard for action and practical thought.

My arguments are not the same as Hume’s. There is at most a loose correspondence between the rationalism of the British moralists and the kind of rationalism that figures in this book—a kind that does not confuse practical with theoretical reason. Nor do I sympathize with Hume’s refusal to speak of “reason” in ethics if, as he concludes, a rationalist account is false. To believe in practical reason is simply to believe that there are things that we should do, and reasons for doing them—and however

³⁴ See Hampton (1995), Millgram (1995), and Korsgaard (1997).

³⁵ I have done so in Setiya (2004c).

he would put the point, Hume surely agrees with that. His scruples about “reason, in a strict and philosophical sense” tend to obscure the best insight of his argument, that there is nothing for reason to *be*, in ethics, if it is not a matter of good character. That is my view too. Like Hume, on this admittedly tendentious reading, I argue that we must interpret the ethical “should” in terms of ethical virtue, because no other interpretation makes sense.