

PRACTICAL KNOWLEDGE

Selected Essays

Kieran Setiya

OXFORD
UNIVERSITY PRESS

INTRODUCTION

This book has two themes: the nature of intentional action and the foundations of ethics. What is it to act for reasons and so to act intentionally? And why are certain facts reasons to act in one way or another, considerations that count in favour of doing so? Such reasons fix what we ought to do.

Each question is of interest in itself and some of the essays deal with one to the exclusion of the other. But the themes are closely related. In particular, there is an approach to ethics I have called ‘ethical rationalism’, which aims to derive the normative facts—what there is reason for us to do—from the nature of agency or the will.¹ According to the rationalist or ‘constitutivist’, the standards of practical reason are explained by what it is to act intentionally, or to have the capacity to do so. In one way or another, action theory is the basis of ethics.

My relationship to this approach is complicated and it plays a special role in the essays to come. Unlike some, I think it is possible to construct a compelling argument from premises in the philosophy of action to ethical conclusions: from the metaphysics of agency to the norms of practical reason.² This strategy has much wider application than is often assumed. While the rationalist approach has been associated with Kant, who aimed to derive the moral law from the idea that we act ‘under the idea of freedom’, a less ambitious rationalist might derive the norm of means–end efficiency from the role of desire in motivation,

1. This terminology is introduced in Setiya 2007.

2. For accounts of this argument, see ‘Murdoch on the Sovereignty of Good’ (Setiya 2013a; available online at <http://quod.lib.umich.edu/p/phimp/3521354.0013.009/1>), ‘Intention, Plans, and Ethical Rationalism’ (this volume: Ch. 10) and ‘Akrasia and the Constitution of Agency’ (this volume: Ch. 11).

or the irrationality of *akrasia* from the idea that we act ‘under the guise of the good’.³

Despite a vivid sense of the power of ethical rationalism, and of its metaphysical and epistemic virtues, I do not in the end accept it. Although the action theory assumed by the ethical rationalist need not be extravagant—far from it—I think the nature of intentional action, and of acting for reasons, is more minimal or impoverished than the argument requires. One purpose of engaging in the details of action theory, as in Part I of this book, is to support this claim. It is distinctive of my approach that I do so while defending a conception of agency that is in certain ways demanding, a conception inspired by remarks on practical knowledge in Elizabeth Anscombe’s pioneering book, *Intention*. By ‘practical knowledge’ Anscombe means our distinctive knowledge of what we are doing when we are doing something intentionally, and of why we are doing it. One of my central claims is that we cannot explain such knowledge, which many find puzzling, without appeal to practical knowledge in a second sense: knowledge how to do what one intentionally does. ‘Practical knowledge’ can be used in a third way, for knowledge of practical reason, knowledge that is ethical in the broadest sense of the term. It is in this sense that Part II is concerned with ethics and, directly or indirectly, with knowledge of what to do. The title of the book thus applies, in one way or another, to everything contained in it.

In the rest of this introduction, I sketch in more detail how I think about the project of action theory, how my conception of practical knowledge has evolved over the last ten years, and how reflection on agency has implications for ethics. It may be useful to state in advance, without elaboration, some theses I defend. In Part I:

The idea of practical knowledge—knowledge of what one is doing or what one is going to do that does not rest on sufficient prior evidence—is central to our understanding of intentional action. The capacity to act for reasons is the capacity for practical knowledge.

Such knowledge rests on, and is partly explained by, practical knowledge in the second sense, of knowing how.

It does not rest on practical knowledge in the third sense: knowledge of ethical facts. More generally, in acting for reasons, we do not act ‘under the guise of the good’: we need not represent our action as a

3. The Kantian strategy has been pursued by Christine Korsgaard (1996, 2009); alternatives are considered in Setiya, ‘Intentions, Plans’ and ‘Akrasia’, and in *Reasons without Rationalism* (Setiya 2007: Part Two).

good thing to do, or the grounds on which we are acting as normative reasons that support it.

Practical knowledge is knowledge in intention, where intention involves, but is not reducible to, belief or partial belief.

And in Part II:

Some of what we regard as practical reason is the application of epistemology to beliefs that figure in our intentions. In particular, this is true of the instrumental principle, that one must intend the necessary means to one's ends.

A normative reason is a premise of sound practical thinking. In this sense, reasons are considerations that move us insofar as we are practically rational.

There is a valid argument from the metaphysics of agency to the norms of practical reason, of the sort invoked by ethical rationalists.

But ethical rationalism is false: the nature of agency is too thin to provide its premises.

The standards of practical reason are standards of ethical virtue, applied to practical thought.

In relation to the last two theses, the defence offered here is partial: further arguments appear in *Reasons without Rationalism*, to which this collection is at once a sequel and a preface. The essays that follow are independent of that book, but they deal with related topics. My hope is that reading them will give a clearer sense of the difficulty, and the urgency, of its project: to make sense of rational agency and reasons to act outside the context of ethical rationalism. At the same time, I hope they make progress with some of the most intriguing puzzles in the philosophy of action, quite apart from their connection with ethics. It is to those issues that I now turn.

1. What Is Action Theory?

Action theory is concerned in the first instance with what it is to act for reasons. It aims to understand the kind of explanation of what someone is doing that cites the reasons for which she is doing it. Explanations of this sort are

often teleological: 'A is buying fish in order to cook dinner'. But they also take non-teleological forms, as when we state the fact, or putative fact, that is the reason for which someone acts: 'A is returning the book on the ground that he promised to; that is among his reasons for doing it'. When an explanation of either kind is true, it follows that A is acting intentionally. The converse implication is less clear: Anscombe disputes it; Donald Davidson responds.⁴

More important for our purposes is the well-marked ambiguity of 'reason', a term that appears both in statements of the reason for which someone acts, like those above, and in statements of the reasons there are for acting in one way or another, considerations that count pro and con: 'The fact that his friend is in need is a reason for A to help'. Philosophers call the latter 'normative reasons'. The logic of normative reasons is quite different from that of reasons-for-which. When A is φ -ing on the ground that p , it follows that A is φ -ing, and arguably that A believes that p ; it is at least doubtful whether it follows that p .⁵ When the fact that p is a reason for A to φ , it follows that p , but not that A believes that p or that she is φ -ing. Other connections are in dispute. Does it follow, when A is φ -ing on the ground that p , that she represents the fact that p as a normative reason to φ ? Some philosophers say yes; I argue that the answer is no.⁶

What I want to address now is not that question but a more abstract one, about the aims and ambitions of action theory. In my view, the principal aim can be stated quite simply. We want to know if the following principles can be completed without circularity, and if so, how:

To φ intentionally is to φ . . .

To φ on the ground that p is to φ . . .

Our attempts at a theory of what it is to act intentionally, or to act on the ground that p , ought to explain how these phenomena relate to others: why acting on the ground that p requires the belief that p , assuming it does, how it involves the agent's intentions, what intentions are, and so on. These are among the issues addressed by the essays in Part I.

Unfortunately, we cannot leave the subject here. For the philosophical treatment of action is often introduced in ways that conflict with mine, ways that import demands extraneous to action theory, or that leave its object obscure. Most prominent here is the invocation of Wittgenstein, who asked,

4. Anscombe 1963: 25; Davidson 1963: 6.

5. See Dancy 2000: 132. I expand on this in Setiya 2011: 132–134.

6. In Setiya, 'Sympathy for the Devil' (this volume: Ch. 3) and 'Akrasia'.

in the *Investigations*, ‘What is left over if I subtract the fact that my arm goes up from the fact that I raise my arm?’ (Wittgenstein 1953: §612). In my view, this question is seriously misleading.⁷ It is unhelpful, first, in taking a ‘subtractive’ form, which suggests an additive theory: what is left over when I subtract the fact that my arm goes up from the fact that I raise my arm is X, so raising my arm consists in X plus my arm going up. Looking for theories that have this shape is arbitrarily restrictive. Compare a simple approach on which I raise my arm just in case my arm goes up because I intend it to. If you subtract the fact that my arm goes up, what is left is my intention, causing nothing. You cannot construct an arm-raising, even on this simple approach, just by adding intention to my arm’s going up: you need the causal relation.

This defect is superficial; we need not assume that the account of what it is for me to raise my arm will take a conjunctive form. Instead, we can ask whether and how it is possible to complete this formula without circularity:

For me to raise my arm is for my arm to go up . . .

But this, too, is unhelpful, because it is too general. We can ask a similar question about the application of any transitive verb. Is there a non-circular completion of principles like these?

For the flower to open its petals is for the petals to open . . .

For the fire to melt the ice is for the ice to melt . . .

What fills the ellipses may be a further conjunct, a causal explanation, or something else. The project of spelling it out is not specific to intentional action, nor is it clear what motivates it. Why think that the application of transitive verbs can be explained in terms of their intransitive counterparts? And why suppose that the question is philosophically urgent? Is there some basis for the primacy of the intransitive? A puzzle in the metaphysics of transitive verbs?

A more radical but more principled approach would aim at a reduction of dynamic phenomena in general. The contrast between static and dynamic properties corresponds to the linguistic contrast between verbs that take progressive or perfective aspect and ones that do not.⁸ Some verbs have two

7. As Wittgenstein would agree, though for different reasons: ‘When I raise my arm I do not usually *try* to raise it’ (Wittgenstein 1953: §622). There is no inner state—trying, willing, intending—whose presence is a condition of raising my arm.

8. A classic treatment is Comrie 1976, though the distinction has philosophical roots; see Vendler 1957; Kenny 1963: 171–186. Later discussions include Mourelatos 1978, Graham

forms, one progressive—‘The floor was shaking’; ‘He was buying a house’—the other perfective, indicating completion or the fact that something happened: ‘The floor shook’; ‘He bought a house’. Others admit no such distinction: ‘The fruit was red’; ‘She knew everything’. These sentences do not report a completed act or event, but a state or condition that something was in. When verbs of the first kind are used in the present tense, they either have progressive aspect—‘The floor is shaking’; ‘He is buying a house’—or they are habitual, indicating a repeated or serial action: ‘The floor shakes’; ‘He buys houses’. Outside of special contexts, like certain forms of narrative, there is no present perfective. Verbs of the second kind, which admit no distinction of progressive and perfective aspect, have a non-habitual use in the simple present: ‘The fruit is red’; ‘She knows everything’.

Though it is introduced linguistically, the distinction here is metaphysical. Some of the things we predicate of objects can be instantiated ‘perfectively’ and in that sense done, while others cannot. There is no standard terminology for this distinction. We can use ‘state’ for properties that lack perfective instantiation. But there is no obvious term for the rest. It is tempting to call what can be done in the perfective sense an act. But in this sense, acts can be performed by inanimate objects, like the flower or the fire, that fall outside the scope of action theory. We might try ‘event’. But there are problems here, too. What we mean to identify, in contrast with states, are things predicated of or instantiated by objects, picked out by verbs like ‘shake’ and ‘buy’. ‘Event’ is typically used, instead, for the referents of noun phrases like ‘the shaking of the floor’, ‘his purchase of a house’. Though there is a close relation here, events in this sense are not our primary topic.⁹ We are interested in what it is for agents to do things, to instantiate properties of certain kinds. Because I cannot think of a noun to contrast with ‘state’ that is neither misleading nor arbitrary, I use the adjective ‘dynamic’ for the properties in question.

With this background, we can locate a possible project, of explaining what it is to instantiate a dynamic property—to shake, or buy something—in terms of states of objects and relations among them. Perhaps there is metaphysical pressure to think of reality as fundamentally static. But although it may be more principled than the ‘primacy of the intransitive’, it is clear that this project, too, has no essential place in action theory: in an account of what it is to

1980, and Galton 1984. I explore the distinction, and its relation to epistemic agency, in Setiya 2013b.

9. See Hornsby 1997: 87–92 on actions as events and as things done.

act for reasons. It is a project in general metaphysics that action theorists may or may not embrace.

What, then, is the project of action theory? It is unhelpful to introduce it by citing the difference between things I *do* and things that merely happen to me, as if emphasizing that word is enough to specify our topic. In one sense, opening the petals is a thing the flower does, and opening is what the petals do: these are doings in that they are dynamic properties. What is more, it is clear that when the doctor taps my knee, I kick my leg. That is something I do, not something that just happens to me, even when I do it by reflex. It is very different when the doctor lifts my leg to examine it more closely. That merely happens to me: I do not lift my leg; he does. But this is not the line that action theorists want to draw. When they distinguish the actions that interest them from mere happenings, kicking my leg by reflex is meant to fall in the second class.

Nor does it help to emphasize the word 'I': 'What is the difference between things I do and things that merely happen to me?' In ordinary terms, it is clear that I kick my leg as a matter of reflex action. The sense in which it wasn't really me, if any, is not something of which we have a pre-theoretic grasp, to be illuminated by the philosopher, but a fragment of tendentious philosophy, on which acting for reasons is explained in terms of the agent's identity. That might be right, but we should not appeal to it in specifying what we want to explain. The same is true of 'identification', which is a technical term in need of definition. Once defined, it is a term we may use in giving a philosophical account of agency, but we need first to locate the target of that account.

To say that our topic is 'action', unqualified, is not to make progress, since the term could apply to the doings of flowers and fires, and to kicking my leg by reflex. We need to identify a special class of actions in this encompassing sense. In my view, the right way to do so is the one with which I began. We want to know if the following principles can be completed without circularity, and if so, how:

To raise my arm intentionally is to raise my arm . . .

To raise my arm on the ground that *p* is to raise my arm . . .

The action theorist may take for granted the existence of dynamic phenomena, like raising my arm, in explaining what it is to do so intentionally, or for reasons. She need not attempt to answer Wittgenstein's question in any form. There is nothing illicit about this. It may turn out, in the course of understanding intentional action, that we need a reductive account of some

dynamic properties, or all of them. But we should not assume this from the start. It would be a substantial discovery that action theorists must, or can, take on such metaphysical ambitions.

Here I agree with the conclusion, if not the argument, of a notorious passage in Anscombe's *Intention*, §19:

[In] describing intentional actions as such, it will be a mistake to look for *the* fundamental description of what occurs—such as the movements of muscles or molecules—and then think of intention as something, perhaps very complicated, which qualifies this. The only events to consider are intentional actions themselves, and to call an action intentional is to say it is intentional under some description that we give (or could give) of it. (Anscombe 1963: 29)

Referring to actions 'under descriptions' is Anscombe's way of focusing attention not on events picked out by noun phrases—'Kieran's raising of his arm'—but on properties instantiated by agents. In our terms, her point is that we should take for granted the description of what an agent is doing when he is φ -ing intentionally and investigate how he is doing it. That is what I have been urging on methodological grounds, not by arguing that it is impossible to approach the topic with further reductive aims, but that it is unnecessary.

As it happens, Anscombe believes that the search for a non-circular completion of our formula, 'To φ intentionally is to φ . . .', is hopeless. But this requires a further argument. Despite appearances, the conclusion of §19, that 'an action is not called "intentional" in virtue of any extra feature which exists when it is performed' (Anscombe 1963: 28) is not the conclusion she later reports in similar terms, that being the execution of intention is not 'a mere *extra* feature of events whose description would otherwise be the same' (Anscombe 1963: 88). The conclusion of §19 is that *if* there is a non-circular account of what it is to φ intentionally, it will be in terms of φ -ing, not in terms of 'preintentional movements' accompanied by some special feature, *I*. The conclusion of the later argument, in §§46–48, is that there is no such account. It is addressed to a view that tempts us when the description of what someone is doing could apply to behaviour that is intentional or not: 'What is the difference between merely φ -ing and φ -ing intentionally, or on the ground that *p*?' What is the extra feature of this event, whose description would otherwise be the same?

Anscombe's reason for doubting that there is any such feature turns on the involvement of practical knowledge, and knowledge of reasons, in intentional

action.¹⁰ Although I believe that there is real insight in Anscombe's appeal to such knowledge, I do not agree with her about its nature or scope, and I do not believe that it precludes a non-circular theory of intentional action. In the next two sections, I say more about my commitments in this area, how they have shifted over time, and how they are traced through the essays in Part I.

2. Knowledge in Intention

Begin with a simple claim:

(K) When A is φ -ing intentionally, A knows that he is φ -ing.

In *Thought and Action*, Stuart Hampshire wrote that 'if a man is doing something without knowing that he is doing it, then it must be true that he is not doing it intentionally'; 'doing something [...] intentionally [...] entails knowing what one is doing' (Hampshire 1959: 95, 102). Two years earlier, Anscombe published the first edition of *Intention*, holding that intentional action is that 'to which a certain sense of the question "Why?" is given application; the sense is of course that in which the answer, if positive, gives a reason for acting'; this question is 'refused application by the answer: "I was not aware I was doing that"' (Anscombe 1963: 9, 11). Anscombe went on to classify knowledge of our intentional actions as 'knowledge without observation', meaning not only that it is not a matter of perceiving what we are doing, but that it is not inferred from other facts we know about ourselves.¹¹ If K is a necessary truth, there is pressure to reject an inferential model of the knowledge it ascribes: what could prevent me from forming and executing the intention to φ , thus φ -ing intentionally, without making the relevant inference?¹²

Much of the interest and the controversy around such knowledge is epistemic. How is it possible, critics ask, for me to know what I am doing except on the basis of sufficient prior evidence?¹³ Difficulties here might lead us to question Anscombe's insight, or Hampshire's, in stating K. But there are more mundane objections to their approach. For there are apparently obvious, compelling exceptions to K. Perhaps the most notorious is due to Davidson, who

10. See 'Anscombe on Practical Knowledge' (this volume: Ch. 6).

11. This is made explicit at Anscombe 1963: 50.

12. I make this argument at greater length in 'Practical Knowledge' (this volume: Ch. 1).

13. See, especially, Grice 1971; Langton 2004; Paul 2009.

imagines a carbon-copier trying to make ten copies at once.¹⁴ If he is succeeding, Davidson suggests, the carbon-copier is making ten copies intentionally. But he may not know, or even believe, that the copies are going through so many times. Davidson concludes: ‘It is a mistake to suppose that if an agent is doing something intentionally, he must know that he is doing it’ (Davidson 1978: 91).

In his first discussion of this problem, Davidson concedes that, while the carbon-copier may not know that he is making ten copies, what he is doing ‘is known to him under some description’ (Davidson 1971: 50). Inspired by this, I once proposed a replacement for K:

(B) When A is φ -ing intentionally, A believes that he is φ -ing, or else he is φ -ing by doing other things, in which he does believe.¹⁵

Thus, the carbon-copier is making ten copies by pressing hard on the carbon-paper, he is pressing hard intentionally, and he believes that he is doing so. Since beliefs of this kind are necessarily present in intentional action, there is the same pressure to deny that they are formed by inference. The revision preserves what is most interesting, and most puzzling, in K.

But it does not go far enough. Suppose, for instance, that I have been recently paralysed, and I attempt to clench my fist, under anaesthetic, behind my back. If I succeed, I am clenching my fist intentionally, but I do not know, or believe, that I am doing so. Nor do I clench my fist by taking further means that meet this condition: means I know, or believe, I am taking. This prompts a further revision:

(C) When A is φ -ing intentionally, A is more confident that he is φ -ing than he would otherwise be.¹⁶

This condition is met in recent paralysis. But the examples do not end. Thus Sarah Paul imagines a case in which I am less confident that I am φ -ing when I do so intentionally than otherwise. She appeals to the neurological disorder ‘Alien Hand Syndrome’, in which I am disposed to perform elaborate actions

14. Davidson 1971: 50, 1978: 91–92.

15. Adapted from ‘Belief’, in Setiya 2007: 26; see also ‘Practical Knowledge’ and ‘Practical Knowledge Revisited’ (this volume: Ch. 2).

16. This formula adapts and simplifies the treatment in ‘Practical Knowledge’; see also Pears 1985: 79–82.

unintentionally (Paul 2009). But we need not go so far. If I try to breathe steadily, my confidence that I am doing so may be less than when I trust my breathing to my autonomic system.¹⁷ Still, if I am breathing steadily, I am doing so intentionally. Whether C is refuted by this case depends on the relevant contrast, the situation to which we look in determining how confident I would otherwise be. Is it a situation in which I do not try to breathe steadily and my autonomic system kicks in, or one that eliminates my attempt without replacing it with another potential cause? Is there a principled way to decide?

One might respond to these problems by revisiting K.¹⁸ When we reflect on how little is involved in an action's being in progress—that one can be φ -ing but never φ , or even get close to doing so—can we credit A with knowledge that he is φ -ing even in the cases described above? If it is enough to be in the process of φ -ing that one take means by which one intends to φ , the carbon-copier can know that he is in the process of making ten copies, even if he does not know that he will—as you might know that you are building a house that in fact you will never complete. If it is enough that one intend to φ , I can know I am in the process of clenching my fist even when I doubt that it will move.

I do not find these options credible. Anscombe is right to reject the 'false avenue of escape' (from doubts about practical knowledge) on which 'I really "do" in the intentional sense whatever I think I am doing' (Anscombe 1963: 52). If I intend to clench my fist and it fails to move, an observer would flatly deny that I am clenching it. Nor is there some hidden ambiguity here, 'two objects of knowledge', practical and theoretical, picked out by the same words (Anscombe 1963: 57). The proposition I know when I know that I am clenching my fist because I am doing so intentionally is the one an observer knows when he sees it begin to move; and it is one whose truth I cannot know when I believe I am still paralysed.

A better way to save the idea of practical knowledge is not to weaken its content but to turn from instances of such knowledge to the capacities that afford it.¹⁹ The capacity to act for reasons, and so to act intentionally, is a capacity for knowledge of what one is doing and why. Like other cognitive capacities, this one can be exercised imperfectly, issuing in mere belief instead

17. See 'Practical Knowledge Revisited.'

18. Perhaps inspired by Falvey 2000 or Thompson 2008: Part One. I consider this response more fully in 'Anscombe on Practical Knowledge'; see also the notes on Thompson below.

19. As in 'Knowledge of Intention' (this volume: Ch. 4) and 'Knowing How' (this volume: Ch. 5).

of knowledge, or in partial belief, a degree of confidence that p . It is this capacity that action theory needs to comprehend. In my view, it should do so by picturing intention as a cognitive state, one that amounts to knowledge, or belief, in many cases, but which is consistent with serious doubt.

This way of putting things prompts several questions. First, it is one thing to insulate the alleged insight about intentional action and knowing what one is doing from obvious counterexamples. It is another to argue for its truth. Why think of the capacity for intentional action as a cognitive capacity at all? Second, there is still the question how such knowledge is possible, how we can know what we are doing except on the basis of sufficient prior evidence. And third, there is the need to say more about the kind of capacity that constitutes the will: what makes this capacity practical as well as epistemic? Let me take these points in turn.

As to the first, there are several reasons. For one, the default expression of my intention in acting takes the form of an assertion, ‘I am ϕ -ing’—or in the case of prospective intention, ‘I am going to ϕ ’—that looks like the expression of belief. This comes out in teleological explanations of action that take what Michael Thompson calls a ‘naïve’ form.²⁰ Asked, ‘Why are you boiling water?’, I reply that I am making tea, since I am boiling water in order to make tea. This makes sense if the assertion, ‘I am making tea’, expresses the intention with which I am boiling water, not some further mental state. In addition, there are indirect arguments for the cognitive conception of the will, some of which appear below. There is an argument from the nature of instrumental reason, an argument from self-knowledge of intention, and an argument from the epistemology of knowing how.²¹ Taken together, these arguments convince me that the cognitive conception is right.

On the second question, I explain the possibility of practical knowledge—knowledge in intention of what one is doing or what one is going to do—by appeal to knowledge how. This view occupies an intermediate place between two extremes. According to the practical knowledge sceptic, one is never justified in forming a belief without sufficient prior evidence, as one does in forming an intention on the cognitive account.²² On the other, permissive extreme, one is justified in forming a belief so long as one knows that the

20. In Thompson 2008: Part One.

21. See, respectively, ‘Cognitivism about Instrumental Reason’ (this volume: Ch. 8), ‘Knowledge of Intention’, and ‘Knowing How’.

22. See the authors cited in note 13.

belief will be true, and supported by evidence, once formed.²³ I agree with the sceptic in finding this insufficient. The force of the complaint is clear in connection with knowledge. When you come to know that p , the truth of your belief is credited to you: it turns on dispositions or capacities whose operation makes the truth of that belief no accident. This condition is absent from the permissive view, which is thus too liberal. But the sceptic's position is too extreme. For the condition can be met in at least two ways: by forming beliefs on the basis of sufficient prior evidence, or by forming beliefs one is disposed to make true. One is justified in forming the intention to φ , with its correlative belief, only if one knows how to φ , where knowing how to φ is a disposition to execute that intention whose operation makes the truth of one's belief no accident. Since knowledge how comes by degree, the degree of reliability in one's disposition, we should expect a similar gradation in the strength of the beliefs one is entitled to form.

The details of this sketch are spelled out in Part I, though it leaves some questions open. One, in particular, I would like to close. In coming to know that one is φ -ing, or that one is going to φ , where this is knowledge in intention, must one know, or be in a position to know, that one is able to φ ? I think the answer must be yes. We would otherwise license an illicit form of 'bootstrapping' in which one decides to φ , comes to know that one is φ -ing on the basis of knowing how, and then infers, without evidence, that one is able to φ .

What prevents such bootstrapping is the demand that one know, or be in a position to know, that one is able to φ . But this demand is potentially problematic.²⁴ In some cases, one has empirical knowledge that one is able to φ , as for instance on the basis of past attempts. But not in every case. What happens when one is not in a position to know in advance that one is able to φ ? How can it be rational to attempt a new intentional action—say, riding a bike—for the very first time, if my intention in acting involves the belief that I am doing it? In some cases, the answer may be that I do not intend to φ but to take some possible means: to push the pedals, hold the handlebars, and so on. But this is not a general solution. For in other cases, it is clear that I intend to φ even though I am not in a position to know that I can. (I do not merely intend to try, since my intention is not fulfilled if I try but fail.) This must be true when I learn new basic actions, not performed by taking further means.

23. See Harman 1976: 164n8; Velleman 1989: 56–64.

24. The problem is raised, but not resolved, in 'Practical Knowledge Revisited' (this volume: 68). As I note in that discussion, the issue here is parallel to problems of bootstrapping and 'easy knowledge' that appear in other domains.

The key is to remind ourselves that belief comes by degree. I may not know, or believe, that I am acting as I intend, though I have some confidence that I am. Likewise, while I may not be in a position to know that I am able to φ when I begin to learn, I must already be entitled to some degree of confidence that I am able to φ . This is less peculiar than it seems. What I am in a position to know is, in effect, that I might be able to φ , that there is some prospect of success. It is this minimal entitlement, together with knowledge how, that makes it rational to try. As to where the entitlement originates, there are two options. One is that, in order to be rational in attempting to φ , one must have at least a shred of empirical evidence that one can do it, perhaps from one's success at related activities, or one's general ability to control to one's body. The other is that one's confidence is justified a priori. Either way, the possibility of learning how can be made consistent with a cognitive theory of intention if, and only if, we allow for partial belief.

The final question raised above is how the capacity for practical knowledge differs from other cognitive capacities. What makes it distinctively practical? Part of the answer is contained in my remarks on knowing how. When one knows that p , it is not an accident that one's belief is true. In the case of theoretical knowledge, this condition is met by the origin of one's belief, its being formed by a method whose reliability is not an accident, as by appeal to sufficient prior evidence. In the case of practical knowledge, the condition is satisfied in part by one's disposition to act on the belief that figures in one's intention, a disposition that constitutes knowledge how.

This contrast—between two ways in which knowledge can be secured, practical and theoretical—is reminiscent of Anscombe on mistakes of performance and of judgement.²⁵ But it is quite different. Anscombe's point is not about the ground of non-accidental truth but about the locus of error. Anscombe distinguishes the standard of mistakes imposed by intention from the one applicable to belief. When you believe p and p is false, your belief is mistaken. That is a mistake of judgement. For Anscombe, there is no mistake of judgement when you are not acting as you intend, even though you think you are φ -ing and the

25. Anscombe 1963: §32. This section is often cited as the source of a metaphor, that attitudes can be distinguished by their 'direction of fit', with belief on one side, and desire on the other. These states relate to the world in opposing ways, beliefs being meant to fit the world, desires to make the world fit them. As Kim Frost has argued, this is both a misreading of Anscombe, whose discussion is not metaphorical or about mental states in general, and dubious in itself; see Frost 2014.

proposition that you are φ -ing is false.²⁶ There may be a mistake of judgement in the background, when your failure rests on a false belief about means to ends; but that is a separate matter. Moreover, when ‘a man is *simply* not doing what he says’, in that his failure does not rest on false beliefs, ‘the mistake is not one of judgment but of performance’ (Anscombe 1963: 57).

The condition just described is sufficient for performance error. Is it also necessary? Anscombe does not say. But she seems right to insist that one’s intention in acting sets a standard for what one does, in relation to which one can make mistakes. In this respect, the capacity for practical knowledge differs sharply from other cognitive capacities, a difference we can add to the one described above. It is much less clear that Anscombe is right on the negative point, that there is no mistake of judgement when this capacity misfires and one is not acting as one intends. On the more natural view, intention involves belief, or partial belief, and is subject to the same condition of error. When I think I am pressing button A, because that is what I intend, but I am pressing button B, I mistakenly believe that I am pressing button A, and I make a mistake in pressing button B. I make mistakes of judgement *and* performance. In support of this we can cite the fact that knowledge implies belief, so that intention involves belief when one has knowledge in intention, and the fact that one can make inferences from practical knowledge, or would-be knowledge, in just the way one does from other beliefs.²⁷

At the same time, intention is distinctive not just because its standing as knowledge is secured in a distinctive way, and because it sets a standard for mistakes of performance, but because it is in the nature of intention to motivate action. Intending involves the kind of wanting whose ‘primitive sign’, in Anscombe’s words, is ‘trying to get.’²⁸ What this means is, roughly, that when

26. Anscombe is, if anything, more explicit in the case of prospective intention: ‘If I do not do what I said I would, I am not supposed to have made a mistake, or even necessarily to have lied; so it seems that the truth of a statement of intention is not a matter of my doing what I said. But why should we not say: this only shows that there are other ways of saying what is not true, besides lying and being mistaken’ (Anscombe 1963: 4).

27. It is a good question why Anscombe denies that there is a mistake of judgement when I am not acting as I intend. Perhaps she doubts that it would be rational to form an attitude to p that is in error when p is false unless one has sufficient evidence that p is true. That principle would count against the forming of intentions, if intention involves belief. But I do not see the force of the principle, once we accept the reality of practical knowledge.

28. Anscombe 1963: 68. To think of intention as both cognitive and motivating or ‘desire-like’ (in the terminology of Setiya 2007) is to avoid the problem of Parfit’s insomniac (discussed in Harman 1976: §III), whose belief that he will stay awake is self-fulfilling. No matter what we add to the content of this belief—perhaps he believes that he will stay awake because he so believes—it does not count as the intention to stay awake.

one intends to φ and one can do so directly, without taking further means, one is disposed to φ in execution of one's intention; if one is capable of φ -ing, but only indirectly, by taking further means, one is disposed to intend the apparent means because one intends to φ . These facts about the 'functional role' of intention might explain why it sets a standard for performance error, as a kind of malfunction. This standard is clearly violated when one is simply not acting as one intends, pressing B when one intends to press A. It is less clear whether it is violated when one's failure to act as one intends derives from false beliefs about means, a point reflected in Anscombe's account.

When I execute my intention, that explains what I am doing: I am φ -ing because I so intend. Is this a causal explanation? In an essay written after *Intention*, Anscombe emphatically says no: it is a 'mistake [. . .] to think that the relation of *being done in execution of a certain intention*, or *being done intentionally*, is a causal relation between act and intention' (Anscombe 1983: 95). But the issue is obscure. Anscombe assumes that the relata of causation are states, where 'a state is supposed to be something holding of its subject here and now, or over a period of time, without reference to anything outside that of which it holds or the time at which it holds' (Anscombe 1983: 99). Intending is not a state in this sense, since whether I intend to visit the bank depends not only on what is true of me, here and now, but on whether I am acquainted with banks, that is, on facts about my past environment. An intrinsic duplicate of me in a world without banks could not intend to visit one. The proper response to this observation is not to deny that intentions are causes, but to resist the view that causation and causal explanation can only advert to 'narrow' or 'local' states, in Anscombe's sense.²⁹

The best way to make sense of the explanation of action by intention is to think about the manifestation of dispositions or the activation of powers.³⁰ The glass breaks when struck because it is fragile, in that it is disposed to break when struck: it manifests that disposition. The radioactive isotope decays because it is disposed to do so: its disposition to decay is realized. In a similar way, intentions interact with knowledge how, a dispositional state. To manifest know-how is to execute one's intention in action; to act intentionally is

29. See Yablo 1997.

30. Here I agree with Hyman 2013. We differ in that he treats desire, and so intention, as a disposition to act. I think he omits the pivotal role of knowing how. One can intend to φ but not be disposed to φ because one does not know how to φ . And the disposition one manifests in acting intentionally is a disposition to execute one's intention in action that constitutes knowledge how.

to manifest know-how.³¹ If explaining something as the manifestation of a disposition is causal explanation, so is the explanation of action by intention and knowledge how.

In each case, there is room for ‘causal deviance’. A fragile glass can break when struck in a way that does not manifest its disposition to do so: suppose it is attached to an explosive device, triggered by contact, that would destroy even a sturdy object. When a glass manifests its disposition to break when struck, its breaking must be caused by being struck ‘in the right way’. Similarly, I may blink because I intend to do so without manifesting knowledge how to blink: suppose my intention makes me nervous and I am caused to blink involuntarily.³² Then I do not blink intentionally or in execution of my intention. For that to happen, my intention must explain my blinking ‘in the right way’. What counts as the right way depends on the sort of disposition involved. In connection with knowing how, it is useful to distinguish two sorts of deviance: basic and non-basic.³³ A basic intentional action is one that is not performed by means of some other intentional action. Here the crucial concept is that of guidance: intention and know-how not only cause but guide one’s behaviour. Understanding what it is to manifest know-how is understanding the way in which it guides intentional action, correcting for perceived divergence from what the agent intends. When one performs an intentional action by means of others, to manifest knowledge how is to act in accordance with one’s plan; those who know how are disposed to formulate suitable plans and execute them in action.

There is one more source of resistance to intention as cause of intentional action, which is Aristotle’s claim that action itself is the conclusion of practical reasoning. Those who find this claim attractive recoil from a picture on which practical thought terminates with intention, a mental state distinct from, and productive of, intentional action. This picture is vividly expressed by John Broome:

[An] action—at least a physical one—requires more than reasoning ability; it requires physical ability too. Intending to act is as close to

31. Since know-how comes by degree, so does knowledge in intention: I am entitled to believe that I am φ -ing, when that is what I intend, only so far as I know how to φ and am entitled to be confident of my ability.

32. The case is adapted from Davidson 1973: 79.

33. See Davidson 1973: 78–79; Setiya 2007: 31–32.

acting as reasoning alone can get us, so we should take practical reasoning to be reasoning that concludes in an intention. (Broome 2002: 83)

What is at stake in Aristotle's claim? Presumably, it is agreed on all sides that there is such a thing as acting for a reason: reasons can attach to and inform what one is doing, not just what one intends to do. What is more, it is not just physical ability but also know-how that informs the causation of action by intention. How could the relation of reason to action be closer than this?

There are two ways. First, there is Michael Thompson's idea that intending to φ is a matter of being embarked on intentional action.³⁴ If one intends to φ , one is already in progress, though perhaps at an early or liminal stage. It follows that, if practical thought can reach as far as the intention to φ , it can reach as far as φ -ing itself. There is more to say on behalf of this view; but I am sceptical about it.³⁵ One can intend to φ without even incipiently φ -ing. Suppose I intend to perform a basic intentional action that I don't know how to perform. Am I bound to be in progress towards doing it? Alternatively, suppose I intend to do the impossible. I intend to be walking home by the shortest route but I have taken a serious wrong turn. Is there any sense in which I am in progress to walking home by the shortest route? There is a further difficulty for 'intending as doing' which turns on the logical complexity of objects of intention. I can intend not only to walk home but also to walk home if there are no taxis, not to walk home, to press button A or button B, and so on. If intention is an attitude, we can put such complexity in its object. If intending is being in progress, there must be intentional actions corresponding to arbitrarily complex concatenations of intentional actions; but it is quite unclear what these could be.

The second way to close the gap between intention and action is to go disjunctive.³⁶ Intention is not a factor or constituent of intentional action that might cause its execution but a condition that takes two forms: acting intentionally and mere intending. In the cases described above, in which I lack know-how or it is impossible to φ , I merely intend to act. When I act intentionally, I do not merely intend to act, not because a mere intention is effective, but because I am not in that state at all. Practical thought terminates with

34. Thompson 2008: Part One.

35. For more extensive discussion, see my entry on intention in the Stanford Encyclopedia of Philosophy, online at plato.stanford.edu/entries/intention, section 1; and for influential resistance, on the basis of 'pure intending', Davidson 1978.

36. As in Davidson 1978: 99; Rödl 2007: Ch. 1.

intention, that is to say, with mere intending *or* intentional action. In the latter case, Aristotle is right: the conclusion of practical reasoning is action itself.

One challenge for any disjunctive view is to say what unifies the disjuncts that we gather by a single name. There is a particular puzzle here in that the alleged disjuncts of intention seem to be of different metaphysical kinds: acting intentionally is dynamic, mere intending is static. What brings them together? The natural response, in the present context, is to identify intentional action with practical knowledge and to deny that knowledge reduces to belief.³⁷ Instead of being a constituent of knowledge, along with other factors, belief takes two forms: knowing that p and merely believing that p , the first of which is not explained in terms of the second. Likewise, intention is not a constituent of intentional action but a condition that takes two forms: practical knowledge of what one is doing and merely intending to act, the first of which is not explained in terms of the second. When I have knowledge in intention that I am φ -ing, it follows that I am φ -ing: knowledge entails and does not cause intentional action. Nor is it partly composed of intention, or practical belief, which could be cited as a cause.

A question for advocates of this approach is how they can distinguish practical knowledge from other forms of cognition without appeal to causal conceptions of intending and knowing how. But there is a more basic problem, if my arguments are right, namely that the contrast between practical knowledge and mere belief does not line up with the contrast between intentional action and mere intending. In some cases of partial belief, like that of recent paralysis, one executes one's intention and so acts intentionally, in the absence of practical knowledge. If intention is a cause of action here—we cannot complain that action is entailed, not caused, by intention as practical knowledge—why not elsewhere?

The upshot is a theory of intention as a mental state that involves both desire and belief or partial belief. The content of intention in action is the proposition that one is φ -ing; the content of prospective intention is the proposition that one is going to φ . Intention motivates action by way of knowledge how, and sets a standard for mistakes of performance in what one does. The same state is present when one fails to act as one intends. When one acts intentionally, intending is no less a cause of action than being struck is a cause of breaking in the fragile glass. In each case, a causal power is manifested: a

37. A theme of Williamson 2000. I mention this view briefly, and agnostically, in a footnote to 'Knowing How' (this volume: 151n26). Here I reject it, at least in application to practical knowledge and intention as the cause of action.

disposition to break when struck, or to execute one's intention in action. To act intentionally is to manifest knowledge how.

3. Acting for Reasons

How does this theory of intentional action account for Anscombe's question 'Why?' How does it explain what we do not just intentionally but for reasons? In Anscombe's book, these topics are inseparable. Intentional action is that to which the question 'Why?' is given application, in the special sense that asks for reasons; this question fails to apply when the agent does not know what she is doing or knows it only on the basis of sufficient prior evidence. In effect, the theory sketched above and developed in Part I builds an account of intentional action on the second condition: on the idea of knowledge in intention and its relation to knowing how. It treats acting for reasons as a special though pervasive case of this phenomenon.

What is involved in acting for a reason? Many philosophers hold that, in acting on the ground that p , one takes the fact that p as a normative reason for what one is doing. One must regard the fact that p as a consideration that counts in favour of one's action. One of the main contentions of my work on this topic is that such philosophers are wrong: it is not a condition of acting for a reason that one represent the ground on which one is acting as a normative reason to act. Nor does it help to weaken the content of the normative proposition (from a claim about reasons to one of approximate rationality) or the attitude one takes to this proposition (from belief to mere seeming). As I argue in 'Sympathy' and 'Akrasia', these conditions are equally flawed. In terminology that is now standard, we do not act intentionally, or for reasons, 'under the guise of the good'.

At the same time, there is a sense in which, in acting for a reason, one takes a consideration as one's reason for acting. The sense is not normative, nor is the attitude in question mere belief. It is rather that, as one intends to φ in φ -ing intentionally, so one intends to act for the reasons for which one acts. We decide not only what to do but why.³⁸ Since reasons figure in the content of our intentions, and intention is a cognitive state, acting for a reason involves belief, or partial belief, about the explanation of one's action. That is why we can respond to Anscombe's question when we are acting for reasons. If all goes well, we know what we are doing intentionally and why we are doing it.

38. For recent versions of this claim, see Wallace 1999: 239–242; Searle 2001: 16; Setiya 2007: 39–49. It draws on the Kantian idea that we will the maxims of our actions.

These remarks are compressed and they raise difficult questions. In particular, what is the content of one's intention in acting for a reason? What is the relation between reasons and causes? And if we have knowledge in intention of our reasons for acting, what is the epistemology of such knowledge?

On the first issue: it is tempting to suppose that the content of my intention and belief in acting on the ground that p is precisely that I am acting on the ground that p . That is what I know when I have knowledge in intention of what I am doing and why. But this has a troubling consequence: if we explain what it is to act for reasons in terms of one's intention in acting, and the content of that intention is that one is acting for a reason, our explanation is circular. The content of the intention cites the very phenomenon we are trying to explain. A version of the same point applies to intentional action, as such, if the content of the intention I execute is not simply that I am φ -ing but that I am doing so intentionally.³⁹ We could not explain what it is to act intentionally, without immediate circularity, in such terms. In 'Anscombe on Practical Knowledge', I argue that circularity of this kind plays a crucial role in the argument of Anscombe's *Intention*.

Can we learn to live with the circle?⁴⁰ The puzzle, from my perspective, is how something could appear as part of an account of its own nature. How can a true account of what it is to act for reasons appeal to acting for reasons? Perhaps we should say, instead, that acting for reasons is primitive: it cannot be explained in other terms. But then the mystery is why it involves intention. Why is it impossible to act for reasons without intending to act for them, if intending is not part of what it is to act for reasons?⁴¹ There would be no problem if the explanation ran the other way: if we could say what it is to intend to φ , or to φ on the ground that p , in terms of intentional action. But we already considered and rejected this approach, in both its forms: Thompson's theory of intending as doing, and the disjunctive theory of intending as intentional action or mere intention. Nor is it clear how to extend these accounts from the intention to φ to the intentions involved in acting for reasons.

Alternatively, we might look for precedents, in which entailments among properties go harmlessly unexplained. They cannot be derived from what it is to have those properties. This might be the case with the relation of

39. As suggested by John Searle (1983: 85–86), George Wilson (1989: 274–275), and Carl Ginet (1990: 34–35).

40. As Anton Ford suggests in Ford 2011: §4.

41. There is a similar puzzle about the entailment from acting on the ground that p to believing that p ; see 'Reasons and Causes' (Setiya 2011).

determinates to determinables, as for instance the relation of determinate shades of colour to the colours of which they are shades.⁴² Being scarlet entails being red, but it seems unpromising to explain being scarlet as being red and . . . what? There is no obvious way to complete this formula. On the face of it, being scarlet is a way of being red that is not to be explained in other terms. I do not think, however, that this provides a useful model for the kind of circularity we are considering. For one thing, just as being scarlet is not explained in terms of being red, so being red is not explained in terms of being scarlet. Both are understood through the dimensions of hue, saturation, and brightness that form the space of visible colours. By contrast, if acting intentionally implies intention, where the content of that intention is that one is acting intentionally, the alleged determinable, intending to ϕ , is explained through a specific determinate. Likewise for intending to act on the ground that p . It is in any case far from clear that the entailments and exclusions of determinates and determinables cannot be derived from non-circular accounts of what they are, ones that appeal to their location in the space of colours. With minimal structure, we can prove that accounts of this kind are available, though they may take unfamiliar forms.⁴³

The moral of these reflections is that there is no safe precedent for the inexplicable entailment of intention by intentional action and acting for reasons, if they are taken as primitive. Should we reconsider the circular account, explaining what it is to act intentionally, or for reasons, in terms of intentions whose content is the phenomenon to be explained? There are other views that take this shape: forms of ‘constructivism’ or ‘response-dependence’ on which facts are said to be constituted by our attitudes towards them. But their circularity is no less puzzling.⁴⁴ The sense of mystery is especially deep when the attitude in question is knowledge, or has been introduced, like intention, in terms of a capacity to know. Anscombe gives alleged examples of this: paying, hiring, marrying are said to be actions ‘which can only be voluntary or

42. The language of ‘determinates’ and ‘determinables’ is due to W. E. Johnson (1921). Whatever has a determinable property has one of its determinates, and when x is a determinate of y , x entails y but not the reverse. Determinable properties differ from mere disjunctions in that their nature involves one or more dimensions of variation, locations in which are occupied by their determinates.

43. See Moss 2012. Here I retract a concession made in Setiya 2011: 140–143.

44. See Johnston 1993; Street 2008. Street tries to avoid circularity, but it is not clear to me how she succeeds. The challenge is to maintain that the relevant judgement is about the very property whose nature is being explained while explaining the judgement in terms that do not mention the property. The risk is that we end up picturing the would-be ‘judgement’ as contentless or non-cognitive, or with a definition like this: to be F_1 is to be judged F_2 in condition C .

intentional' (Anscombe 1963: 85), only with performed with knowledge of what one is doing. But the cases do not convince. These are all things one can do unintentionally, though perhaps one must do so through other intentional actions. And Anscombe's alternative is perplexing. How can facts consist in knowledge of those very facts?⁴⁵

Nor are we compelled to embrace circularity. While it is true that the execution of intention is intentional action, it does not follow that the content of intention involves intentionality. If I intend to smile and I am doing so involuntarily, I am doing what I intend, though not intentionally. Likewise, if I intend to skip breakfast and do so because I forget to eat, my intention is fulfilled, though not by intentional action. What is true is that, as I argue in 'Knowledge of Intention', when I intend to φ , I am in a position to self-ascribe my own intention as part of its content: 'I am φ -ing, or I am going to φ , as I intend'. But it is not essential to the content of intention that it be self-ascribed, or that it represent an action as being performed intentionally.⁴⁶ There is nothing circular in an account of intentional action that cites the intention to act.

Turning to reasons, we can say that the content of my intention in φ -ing on the ground that p is that I am φ -ing because I believe that p . I do not think we should be concerned about the appeal to belief in the content of this intention, with its implication of self-knowledge. Grounds on which one acts must be contents of beliefs of which one is aware as such. You cannot act on the ground that p when you do not know that you believe that p . Nor should we fear circularity. We can explain what it is to act because one believes that p , as this figures in the object of one's intention, without appeal to acting for reasons.

Start with the fact that we are moved by beliefs in ways of which we are not aware as well as those of which we are. Without intending to act on the belief that p , and so without acting on the ground of its content—at least in the self-conscious way we have been exploring—I can act on desires that are influenced by that belief. Here, too, there is a difference between a desire that

45. Does Anscombe infer, mistakenly, from the premise that events of paying, hiring, and marrying involve, and so can be identified with, intentional actions, to the conclusion that they are things one can only do intentionally? Does she confuse actions as events with things done?

46. This modifies the view proposed in Setiya 2007: Part One. Intentions need not refer to themselves, as such, and the argument for self-reference in acting for reasons is about the metaphysics of φ -ing because of one's belief that p , which requires one's intention to play a role, not about the content of that intention. I develop these points below.

is motivated by the belief that p and one it merely causes. There is a question of deviance. In the non-deviant case, I manifest a disposition that is sensitive to my attitude and its content. It is important that what I manifest is a disposition to φ because I believe that p , and not a disposition to φ because I have some determinable property of which believing that p is a determinate.⁴⁷ (In the latter case, it is the determinable property that motivates my desire.) How is it different when I act on the ground that p ? Do I intend to φ because I believe that p in a sense that requires me to meet these conditions? No, since I can act on the ground that p without being disposed to φ when I believe that p , except when I also intend to do so. The disposition I manifest is a disposition to φ when I believe that p and have the relevant intention. To φ because I believe that p , as I intend when I act on the ground that p , is to φ because I believe that p and because of that intention, in that I manifest a disposition that is sensitive to these attitudes and their contents, and which constitutes knowledge how. This account of what it is to φ because I believe that p , in the relevant sense, does not appeal to acting for reasons or to conditions of non-deviance beyond those involved in the motivation of desire by belief and of action by intention.⁴⁸ The epistemological theory is the same as before: practical knowledge rests on knowing how to execute one's intention—the intention to φ because one believes that p —and on knowing, or being in a position to know, that one is able to do so.

Having set out in abstract terms the conditions of acting on the ground that p , it is useful to distinguish two varieties. One is instrumental: I am acting in pursuit of a further end. In a paradigm case, I am pressing the keys in order to type a sentence, I intend to type the sentence, and I believe that I am doing so. The content of my intention in pressing the keys is that I am doing so because I intend to type the sentence and believe that pressing the keys is a means to that end. In acting on this intention, I rely on self-knowledge—that I intend to type the sentence and believe that pressing the keys is a means to that end—and on knowing how to type the sentence by pressing the keys, a disposition to act on that intention. My reason for acting is that I am going to type the sentence, or that pressing the keys is a means to doing so.

The other variety of acting on the ground that p is non-instrumental: I know how to put a belief into action, though not as a means to an end. In practice,

47. In other words, motivation by belief meets a version of Yablo's (1992) 'proportionality' condition.

48. A more complete account would tease apart the contribution of first- and second-level explanations, as in Skow 2016: Ch. 6.

it may be hard to say whether a given case is instrumental or not. Why am I going for a walk? Because the weather is fine. I intend to act on my belief that the weather is fine and I know how to do so. But perhaps I am acting on a further intention, to get some sun, and the belief that the weather is fine is part of my conception of the means to my end. Elsewhere, this construction is less tempting, as when I help a friend because he is in need. Can acting from friendship be reduced to acting from intentions that involve one's friend? Is friendship just a matter of having certain ends?⁴⁹ Whatever is true of friendship, I see no general argument that acting for reasons can always be reduced to acting with a further intention, or that intention and desire can never be explained, or motivated, by belief alone.⁵⁰

Finally, as well as acting in order to φ or on the ground that p , we can allow for what Rosalind Hursthouse (1991) calls 'arational action': intentional action expressive of emotion, but not done for reasons in the ordinary sense. When I jump for joy or scream in frustration, I intend to jump because I am happy or to shout because I am at my wit's end. If I know that I have the relevant emotion and how to jump, or shout, because of it, I can have knowledge in intention of what I am doing. At the same time, I need not act with the further intention of expressing how I feel: I am not acting in order to express frustration or joy. Arational action thus differs from acting for a reason. It also differs from acting on the bare intention to φ , doing something for no particular reason, since one's intention in acting arationally includes a further explanation of what one is doing. The present account predicts this intermediate case.

For all its complexity, the account of intentional action that emerges from this discussion has a minimalist character. The core phenomenon is the capacity to execute intentions on the basis of knowledge how, and so to know what one is doing. We can explain the nature of this capacity, and thus of intentional action, without appeal to Anscombe's question 'Why?' We can then explain the capacity to act for reasons—on the ground that p , or in order to φ —by appeal to intentions whose content includes an explanation of one's action. So long as you have knowledge of the relevant states of mind, and you know how to act on them, nothing further is required. You need not regard the reasons for which you act as considerations that count in favour of what you are doing. Your motivation need not have an instrumental structure,

49. Michael Stocker (1981) argues, persuasively, that it is not.

50. On this point, see Nagel 1970: 27–30; Wallace 1990: 373–374; Setiya 2007: 100–106.

appealing to desires whose satisfaction you are promoting. Nor must it be approximately rational. If you know how, you can act on grounds irrelevant or perverse, just as you can act, knowingly and intentionally, for no reason at all.

4. From Action to Ethics

Suppose I am right about the metaphysics of intentional action. What implications does this have for moral philosophy? What, if anything, can we learn about normative reasons, and thus about how we should live, from an account of what it is to act intentionally?

Some philosophers hold that the answer is everything, near enough. The standards of practical reason are explained by, and derive from, the nature of agency or the will. This is the rationalist or constitutivist approach. I am sceptical of it, for reasons that emerge below. But I think it has profound attractions and that it rests on an argument with real power. At the same time, rationalism is a threat to the universality of moral and altruistic reasons: if they apply to everyone, they must follow from inescapable aspects of the capacity to act intentionally, or for reasons, at all. It is a task for philosophers who aim to save morality without supplying that proof to explain how *not* to be an ethical rationalist. That is where the action theory sketched above comes in.

But there are more immediate implications, too. The most obvious is that, if intention involves belief in the way I have suggested, intentions are subject to the epistemic norms that govern belief, as such. This observation puts us in the vicinity of what Michael Bratman (1991) calls ‘cognitivism about practical reason’; but it is important to distinguish three grades of cognitivist ambition. The first grade is metaphysical. ‘Cognitivism about intention’ has been used for the claim that intentions involve beliefs about what one is doing or what one is going to do. ‘Strong cognitivists’ add that intention reduces to belief, though it is not clear to me that anyone in fact adopts this view.⁵¹ The second grade of cognitivism holds that some of the requirements that govern intention are explained by the application of epistemic rationality to the beliefs our intentions involve. A third and final grade of cognitivism about practical reason identifies the norms of good practical reasoning, and thus what counts as a reason to act, with norms of epistemic reason.

51. In particular, although Velleman (1989, 2000) identifies particular intentions with particular beliefs, these beliefs count as intentions only in the context of a background disposition to act in ways that make them true.

Perhaps the best way to bring out the differences here, and to explain why they matter, is to describe my own view, which admits the second grade of cognitivism but not the third. Specifically, in ‘Cognitivism about Instrumental Reason,’ I argue that, properly conceived, instrumental incoherence—failing to will the necessary means to one’s intended ends—is a form of epistemic incoherence. If you intend E, and so believe that you will do it, and you believe that intending M is necessary for doing E, but you do not intend M, we can ask: do you believe that you intend M?⁵² If you don’t, you fail to believe an obvious and practically salient consequence of what you believe: an epistemic failure. On the other hand, if you do believe that you intend M, your belief is not only false but also epistemically defective, since, for reasons explored in ‘Knowledge of Intention,’ intending M puts one in a position to know that one intends M by the exercise of reason alone. Either way, in failing to conform to the instrumental principle, you violate the norms of epistemic reason.

It is a difficult question whether and how this result extends beyond the special case of full belief and necessary means. As a first step, we can note that, even if one’s intention for E involves only partial belief, so long as one is sure that one will not do E unless one intends M one had better believe, to some degree, that one intends M. Assuming it is irrational to do so unless one intends M, the derivation above goes through. What about non-necessary means? It is more difficult to formulate strict principles here. Is it irrational not to intend the most effective means to one’s ends? Arguably not, since there may be reasons against them: efficiency is not the only concern. On the other hand, the partial belief involved in intending E, and partial belief that one will not do E unless one intends M, together place epistemic pressure on one’s degree of belief that one intends M, which it may be possible to meet without epistemic failure only by intending M. The details remain obscure.

The idea that epistemic rationality requires means–end coherence is not original to me.⁵³ I differ from other cognitivists in denying that the instrumental principle is a principle of practical reason. It is not that practical reason is here subsumed by epistemology, but that it is displaced. In fact, it is difficulties internal to the practical interpretation of means–end coherence that motivate the second grade of cognitivism and thereby motivate the first. On

52. On the reasons for this formulation, which appeals to intending M, not doing it, see ‘Cognitivism,’ §III (this volume: 194–202).

53. See Harman 1976; Wallace 2001.

the most promising view, it is a failure of practical reason not to [intend M, if one intends E] when one believes that intending M is necessary for E. That intending M is necessary for E is thus a decisive reason to [intend M, if one intends E].⁵⁴ Now suppose that one intends some terrible end and one cannot alter this intention: whatever intentions one forms, one will still intend E. In the corresponding sense, intending M is a necessary means to [intending M, if one intends E]. And suppose that intending M is necessary and sufficient for doing E. It follows that there is decisive reason to intend M. But that seems wrong: one should not intend M, since E will result!

In outline, this is the argument against practical interpretations of the instrumental principle in ‘Cognitivism about Instrumental Reason’. Its final step has been challenged by Jonathan Way (2010: §4). This step depends on a transition from decisive reason for an end, [intending M, if one intends E], to decisive reason for the means, intending M. Way objects that, while means–end transmission may hold for ‘state-given’ reasons for attitudes, ones that turn on properties of that attitude, it does not hold for ‘object-given’ reasons, which turn on properties of its object. Since the reason to [intend M, if you intend E] is object-given, it does not transmit, and the problem goes away.

I think the objection fails. First, the reasons invoked by the practical interpretation of the instrumental principle do not seem object-given. On the formulation above, which Way accepts, this principle appeals to the necessity of intending M, that is, to a property of intention, which is the mark of a state-given reason. But this is a minor point. Presumably, Way’s thought is that it is only the ‘wrong kind’ of state-given reason that transmits; reasons that play a proper role in reasoning to attitudes do not. This fits the case of belief, where the wrong kind of reason for belief—reasons to believe *p* that turn on the benefits of believing *p*—generate reasons to produce that belief, where reasons to believe that turn on evidence of truth do not. The objection, then, is that, while reasons for action may transmit from ends to means, the ‘right kind’ of reasons for attitudes do not.

A second response to Way goes deeper. Even if he is right that reasons for attitudes do not always transmit to necessary means, what I need for my argument is a special case:

54. Given the relationship of practical reasons to practical rationality proposed in *Reasons without Rationalism* (Setiya 2007: 7–14) and further developed in ‘What Is a Reason to Act?’ (this volume: Ch. 9). More on this below.

Intention Transmission: When you have decisive reason for a profile of intentions, and intending to do A is a necessary and sufficient constitutive means to having this profile, you have decisive reason to intend A.

That is why, when you have decisive reason to [intend M, if you intend E], and intending M is a necessary means to [intending M, if you intend E], you have decisive reason to intend M. This principle applies elsewhere. Suppose, for instance, that you have decisive reason to [intend A, if you intend B], since doing B without A would be disastrous, and you know that you will act as you intend. Suppose, further, that you cannot help but intend B; whatever intentions you form, you will act on this one. It follows, by Intention Transmission, that you have decisive reason to intend A. And that seems right.

In short, my argument against the practical interpretation of the instrumental principle survives Way's critique. If the instrumental principle holds, it is not a requirement of practical reason but of epistemology. This is cognitivism of the second grade, not the third. While I believe that certain requirements on intention are explained by its cognitive aspect—along with means–end coherence, there is the demand for consistency in one's intentions, and between intentions and beliefs about the future—I doubt that this dimension gives much insight into practical reasoning or practical rationality. There are several points to make here.

To begin with, the instrumental principle is not itself a guide to reasoning. The principle tells us that it is incoherent to intend E, believe that intending M is necessary for doing E, and not intend M. The incoherence hits only when intending M *now* is necessary for doing E; otherwise, one may coherently wait to form the relevant intention.⁵⁵ On the other hand, it would not be good reasoning to *form* the intention for M on the basis of one's intention for E and the belief that intending M now is necessary for doing E: by the time one gets around to doing this, it is too late! One must already intend M. What the principle describes is a condition of epistemic coherence on intention and belief, not the form of instrumental reasoning. The reasoning by which one conforms to this principle is pre-emptive: one stops intending E or starts intending M before one arrives in an epistemically defective state. How does such reasoning work, and what are its standards? The cognitive character of intention tells us nothing about that.

55. Again, see 'Cognitivism', §III (this volume: 194–202).

It is unfortunately common for philosophers to object to cognitivism by complaining, with Davidson (1978: 95), that ‘reasons for intending to do something are in general quite different from reasons for believing one will do it’, or that the belief involved in one’s intention could not be formed by good theoretical reasoning.⁵⁶ I find these complaints unfortunate because they are not careful to distinguish the third grade of cognitivism from the second grade or the first. Whatever their force against the former, they have no bearing on the latter. On the view that I defend, intentions involve beliefs, but we do not form such beliefs by reasoning from evidence of their truth. The reasoning by which we do so is distinctively practical. Nor should we equate the reasons that should figure in practical thought, normative reasons for action, intention, and desire, with normative reasons for belief.

As I understand it, the idea of a normative reason is the idea of a premise of sound reasoning; reasons for action are premises of sound reasoning to desires or motivations that influence what one does. It is reasons of this kind that determine what one ought to do. I defend this picture in ‘What Is a Reason to Act?’, and it forms the backdrop to the ethical inquiries that occupy the final essays in Part II. These essays are haunted by a seductive strategy. If what we ought to do is fixed by the balance of reasons, and reasons can be understood in terms of sound reasoning, we can answer the principal questions of ethics by explaining what it is to reason well. So far, I agree. But there is a further thought: if ‘good reasoning’ is analogous to ‘good thief’ or ‘good thermometer’, it may be possible to extract the standards of good practical reasoning from an account of practical thought, and thus from the nature of agency, as we can extract the standards for being a good thief or a good thermometer from the metaphysics of the corresponding kinds. This is the project of ethical rationalism or constitutivism about practical reason.

The strongest form of cognitivism is also constitutivist. It aims to extract the standards of practical rationality, or good practical reasoning, from the fact that reasoning to intention is reasoning to beliefs of certain kinds. This is David Velleman’s ambition in *Practical Reflection* and *The Possibility of Practical Reason*. I do not see how it can be done. There might be hope for Velleman’s approach if it were essential to rational agency not only that one know what one is doing but that one aim, perhaps implicitly, at knowledge that affords a richer understanding of oneself. That would support his view that normative reasons for action are ‘considerations that would provide the

56. As in Ross 2009: 267–268.

subject with an explanatory grasp of the behavior for which they are reasons' and that the strength of a reason turns on to the depth of insight it affords (Velleman 2000: 26). But we need not aim at greater comprehension of ourselves in order to act for reasons. Nor does the demand for explanatory depth gain purchase on our intentions simply because they involve beliefs, so that intentions and the reasoning that generates them are better, as such, to the extent that they give a more unified, deeper explanation of what one is about. Epistemic pressure towards the best explanation turns on the assumption that it is more likely to be true than some ad hoc alternative; that pressure does not apply to the beliefs that figure in our intentions, which are true whenever we act on them. Epistemic ideals of understanding and explanation are screened off by the standard of practical knowledge, for which knowledge of ability and the exercise of know-how suffice.⁵⁷

More tempting versions of rationalism turn not on the cognitive import of intention but on the role of desire in motivation, the conditions of autonomy, or the guise of the good. It is from these aspects of agency that they aim to derive the standards of practical reason. I reject them all, though not because I doubt the force of the derivation. In my view, the rationalist argument is difficult to resist: many philosophers unwittingly endorse conceptions of agency that support it. I explain how the argument goes, and how little it assumes, in §3 of 'Intention, Plans, and Ethical Rationalism' (this volume: 237–242). I develop it further, responding to objections by David Enoch (2006), in 'Akrasia and the Constitution of Agency'.

The attractions of rationalism or constitutivism are many. It offers an unmysteriously objective picture of practical reason, in which reasons for action are no more puzzling than standards of good reasoning, which are no more puzzling, in turn, than standards for being a good thermometer or good thief. But it has serious problems, too. One was mentioned above: it poses a threat to the universality of moral or altruistic reasons. If such reasons apply to us all, and rationalism is true, the demand for moral or altruistic reasoning must derive from the nature of agency or the will. That derivation is hard to produce. But perhaps we must try. A second problem is more basic: whatever its validity, the premises of the rationalist argument fail. This is the final moral of the action theory offered in Part I. Its minimalism about all but the cognitive conditions of agency deprives the ethical rationalist of materials with which to work. 'Good reasoning' is not analogous to 'good thermometer';

57. For further objections to Velleman, see Bratman 1991; Setiya 2007: 107–114.

‘good *qua* rational agent’ is not analogous to ‘good *qua* thief’. There is nothing in the nature of agency from which the standards of practical reason could derive.

Some will fear that this objection makes things worse. The result of undermining ethical rationalism is a threat not just to the scope of moral and altruistic reasons but also to the existence of reasons, as such. The upshot is nihilism about practical reason. But I do not believe this. If we reject rationalism, we doubt that the standards for reasoning well in practice are explained by the nature of agency or the will. We need not conclude that there are no standards at all, that there is no such thing as practical reasoning that is good or bad, as such. As I urge in ‘Hume on Practical Reason’ (this volume: Ch. 12), we can understand good practical reasoning as reasoning that manifests good dispositions of practical thought, and good dispositions as ones that are good as traits of character. What we learn from the failure of ethical rationalism is that there is nothing about the dispositions that bear on practical reasoning, among the dispositions that constitute our character, that explains why they are subject to proprietary norms. It does not follow that they are not subject to standards as the kind of thing they are.

In *Reasons without Rationalism*, I argue for what I take to be Hume’s conclusion—a ‘virtue theory’ of practical reason—without reference to the idiosyncrasies of his philosophy of mind. This brings out what Hume shares with Aristotle, despite the radical difference in their metaphysics: the standards of practical reason are those of ethical virtue, applied to practical thought. It follows that there is no room for the question ‘Why be moral?’ on a certain understanding of what it asks. There is no room to concede that justice and benevolence are virtues of character but deny that the facts to which they make one responsive are normative reasons to act. The only question is what the virtues are.

This question points towards the metaphysics and epistemology of ethical virtue, and of ethics in general, and so beyond the boundaries of action theory. How must we conceive the basic facts about virtue and practical reason so as make sense of objective ethical knowledge? If we treat them as primitive, we make our access to such facts mysterious—though it is not easy to say, with precision, what the mystery is.⁵⁸ More promising views appeal to human nature as the standard of ethical virtue, as in the neo-Aristotelian naturalism of Philippa Foot (2001). Or they share the spirit of ethical rationalism, as in

58. I attempt to do so in Setiya 2012.

forms of constitutivism that look beyond the nature of agency to the nature of the attitudes and concepts that play a role in practical thought.⁵⁹ Are the standards for thinking well somehow contained in the conditions for having these attitudes and possessing these concepts? I am dubious about both views: the neo-Aristotelian idea that human beings are by nature virtuous, and the idea that one cannot possess an attitude or concept unless one approximates the standards of reason that govern it, the standards for using it well. Nor do I think it is enough to appeal to a ‘social externalism’ on which one can possess a concept vicariously if one acquires it from a linguistic or conceptual community whose members meet this condition. A whole society could be disposed to go astray. But I think there is hope for a synthesis of these ideas. If ethical thought is anchored in the world, it must be through the conditions of ethical concept-possession, where these turn not on individual dispositions or on social norms, but on the kind of beings we are.

This is admittedly too brief to carry conviction. It is not the topic of this book, but another, *Knowing Right From Wrong*, which is addressed to practical knowledge in the third sense: knowledge of ethical facts. The essays in this volume are about the second sense and the first, about knowing how and knowledge in intention. And they are about the intersection of action theory and ethics. Moral philosophers should think about the metaphysics of agency not because it is the foundation of ethics, but in order to see that it is not.⁶⁰

References

- Anscombe, G. E. M. 1963. *Intention*. 2nd ed. Oxford: Blackwell.
- Anscombe, G. E. M. 1983. The causation of action. Reprinted in M. Geach and L. Gormally, eds., *Human Life, Action and Ethics*. Exeter: Imprint Academic, 2005: 89–108.
- Bratman, M. 1991. Cognitivism about practical reason. *Ethics* 102: 117–128.
- Broome, J. 2002. Practical reasoning. In J. L. Bermudez and A. Millar, eds., *Reason in Nature: New Essays on the Theory of Rationality*. Oxford: Oxford University Press, 2002: 85–111.
- Comrie, B. 1976. *Aspect*. Cambridge: Cambridge University Press.
- Dancy, J. 2000. *Practical Reality*. Oxford: Oxford University Press.
- Davidson, D. 1963. Actions, reasons, and causes. Reprinted in D. Davidson, *Essays on Actions and Events*. Oxford: Oxford University Press, 1980: 3–19.

59. For different versions of this approach, see Murdoch 1970 and Smith 2010.

60. Special thanks to Sarah Paul for comments on an earlier draft.

- Davidson, D. 1971. Agency. Reprinted in D. Davidson, *Essays on Actions and Events*. Oxford: Oxford University Press, 1980: 43–61.
- Davidson, D. 1973. Freedom to act. Reprinted in D. Davidson, *Essays on Actions and Events*. Oxford: Oxford University Press, 1980: 63–81.
- Davidson, D. 1978. Intending. Reprinted in D. Davidson, *Essays on Actions and Events*. Oxford: Oxford University Press, 1980: 83–102.
- Enoch, D. 2006. Agency, shmagency. *Philosophical Review* 115: 169–198.
- Falvey, K. 2000. Knowledge in intention. *Philosophical Studies* 99: 21–44.
- Ford, A. 2011. Action and generality. In A. Ford, J. Hornsby, and F. Stoutland, eds. *Essays on Anscombe's Intention*. Cambridge, MA: Harvard University Press, 2011: 76–104.
- Foot, P. 2001. *Natural Goodness*. Oxford: Oxford University Press.
- Frost, K. 2014. On the very idea of direction of fit. *Philosophical Review* 123: 429–484.
- Galton, A. 1984. *The Logic of Aspect*. Oxford: Oxford University Press.
- Ginet, C. 1990. *On Action*. Cambridge: Cambridge University Press.
- Graham, D. 1980. States and performances: Aristotle's test. *Philosophical Quarterly* 30: 117–130.
- Grice, H. P. 1971. Intention and uncertainty. *Proceedings of the British Academy* 57: 267–279.
- Hampshire, S. 1959. *Thought and Action*. Notre Dame, IN: Notre Dame University Press.
- Harman, G. 1976. Practical reasoning. *Review of Metaphysics* 29: 431–463.
- Hornsby, J. 1997. *Simple Mindedness*. Cambridge, MA: Harvard University Press.
- Hursthouse, R. 1991. Arational actions. *Journal of Philosophy* 91: 57–68.
- Hyman, J. 2013. Desires, dispositions and deviant causal chains. *Philosophy* 89: 83–112.
- Johnson, W. E. 1921. *Logic, Part I*. Cambridge: Cambridge University Press.
- Johnston, M. 1993. Objectivity refigured: Pragmatism without verificationism. In J. Haldane and C. Wright, eds. *Reality, Representation, and Projection*. Oxford: Oxford University Press, 1993: 85–130.
- Kenny, A. 1963. *Action, Emotion, and Will*. London: Routledge.
- Korsgaard, C. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Korsgaard, C. 2009. *Self-Constitution*. Oxford: Oxford University Press.
- Langton, R. 2004. Intention as faith. In J. Hyman and H. Steward, eds. *Agency and Action*. Cambridge: Cambridge University Press, 2004: 243–258.
- Moss, S. 2012. Solving the color incompatibility problem. *Journal of Philosophical Logic* 41: 841–851.
- Mourelatos, A. 1978. Events, processes, and states. *Linguistics and Philosophy* 2: 415–434.
- Murdoch, I. 1970. *The Sovereignty of Good*. London: Routledge.
- Nagel, T. 1970. *The Possibility of Altruism*. Princeton, NJ: Princeton University Press.
- Paul, S. 2009. Intention, belief, and wishful thinking. *Ethics* 119: 546–557.
- Pears, D. 1985. Intention and belief. In B. Vermazen and M. B. Hintikka, eds., *Essays on Davidson: Actions and Events*. Oxford: Oxford University Press, 1985: 75–88.
- Rödl, S. 2007. *Self-Consciousness*. Cambridge, MA: Harvard University Press.

- Ross, J. 2009. How to be a cognitivist about practical reason. *Oxford Studies in Metaethics* 4: 243–281.
- Searle, J. 1983. *Intentionality*. Cambridge: Cambridge University Press.
- Searle, J. 2001. *Rationality in Action*. Cambridge, MA: MIT Press.
- Setiya, K. 2007. *Reasons without Rationalism*. Princeton, NJ: Princeton University Press.
- Setiya, K. 2011. Reasons and causes. *European Journal of Philosophy* 19: 129–157.
- Setiya, K. 2012. *Knowing Right From Wrong*. Oxford: Oxford University Press.
- Setiya, K. 2013a. Murdoch on the sovereignty of good. *Philosophers' Imprint* 13: 1–21.
- Setiya, K. 2013b. Epistemic agency: some doubts. *Philosophical Issues* 23: 179–198.
- Skow, B. 2016. *Reasons Why*. Oxford: Oxford University Press.
- Smith, M. 2010. Beyond the error theory. Reprinted in K. Setiya and H. Paakkunainen, eds., *Internal Reasons: Contemporary Readings*. Cambridge, MA: MIT Press, 2012: 303–328.
- Stocker, M. 1981. Values and purposes: The limits of teleology and the ends of friendship. *Journal of Philosophy* 78: 747–765.
- Street, S. 2008. Constructivism about reasons. *Oxford Studies in Metaethics* 3: 207–245.
- Thompson, M. 2008. *Life and Action*. Cambridge, MA: Harvard University Press.
- Velleman, J. D. 1989. *Practical Reflection*. Princeton, NJ: Princeton University Press.
- Velleman, J. D. 2000. Introduction. In J. D. Velleman, *The Possibility of Practical Reason*. Oxford: Oxford University Press, 2000: 1–31.
- Vendler, Z. 1957. Verbs and times. *Philosophical Review* 56: 97–121.
- Wallace, R. J. 1990. How to argue about practical reason. *Mind* 99: 355–385.
- Wallace, R. J. 1999. Three conceptions of rational agency. *Ethical Theory and Moral Practice* 2: 217–242.
- Wallace, R. J. 2001. Normativity, commitment, and instrumental reason. *Philosophers' Imprint* 1: 1–26.
- Way, J. 2010. Defending the wide-scope instrumental principle. *Philosophical Studies* 147: 213–33.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wilson, G. 1989. *The Intentionality of Human Action*. Stanford, CA: Stanford University Press.
- Wittgenstein, L. 1953. *Philosophical Investigations*. Translated by G. E. M. Anscombe. Oxford: Blackwell.
- Yablo, S. 1992. Mental causation. *Philosophical Review* 101: 245–280.
- Yablo, S. 1997. Wide causation. *Philosophical Perspectives* 11: 251–281.