


Action, Intention and Knowledge

Lucy Campbell
Homerton College, Cambridge



June 2015

This dissertation is submitted for the degree of
Doctor of Philosophy.

Acknowledgements

For financial support throughout my PhD I would like to thank the University of Cambridge, the Cambridge Philosophy Faculty and an anonymous donor without whose generosity this Thesis would almost certainly not have been completed. I would also like to thank my college, Homerton, who have made money available for me to attend conferences and other academic meetings away from Cambridge; this has been a great help.

Parts of the work in this Thesis, and their antecedents, have been presented in Oxford, Cambridge and King's College London, and I am indebted to audiences in each of these places for their time and attention and the helpful discussions which ensued. I would also like to thank the members of a reading group on action which was held in the Cambridge Philosophy Faculty during 2013, and in particular the stalwarts; Arif Ahmed, Jonathan Birch, Ali Boyle, Alexander Greenberg, Jane Heal and Thomas Land. Thanks also go to the King's College London Philosophy Faculty for preparing me for the PhD, in particular to M.M. McCabe, to Shalom Lappin and to Mark Textor.

For some excellent philosophical conversations, both recent and historical, about issues considered in the Thesis but also other interesting things, I thank Maria Alvarez, Luke Brunning, Mike Campbell, Katherine Cecil, Harry Cook, Chris Cowie, Ryan Cox, Tim Crane, Alex Davies, Fee Doherty, Alison Fernandes, Alexander Greenberg, Nathan Hauthaler, Ben Jeffery, John Maier, M.M. McCabe, Rory O'Connell, Nikita Perepelov, Emily Platten, "The Peters" Ridley and Sutton, Lukas Skiba, Henry Tyrrell, Ellisif Wasmuth, Jen Wright and John D. Wright. Chris Cowie, Ryan Cox, Tim Crane, Alison Fernandes, Alexander Greenberg and John Maier deserve particular thanks for reading and commenting on previous incarnations of various parts of the Thesis.

I am very lucky to have been supervised from October 2011 by Jane Heal, and from October 2013 by Richard Holton. My meetings with Jane and Richard have been both enjoyable and incredibly useful, and the help they have given me far outstretches what I have been able to indicate in the occasional footnote. Jane in particular, who has supervised me throughout my entire PhD, has given me an extraordinary amount of her time, often meeting me once-weekly for months-long stretches. I am particularly indebted to Jane for keeping me on as a supervisee long after her retirement in 2012. She has read and re-read (and re-re-read) several drafts of each chapter as well as several other pieces, remaining at once encouraging and rigorously critical. And she has been a great help during those times when work wasn't going so well and I was in low spirits. Her kindness, understanding and patience have been as important to my finishing this Thesis as has her excellent supervision.

For friendship, fun and moral support, I would like to thank all of my friends, but especially those, old and new, who I have been lucky enough to have been in contact with not-too-seldom over the past four years; Dan Brigham, Chris Cowie, Julia Felder, Alexander Greenberg, Mynn Khine, M.M. McCabe, Peter Ridley, Jenni Sidey, Lukas Skiba, Henry Tyrrell, Isobel Urquhart, Ellisif Wasmuth and Mike Withey, and amongst those, to Alexander, Chris, Julia, Ellisif and Jenni in particular for so often keeping me company in the pub and occasionally elsewhere. Thanks to my parents for their ongoing

support, and most of all to Alexander Greenberg, who has kept me as close to sane as is reasonable to hope for.

Table of Contents

<i>Introduction</i>	i
Part One: Motivating Intentionalism	1
Chapter One: Practical Knowledge	3
Introduction	3
1. Practical Knowledge's Content	4
2. Practical Knowledge's Apparent Features	6
3. The Special Relationship between Practical Knowledge and Intentional Action	11
4. Further Remarks on Practical Knowledge	15
5. Conclusion	18
Chapter Two: Alternatives to Intentionalism	19
Introduction	19
1. Lucy O'Brien's Consciousness-Based Account	20
2. Sarah Paul's Inferentialist Account	32
<i>Concluding Part One</i>	41
Part Two: Cognitivist Intentionalism	43
Chapter Three: Cognitivist Intentionalism	45
Introduction	45
1. Velleman's Internalist Cognitivist Intentionalism	45
2. Setiya's Externalist Cognitivist Intentionalism	56
Chapter Four: The Content of Intention	63
Introduction	63
1. Preliminaries	64
2. An Argument Against Propositionalism	68
3. Three Responses and Conclusion	74
<i>Concluding Part Two</i>	81
Part Three: A Better Intentionalism	83
Chapter Five: Anscombe's Non-Cognitivist Intentionalism	84
Introduction	84
1. Intentional Action, Practical Knowledge, 'Why?'	86
2. Merits and Clarifications	96
3. Three <i>Prima Facie</i> Problems	106
Chapter Six: Non-Propositionalist Intentionalism	111
Introduction	111
1. Non-Propositionalist Intentionalism	114
2. The Why Knowledge? Problem: Clarifications	120
3. Why Practical Knowledge is <i>Knowledge</i>	127
4. Knowledge and Mental Engagement	134
5. Conclusion	141
<i>Conclusion</i>	143
<i>Bibliography</i>	147

Introduction

a. The Thesis to be advanced: Non-Propositionalist Intentionalism

This Thesis recommends a way of thinking about the knowledge we have of our own intentional actions or, as I will call it, *practical knowledge*. On this view, one's knowledge of what one intentionally does is constituted by the intention involved in doing it, on the condition that the intention is executed. Intentions are understood as differing from beliefs both normatively and semantically. They differ normatively in that they set standards on how things are to go or to be going, whereas beliefs have standards set for them by how things are, were or will be. And they differ semantically because whereas beliefs are propositional attitudes, intentions are not, or at least are not fundamentally so.

I call this view *Non-Propositionalist Intentionalism about Practical Knowledge*, or *Non-Propositionalist Intentionalism* – *NPI* – for short. The 'Intentionalism' part of the name indicates a view on which practical knowledge is *constituted by intention*, and the *Non-Propositionalist* part indicates that intentions are here understood as – unlike beliefs – not propositional attitudes. *NPI* is a response to the problems I find in various attempts to understand how practical knowledge could be constituted by a true belief held in epistemically favourable circumstances, and is a development of the view I find in Anscombe's *Intention*.

In addition to its key claim – that practical knowledge is constituted by intentions, on the condition that they are executed – *NPI* involves two other commitments, one epistemological and one action-theoretic. The first is a commitment to understanding knowledge *per se* as neutral between its practical and theoretical manifestations, as a certain kind of *mental engagement* with a fact; a kind apt to underwrite certain specifically epistemic capacities, such as the capacity to express and to make inferences from what one knows. Possessing these capacities *can* be a matter of having a belief which is true and held in epistemically favourable circumstances, but it can also be a matter of having an intention which is being or will be executed. Knowledge manifested in the first way is 'theoretical' and knowledge manifested in the second way is 'practical'.

The second commitment is to a non-reductive conception of the relationship between the phenomena (concepts) of *practical knowledge*, *intentional action* and *intention-execution*. These three phenomena are understood as three distinct yet essentially related facets of a single underlying phenomenon. The concepts which pick out each of these three facets are to be conceptually elucidated in terms of one another. Thinking of someone as φ -ing intentionally, as executing their intention to φ , or as having practical knowledge that they are φ -ing is thinking of them as instantiating a single phenomenon under three different aspects.

I argue that *NPI* is the best way to jointly meet several constraints on a good account of practical knowledge, which will be introduced and explained throughout the Thesis.

b. The structure of the discussion

The Thesis falls into three parts; *NPI* is introduced and defended in Part Three. Let me sketch the argumentative route which will take us there.

Part One introduces our target phenomenon - *practical knowledge* - and motivates accepting some version of *Intentionalism*; the view that practical knowledge is somehow constituted by the agent's intention. In Chapter One practical knowledge is described first in terms of its content and then in terms of three philosophically interesting features which appear to characterise knowledge with this content. An initial constraint - the *Apparent Features* constraint - requires an account of practical knowledge either to *accommodate* or to *explain away* each of these apparent features. Practical knowledge and intentional action are shown to be particularly tightly related, and in light of this a second constraint is added: the *Special Relationship* constraint requires an account of practical knowledge to contain a plausible conception of this relationship.

Chapter Two aims to dismiss two general approaches to thinking about practical knowledge by dismissing representative examples. Lucy O'Brien sees practical knowledge as somehow grounded in the agent's *conscious awareness* of acting, whereas Sarah Paul thinks of practical knowledge as *inferred from* the agent's knowledge of her intention. I suggest that both O'Brien's *consciousness-based* account and Paul's *inferentialist* account flout the *Special Relationship* constraint, albeit in different ways. Paul's account, additionally, flouts the *Apparent Features* constraint and a new constraint we are led to by the discussion: the *Why Knowledge?* constraint requires an account of practical knowledge to make it clear why the phenomenon it describes achieves the status of *knowledge*. I try to show that it is the core respective ideas in O'Brien's and Paul's pictures which cause problems for their views; O'Brien's problems stem from thinking of practical knowledge as consciousness-based, and Paul's from thinking of it as inferentially grounded. These two broad approaches are the key alternatives to *Intentionalism*, so rejecting them motivates considering the latter.

Part Two considers and rejects a version of *Intentionalism* I call *Cognitivist Intentionalism* - *CI*. The 'Cognitivism' included in *CI* is the view that intentions are a kind of belief. If we think of 'ordinary' knowledge as constituted by true beliefs held in epistemically favourable circumstances, then we can think of 'practical' knowledge as a special case, one in which the true belief is an intention. This basic structure is consistent with different accounts of what the 'epistemically favourable circumstances' are, in which a true *intention*-belief gets elevated to knowledge-status. I consider an internalist and an externalist version of *CI* in Chapter Three. The internalist version comes from David Velleman and is found lacking in similar ways to Sarah Paul's inferentialist view. I explain that this should come as no surprise because despite very different accounts of practical knowledge's psychology, Paul's and Velleman's *epistemologies* of practical knowledge are very similar. The externalist version - due to Kieran Setiya - overcomes these problems, but does not obviously meet the *Why Knowledge?* constraint, and seems to rely on a suspect claim; that we can form beliefs at will.

Chapter Four argues for a fourth constraint on a good account of practical knowledge. The *Intention-Content* constraint rules out any account of practical knowledge which requires seeing intentions as in all cases propositional attitudes. I assume that *some intentions are* propositional attitudes, and argue that these can be executed only by executing a non-propositional intention; specifically one whose content does not contain

reference to the intender, an intention simply *to* φ . No version of *CI* can meet the *Intention-Content* constraint: if intentions are a kind of belief then they *must* be propositional attitudes in all cases. We thus have a reason to reject *CI* which is independent of problems with the details of any particular version of it. If practical knowledge *is* somehow constituted by intention, it cannot be because those intentions are a kind of belief.

Part Three introduces, develops and defends a better *Intentionalism*. Chapter Five considers a view which I find in Anscombe's *Intention* and label *Non-Cognitivist Intentionalism* – *NCI*. *NCI* holds that practical knowledge is constituted by intentions *rather than* beliefs. Intentions and beliefs differ *normatively*: intentions set standards on how things are to go whereas beliefs have standards set for them by the facts. Because practical and theoretical knowledge are constituted by attitudes of different normative shapes, they themselves differ *formally* in two important ways: practical knowledge is undermined by a mistaken performance; theoretical knowledge by a mistaken judgment, and practical knowledge is 'the cause of what it understands', whereas theoretical knowledge is 'derived from the objects known'. Because it sees practical knowledge as having these formal properties, *NCI* happily meets the *Apparent Features* and *Special Relationship* constraints, and does so in a unified way. Our discussion also leads us to introduce a new constraint. The *Structure-Sensitivity* constraint requires a good account of practical knowledge to be sensitive to intentional action's internal structure, which I explain. *NCI* is shown to meet this constraint very easily too.

But *NCI* faces two serious problems. First, it is unclear whether it meets the *Why Knowledge?* constraint. For *NCI* holds that practical knowledge is constituted by intention, which is *not* a kind of belief, just in case the intention is executed. But it is hard to see how an attitude other than belief could constitute knowledge, and hard to see in particular how *executing an intention* could be an *epistemic*-phenomenon rather than an *agential* one. Second, it is hard to see how *NCI* could meet the *Intention-Content* constraint consistently with a new and final constraint. This – the *Knowledge-Content* constraint – requires an account not to leave it mysterious why practical knowledge is a kind of *propositional* knowledge. Indeed, it might seem that *any Intentionalism* will flout either the *Intention-Content* or the *Knowledge-Content* constraint, depending on whether it sees intentions as propositional or as non-propositional attitudes: how could a non-propositional attitude constitute a kind of *propositional* knowledge?

In Chapter Six I consider a version of *Intentionalism* which contains a conception of intention as a *non-propositional* attitude. *Non-Propositionalist Intentionalism* – *NPI* – trivially meets the *Intention-Content* constraint, and since it *entails* *NCI*, also takes over *NCI*'s capacity to meet the *Apparent Features*, *Special Relationship* and *Structure-Sensitivity* constraints, as argued in Chapter Five. I explain how *NPI* can be developed in such a way that it meets the *Why Knowledge?* and *Knowledge-Content* constraints too. It meets the latter because its core claim is that intentions constitute practical knowledge on the condition that they are executed, and because the execution of the intention brings into being something with propositional shape, *viz.* the fact that one is φ -ing. This fact, according to *NPI*, is the object of the agent's practical knowledge. And it meets the former because someone who is executing their intention possesses a set of capacities which commonsense epistemology links to *knowing that* p , and because we need not think of them as believing that they are acting in order to explain how they have these capacities; we need only think of them as executing their intention. The upshot is that executing an intention can be thought of as an epistemic phenomenon *as well as* an agential one: it

bears all the marks of knowing. Thus I will have made the case for *NPI* as an account of practical knowledge able to meet all six of our constraints, and so an account to take seriously.

c. Our Commonsense-Human methodology

My approach to thinking about practical knowledge is guided by commonsense psychological, epistemological and action-theoretic reflections, and not by discoveries in empirical science or hypotheses from empirically-minded philosophy of mind or theoretical epistemology. My starting assumption is that our commonsense-psychological, -epistemological and -action-theoretic concepts are generally in good working order. Because I think these commonsense concept-types (epistemological, psychological etc.) are inextricably linked with one another, it is useful to have a term for the whole web of them, and I will refer to them collectively as our *commonsense-human concepts*. I will refer to the methodology whose starting-point is a respect for the good working order of these concepts as a *commonsense-human methodology*.

The methodology does not assume that we cannot make or discover *mistakes* within our commonsense-human modes of thought, but it does mean that mistakes can be exposed *as* mistakes only by exposing the ways in which they fail to mesh with other parts of our commonsense-human understanding. One upshot of this is that *most* of the concepts, platitudes and assumptions of commonsense must be held fixed as the background against which to consider any given part of it. Another is that any account of a commonsense-human phenomenon which entails that our commonsense-human understanding of it is systematically mistaken, must provide an error-theory for these mistakes, itself given in commonsense-human terms.

My adherence to this commonsense-human methodology is motivated by the thought that our *understanding* of the kind of creatures we are is essentially bound up with our *being* the kind of creatures we are. From this perspective, eliminativist or revisionist conceptions of commonsense-human concepts are incoherent, since they presuppose that our characteristically human modes of understanding, behaving and communicating could be abandoned, whilst tacitly assuming that *we* – i.e. *this kind of creature* – would survive the abandonment. So I think that taking our commonsense-human concepts seriously is unavoidable, not only psychologically, but logically too. But I won't argue for these points; I mention them by way of explaining why I commit myself to the commonsense-human methodology I am describing.

But I think the methodology can be justified more prosaically too. For our *quarry* – practical knowledge – is itself a commonsense-human phenomenon. My aim in this Thesis is to understand something which we all recognise from our experience as humans, specifically as agents and as knowers, although most of us may not have a name for it. Because our target phenomenon is a commonsense-human one, we should not expect a methodology *other than* the kind I have been describing to deliver an account of it.

d. Why care about practical knowledge?

Understanding practical knowledge is important for various reasons. I will consider just a few.

First, some have thought that understanding our knowledge of our own intentional actions might provide the basis for an understanding of how we know about our own intentional mental states. This can be developed in various ways, but the common thread is to see our relation to these states as implicating our rational agency in some way.¹ The prospects for a view like this aren't my focus here, but clearly understanding practical knowledge is a prerequisite for such a view.

Second, there is the thought that the actions which are the objects of practical knowledge form a morally and legally relevant class. Denying that one *knew* that one was acting is a way of denying *one kind* of moral and/or legal responsibility. (I am careful not to make the (false) claim that there is *no* sense in which we can be responsible for those things we do unknowingly.) If I poisoned the well unknowingly, because the fertilisers I used leached into the groundwater, then I am in a very different moral and legal situation to someone who did so intentionally and knowingly, in order to kill off the villagers. Again, understanding practical knowledge enables more to be said about this special kind of moral and legal responsibility. These issues are not my focus, but they will come up here and there during my discussions.

Third, as we shall see in Chapter One, some have thought that practical knowledge's philosophically interesting features follow from the fact that acting intentionally is in some important sense *dependent* on having practical knowledge. But ordinarily, the facts we know are *independent* of our knowing them. The pangolin would still be endangered even if I didn't know that it was. If practical knowledge does not fit this model, then a rather fundamental feature of our conception of knowledge will have to be re-thought. This is a central thought in Anscombe's *Intention* and I will argue for it and develop it in Chapters Five and Six.

Finally, the idea that acting intentionally is somehow dependent on one's knowledge has action-theoretic implications. For if the idea is right, then in order to understand either of these phenomena – intentional action and practical knowledge – we need to understand them both, and to understand their interrelations. This idea too is central to Anscombe's discussion, and is another which I will develop in Chapters Five and Six.

¹ See e.g. Moran, "Making Up Your Mind"; *Authority and Estrangement* and "Self-Knowledge, 'Transparency', and the Forms of Activity"; O'Brien, "Self-Knowledge, Agency, and Force" and *Self-Knowing Agents*; Peacocke, "Mental Action and Self-Awareness (I)" and "Mental Action and Self-Awareness (II)" and Boyle, "Active Belief" and "'Making Up Your Mind' and the Activity of Reason."

Part One

Motivating Intentionalism

Part One introduces the phenomenon of practical knowledge, develops three initial constraints on an adequate account of it, and motivates taking some version of *Intentionalism* seriously. Chapter One provides an initial sketch of practical knowledge, whereas Chapter Two considers and ultimately rejects two alternatives to *Intentionalism*.

Chapter One:

Practical Knowledge

Introduction

I am currently typing, typing a sentence, and writing a PhD Thesis. I know that I am doing all these things, but how? What is involved in having this knowledge?

Many philosophers have thought that one's knowledge of one's own intentional action is knowledge of a special kind. *Practical knowledge*, as I will call it, appears to have certain philosophically interesting features, features which do not attach to one's knowledge of one's own *non-intentional* actions, or to one's knowledge of the intentional actions of others. This Thesis considers how best to understand practical knowledge, in order to do justice to its apparent possession of these features.

In this introductory Chapter I define a starting-point for this investigation by providing a sketch of the phenomenon of practical knowledge. The sketch must be sparse enough to be more or less neutral between the various accounts of practical knowledge I will consider later on. But it must be full enough for it to be clear *which phenomenon* the expression 'practical knowledge' picks out, and to give an initial sense of what a good account of the phenomenon must explain.

I start in §1 by describing practical knowledge in terms of its content; as knowledge of one's own current or ongoing intentional actions. But there is more to practical knowledge than this; practical knowledge appears to be *first-person authoritative*, *epistemically ungrounded* and *psychologically immediate*, features which I explain in §2. An initial constraint on an account of practical knowledge – the *Apparent Features* constraint – requires an account to either accommodate, or explain away these features.

That knowledge with a certain *content* should display unusual epistemological and psychological features ought to strike us as curious, as I explain in §3. So a second constraint on an account of practical knowledge – the *Special Relationship* constraint – requires it to explain why the special features described in §2 should attach to knowledge with the content described in §1, and to contain a plausible conception of the relationship between intentionally doing something and having practical knowledge that one is doing it. I give some reasons for and against considering that relationship a *conceptual* one, without drawing any conclusions; I will return to this issue in later chapters.

In §4 I introduce some further possibilities for thinking about practical knowledge, possibilities we will consider more decisively in later chapters. We will see that some have been attracted to the idea that practical knowledge can take in *future* as well as *present* intentional actions, and that some have thought that practical knowledge is knowledge not only of one's intentional actions themselves, but also of the *reasons* for which one acts. Again, I make no decisions on these issues, but will come back to them; in particular in Chapter Five. I conclude in §5 by re-capping our two initial constraints on an account of practical knowledge.

1. Practical Knowledge's Content

A preliminary but nevertheless helpful characterisation of our target phenomenon is as follows:

Practical knowledge is the knowledge we have of our own actions when those actions are intentional.

This slogan will form the basis of our sketch of practical knowledge. Drawing out some of the slogan's implications will help us distinguish our target phenomenon from other phenomena which are sometimes given the label 'practical knowledge'.

'Practical knowledge' is sometimes used to refer to knowledge *how* to do something, understood as a certain kind of capacity or ability for acting in a certain way, and often distinguished, following Ryle, from 'knowledge-that', or propositional knowledge.¹ Practical knowledge in our sense is not Rylean know-how, but a kind of *propositional* knowledge. I have practical knowledge *that* I am typing, and although it is very plausible that I could not know *that* I am typing unless I knew *how* (in Ryle's sense) to type, my practical knowledge is not itself knowledge-how, but knowledge-that.²

The distinction between practical knowledge and knowledge-how survives even when Ryle's distinction between know-how and know-that is collapsed. Some have thought that my knowing how to type is my knowing *that*, of some way of acting *w*, *w* is a way of typing (or a way for me to type).³ But my practical knowledge (in our sense) is different from my knowledge how to type, *even if* construed as a kind of knowledge-that. For my current practical knowledge that I am typing is *existential* in a way that propositionally construed knowledge-how is not; it *entails the current existence of an act of typing*, whereas my knowing *how* to type, however understood, does not.

Finally, sometimes moral and prudential knowledge are referred to as 'practical' knowledge. But this knowledge is not existential either. My knowing that it is wrong to go about murdering people (moral knowledge) doesn't entail that I'm doing any murdering, nor does my knowing that poisoning is a good (*i.e.* relatively covert) method of murder.

Our slogan characterises the object of practical knowledge as the agent's *own intentional* actions. That is not to say we can't know about others' actions, or about our own non-intentional ones. Obviously we can. What it means is that these kinds of knowledge are not examples of the phenomenon this Thesis seeks to understand. And although practical knowledge is always knowledge *of intentional action*, this does not mean that it must have contents of the form "I am φ -ing *intentionally*". What it means is that when someone's knowledge that they are φ -ing is *practical knowledge*, they *are* φ -ing intentionally. We will begin to see why this is in §3, when we start to consider the special relationship

¹ Ryle, *The Concept of Mind*. Ch. 2.

² Throughout I will talk of knowing *that* one is φ -ing interchangeably with knowing *what* one is doing. Identifying *knowing what* with a kind of *knowing that* is controversial, but my choice of usage does not commit me to this identification, it serves only to simplify discussion. The phenomenon I am interested in is knowledge of the kind I express in saying "I am φ -ing", and which is properly referred to as my knowledge *that I am φ -ing*. I will also sometimes refer to this phenomenon as ' φ -ing knowingly.' In all cases, the phenomenon I have in mind is a kind of first-person propositional knowledge *that one is φ -ing*.

³ See e.g. Stanley and Williamson, "Knowing How."

which seems to exist between practical knowledge and intentional action. For now, I will speak as if practical knowledge is always knowledge of *current* or *ongoing* intentional action; as if its content is always of the form “I am φ -ing”, but I will come back to the question of practical knowledge and tense in §4(a).

Before moving on I want to make two further comments about the idea that practical knowledge is knowledge of specifically *intentional* action. First, actions are not intentional *per se*, but relative to a description. I am currently typing, but I am also making tapping noises as I hit the keys, and whilst I am *typing* intentionally, my *making tapping noises* is not intentional. It is plausible to think that I am performing one action here which can be described in two different ways; this is the standard view and I will assume it here.⁴ As we might expect given that practical knowledge is particular to intentional action, practical knowledge is also description-relative: whilst I have practical knowledge *that I am typing*, I don’t have practical knowledge *that I am making tapping noises*, although I may well have *some* knowledge of this (the difference between practical knowledge and non-practical knowledge of action will be explored in §2).

Second, I should be clear that in thinking of practical knowledge as knowledge of specifically *intentional* action, I am glossing over a complexity within the literature. For although many have sought to understand the special kind of knowledge we have of a privileged class of our own actions, not everyone has identified this privileged class with our *intentional* actions. Others have thought that we have a special kind of knowledge of our *voluntary*,⁵ or our *full-blooded*,⁶ or our *autonomous*,⁷ or our *deliberate*⁸ actions. Plausibly, these are all slightly different notions, each picking out a class of action which is privileged in its own characteristic way.⁹ But they do seem to have something in common; they all seem to pick out actions which are related in a privileged way to their agents’ *mind, rationality* or *personhood*.¹⁰

Why choose to identify the privileged class as that of our *intentional* actions? There are two reasons. First, my study is located within a tradition – begun in Anscombe’s *Intention* – in which the privileged class *is* most commonly identified with *intentional* action. And second, my focus in this Thesis is on a certain *kind* of account of practical knowledge; an account on which practical knowledge is in some sense constituted by the agent’s *intention*. So it is natural to settle on the class of *intentional* actions as those which are taken in by practical knowledge.¹¹

⁴ See e.g. Anscombe, *Intention*, 11–12; and “Under a Description”; Davidson, “Agency,” 59; and “Problems in the Explanation of Action,” 104. Ursula Coope suggests this view goes back at least to Aristotle (Coope, “Aristotle on Action,” 109.). Cf. Goldman, *A Theory of Human Action*, 1–10. There is philosophy to be done in explaining when two descriptions attach to the same action, and when they pick out distinct actions, as they sometimes do, but I won’t be doing it here.

⁵ Hampshire, *Freedom of the Individual*, 54.

⁶ Velleman, *Practical Reflection*, 5.

⁷ Velleman, “Precis of ‘The Possibility of Practical Reason,’” 4.

⁸ O’Brien, *Self-Knowing Agents*, 161.

⁹ For a discussion of how these kind of adverb might be distinguished, see Hyman, *Action, Knowledge and Will*. Ch. 1.

¹⁰ Velleman, *The Possibility of Practical Reason*, 4.

¹¹ Not everyone agrees on this. Velleman and Hampshire, for example, think that practical knowledge is *constituted by* intention, but is *knowledge of* full-blooded or autonomous action (Velleman) or voluntary action (Hampshire).

2. Practical Knowledge's Apparent Features

We have just identified our target phenomenon in terms of its content. But why should knowledge *with a certain content* be worthy of philosophical investigation? We don't care about the distinction between meteorological and non-meteorological knowledge, or about people's knowledge about their own shoes as compared to their knowledge of the shoes of others. So why is knowledge about *our own intentional actions* interesting in a way that knowledge about our own non-intentional actions, or about the intentional actions of others, is not? The answer is that this knowledge seems to have some curious features. I will distinguish three, explaining each in turn.

a. First-Person Authority

First person authority is a property usually associated with psychological self-knowledge; knowledge of one's own intentional and sensory states. In this connection, the expression 'first person authority' – henceforth 'FPA' – has various connotations; of a person's having *privileged epistemic access* to her mental states, of her being *the best person to ask* about her own mental goings-on, of her statements about her own mind being *presumed true*, *if presumed sincere*, and so on. And there is disagreement over exactly how to understand FPA; over which of these and other factors are relevant. To separate out the various properties associated with FPA would take us too far afield, but luckily we need not do this in order to make use of the notion in our understanding of practical knowledge.

Instead we can start by understanding practical knowledge's apparent FPA rather loosely; as involving some kind of *first-/third-person asymmetry*. There is some kind of difference between a person's capacity to know what she is intentionally doing and anyone else's capacity to know it. This does not mean that we are *infallible* about our own intentional actions. I can go wrong at least in the sense that I might think I am intentionally doing something which I am not – as when I think I'm kicking Rory under the table but I am in fact kicking Jen. What it means is that when I *do* have practical knowledge that I am φ -ing, my epistemic relation to the fact that I am φ -ing is in some sense of a different order than yours could be.

I leave open here exactly what this sense is. All of those whose accounts of practical knowledge we will consider in the coming chapters – our protagonists, as I will call them – agree that practical knowledge displays FPA. A good account of practical knowledge will contain a plausible explanation of what this amounts to. I will come back to this in (d).

b. Epistemic Ungroundedness

Practical knowledge also appears not to be epistemically grounded in any of the usual ways in which knowledge of contingent facts is grounded. The 'usual ways' are memory, testimony, perception and inference, including inference from general knowledge and from any of the other three kinds of ground.

We can rule out memory and testimony as grounding practical knowledge very easily. Memory can furnish us only with knowledge of the past, or function to maintain knowledge previously learned. But practical knowledge is never knowledge of *past*

intentional action. Testimony often grounds our knowledge of others' intentional actions, as when Henry tells me on the phone that he's starting a new painting; or our knowledge of our own *non*-intentional ones, as when later, whilst visiting his studio, Henry tells me that I'm standing on a tube of paint. But Henry doesn't appear able to inform me that I am – intentionally – wiping the paint off my shoe. If I am wiping off the paint intentionally, I don't seem to need someone else to tell me I am. And my practical knowledge can't be based on *my own* testimony, on pain of circularity.

Sometimes we know about others' intentional actions and our own non-intentional ones by perception or observation.¹² (I will use these expressions interchangeably throughout the discussion of this Thesis.¹³) But I don't seem to know about my own intentional actions in this way either.¹⁴ Consider my knowledge that I am typing these words. Yes, I can see and feel my fingers moving, hear and feel them hitting the keys and see these words appearing on the screen, but my knowledge that I am typing these words doesn't seem *based on* these observations. After all, if I didn't know what I was typing until I saw the words appear on the screen, it would come as a surprise to me that I ended up writing just these words. And not only is it *not a surprise*, but it's hard to see how I could *be* intentionally writing a sentence if I had to rely on perception to tell me what I was writing. Having written "Not only is it *not a surprise*," for example, I would not yet know which sentence these words were the beginning of. In such a scenario it is hard to see how writing a sentence could be any less than a miracle.

A further reason to doubt that perception grounds practical knowledge is illustrated nicely by Kevin Falvey:

Suppose [...] that I have a nosy neighbor who spies on me using a telescope trained on my apartment. If I am sitting on the couch one afternoon reading a book, she would be in a position to see that I am doing so, but she would not be in a position to see, what might be perfectly true [...], that I am making bread.¹⁵

Whilst sitting on the couch, Falvey knows that he is making bread in the absence of *any* current relevant perceptible clues. Nor does his knowledge seem to be based on *earlier* clues which his neighbour wasn't party to. To see this, assume that his nosy neighbour has been watching him all day, has seen him kneading the dough and putting it on the radiator to rise. It is still possible for Falvey to have knowledge which his neighbour lacks; for adjust the example so that Falvey is *not* making bread at all, not really; instead he's experimenting with a new low-yeast recipe he's found (he doubts it will yield a dough that rises properly). Things *look* just as they did in the previous case, and the neighbour judges as before, that Falvey is making bread. But he's not – he's doing an experiment. And *Falvey* knows this despite the fact that the same (or: relevantly similar) *perceptual* information is available to him as to his neighbour.

¹² A piece of knowledge might be *grounded in perception* without being *inferred from perceptual beliefs* (see Roessler, "Intentional Action and Self-Awareness"; Grünbaum, "Perception and Non-Inferential Knowledge of Action."). I consider an inferential conception of practical knowledge below.

¹³ This is standard in the literature, although not ubiquitous; John Schwenkler and Johannes Roessler both think that although practical knowledge is *perceptually* based without being *observationally* based (Schwenkler, "Perception and Practical Knowledge"; Roessler, "Intentional Action and Self-Awareness."). Both Schwenkler and Roessler make a strong case for practical and perceptual knowledge being *intertwined* in interesting ways, but I don't think either makes a convincing case for the claim that practical knowledge is epistemically *grounded* on perception of what one is doing, as that claim is understood here. I won't have time to consider their discussions.

¹⁴ Cf. Proust, "Perceiving Intentions."

¹⁵ Falvey, "Knowledge in Intention," 25.

It might be pointed out here that actually Falvey has access to perceptual clues which his neighbour lacks, *viz.* the deliverances of proprioception.¹⁶ But it is implausible that Falvey's practical knowledge is based on proprioception. For in the initial example he knows that he *is making bread* (and not doing an experiment) and in the re-jigged version he knows that he is *doing an experiment* (and not making bread). Yet we can stipulate that his bodily movements, and so his proprioceptive information, up to and including his sitting on the couch, are just the same. Sure, his bodily movements will differ between the cases later on (in one case he takes the dough to the oven, in the other he takes it to the bin), but this is irrelevant to the knowledge he has *whilst sitting on the couch* of what he is doing.

It is important to be clear about what is being denied when we say that practical knowledge is not perceptually grounded. The idea is that it is not justified by reference to the agent's perception of her action. This does not rule out that perception is a *precondition* for having practical knowledge. It seems very plausible that certain intentional actions cannot be performed in the absence of any sensory feedback. Even rather simple things like pointing to the right seem hard to imagine being able to do in *complete* sensory (including proprioceptive) deprivation,¹⁷ and more complex and world-involving actions like cycling to a friend's house, playing darts or building a tree-house would be out of the question.

But that the capacity for perception is presupposed by the capacity for practical knowledge does not mean that practical knowledge is epistemically grounded in or justified by perception. Our judgments about what we are intentionally doing do not seem to be based on perceptual information, in the way that our judgments about what others are doing might be so based. Asked how I know that Jess is cycling, I might say that I can see that she is. Asked how I know that *I* am cycling, I won't say that I can *see* that *I* am – or that I can *feel* that I am, for that matter, either.

Finally, we sometimes come to know about others' intentional actions on the basis of inferences from various bits of general knowledge about their tendencies or capacities, or else on the basis of inferences from perceptual clues, testimony or both of these combined with general knowledge. I can hear a tapping coming from Henry's room, and I know that he meant to get on with some work today, which would involve typing and so tapping. I infer that he's typing, and my judgment might constitute knowledge in virtue (*inter alia*) of being so grounded. But although in my own case I can hear a tapping noise, and I know that *I too* meant to get on with some work today, which would involve typing, so tapping, it isn't *prima facie* plausible that my knowledge that I am typing is based on these kinds of consideration.

It might be agreed that my knowledge that I am typing is not based on the same kind of premises as my knowledge that Henry is, yet be maintained that my knowledge is nonetheless inferentially grounded. Some have thought that my knowledge that I am typing is inferred from (*inter alia*) my knowledge of my own mind, specifically, my knowledge of my own *intention*. I will consider a view like this (from Sarah Paul) in more detail in Chapter Two. The important point here is that this isn't how things *appear*. For in general when someone has inferentially grounded knowledge, she will be able to answer a

¹⁶ For a proprioceptive conception of practical knowledge see Dokic, "The Sense of Ownership: An Analogy between Sensation and Action."

¹⁷ Cf. Peacocke, "Mental Action and Self-Awareness (I)," 359.

“How do you know?” question by citing one or more of the inference’s premises, yet asked “How do you know you’re typing?” I would *not* naturally respond by saying that I intend to be doing so.

Practical knowledge seems unable to be epistemically grounded in memory or testimony. And it does not appear to be grounded in perception or in inference. And this is so despite the fact that what is known in practical knowledge – facts about a person’s ongoing intentional actions – *are knowable* in the usual ways. For *someone else* can know that I am typing *only* on the basis of testimony (mine or someone else’s), perception or inference. And *I* am just another *she* from *her* perspective.¹⁸

c. Psychological Immediacy

Practical knowledge seems to have a further philosophically interesting feature. We are often said to be able to say *straight-off* what we are intentionally doing, where ‘straight-off’ means something like ‘without having to think about what one is doing’.¹⁹ I will refer to this feature as ‘psychological immediacy’.

Initially this might not seem to distinguish practical knowledge from knowledge of others’ intentional actions, or of one’s own non-intentional ones, for often one does not need to do any thinking prior to saying what someone else is doing intentionally, or what one is non-intentionally doing. If asked what you are doing I might be able to reply straight-off “She’s putting salt on her chips”, and whilst accidentally knocking a stack of plates out of the cupboard I might be asked “Who’s making that racket?”, responding ‘straight-off’ “I am!” But practical knowledge does differ from these kinds of case in two important ways.

First, whereas there *are* examples in which I can say ‘straight-off’ – without consideration – what someone else is intentionally doing, or what I am non-intentionally doing, I needn’t be able to. If I am not wearing my contact lenses and am asked at lunch what you are doing I might have to peer right in at you in order to be able to say that you are putting salt on your chips. And if asked, whilst recounting the amusing tale of the time Dan forgot his passport and missed his plane to Oslo, “Do you think you might be upsetting Dan?” I may have to consider things a little before coming out with “Oh dear, I think I am.”

To see the second difference we need to say more about *psychological immediacy*. Compare my being able to say straight-off that you are (intentionally) putting salt on your chips and my being able to say straight-off that *I* am (intentionally) putting salt on *my* chips. My capacity to say straight-off what you are doing depends on my having previously *noticed* or otherwise discovered that you are salting your chips. But my being able to say straight-off that *I* am salting *my* chips doesn’t seem to require any prior noticing or discovery. The same difference holds between my capacity to say straight-off what I am doing intentionally and what I am doing non-intentionally. Say I am non-intentionally writing the words ‘casual relationship’ (I meant to be writing ‘causal relationship’). I might be able to say straight-off that I am doing so once I notice my mistake, but if I am writing

¹⁸ See also O’Brien, *Self-Knowing Agents*, 156.

¹⁹ See e.g. Paul, “How We Know What We’re Doing,” 10; Falvey, “Knowledge in Intention,” 22; O’Brien, *Self-Knowing Agents*, 168–171.

‘casual relationship’ intentionally, I can say that I am doing so straight-off *without* having to *notice* or discover that I am.²⁰

‘Noticing’ and ‘discovering’ are not the only relevant concepts here. Sometimes my knowledge of what you are intentionally doing or of what I am doing non-intentionally will follow on from *finding*, *working* or *figuring out* what’s being done. Part of the idea behind the seeming psychological immediacy of practical knowledge is that having practical knowledge of what one is doing seems inconsistent with *finding*, *working* or *figuring out* what one is doing.

Practical knowledge’s apparent psychological immediacy may seem to go hand-in-hand with its apparent epistemic ungroundedness.²¹ Discovering (etc.) that something is the case seems to require discovering *via some epistemic route* or other, but we said above that practical knowledge does not seem to be epistemically grounded in any of the usual ways. But psychological immediacy and epistemic ungroundedness *do* seem able to come apart. To bring this out consider my (non-practical) knowledge that my eyes are open. It is at least arguable that the basis on which I know this is perceptual, that it is in some sense based either on the perceptions which I have given that they are open, or on proprioceptive awareness, or a combination of the two. So it is at least arguable that my knowledge that my eyes are open is *epistemically grounded*, in part on perception. But it doesn’t look like something I have *discovered* or *noticed* (etc.) to be the case.

The example might be controversial. One complaint I can imagine is on Wittgensteinian grounds; *that one’s eyes are open* doesn’t look like *knowledge* in the same way that (e.g.) *that there’s a chaffinch in the sycamore* can be knowledge. But I don’t need this example to do very much work. The only point is that it is *not obvious* that practical knowledge’s apparent *psychological immediacy* go hand-in-hand, so that it is worth thinking of these as two distinct apparent properties of practical knowledge. That these properties are distinct will be important in our discussions of Sarah Paul’s and in particular of David Velleman’s conceptions of practical knowledge in Chapters Two and Three.

d. The Apparent Features constraint

We have seen that practical knowledge at least *appears* to be *first-person authoritative*, *epistemically ungrounded* and *psychologically immediate*. Having made these observations, we can state an initial constraint on an adequate account of practical knowledge:

The Apparent Features Constraint

An adequate account of practical knowledge will either accommodate, or explain away, all three of practical knowledge’s apparent features.

Before moving on, a few comments about the *Apparent Features* constraint are in order.

First, the *Apparent Features* constraint does not require a given account to either accommodate (treat as veridical) all of practical knowledge’s apparent features or to

²⁰ See e.g. Velleman, *Practical Reflection*, 47; Anscombe, *Intention*, 51; O’Brien, *Self-Knowing Agents*.

²¹ The idea that we can say ‘straight-off’ what we are intentionally doing is often treated as interchangeable with the idea that the knowledge doesn’t seem to rely on observation or evidence (e.g. Falvey, “Knowledge in Intention,” 22–23.). I have no complaint with using ‘straight-off’ in this way, but it is not the way I will be using it.

explain them all away. It requires that for each feature, a given account must *either* accommodate it, *or* explain it away. Most of the accounts we will consider seek to accommodate some of the features and explain others away.

Second, not all of practical knowledge's apparent features are disputed by our protagonists. All of our protagonists seek to *accommodate* the appearance that practical knowledge displays *FPA*, although they understand this feature differently, and not all as plausibly as each other. Accommodating practical knowledge's apparent *FPA* will involve giving a plausible characterisation of what it involves. Our protagonists also all agree that practical knowledge is *psychologically immediate*, although again, they give different accounts of why this is.

It is practical knowledge's apparent *epistemic ungroundedness* which attracts the most argument, although even here there is a lot of agreement. All of our protagonists accept that practical knowledge is not based on memory, testimony or perception. The disagreements are as follows: O'Brien, Velleman and Setiya think that practical knowledge *is* epistemically grounded, but not in any of the ways discussed in (b). Sarah Paul, as I have already mentioned, thinks that practical knowledge is inferentially grounded, despite appearing not to be. And Anscombe thinks that practical knowledge is not epistemically grounded at all, an approach I develop myself in Chapter Six.

Finally, opting to *explain away* any of practical knowledge's apparent features brings with it an explanatory demand, stemming from our commonsense-human methodology as explained in my introduction (p. iv). If an account suggests that we are systematically subject to misleading appearances about practical knowledge, it also must say something – in commonsense-human terms – about why this should be; about why *this particular* region of our commonsense-human understanding should lead us astray. In the end I will try to describe a conception of practical knowledge (heavily influenced by what I take to be Anscombe's) which accepts all of its apparent features at face value.

3. *The Special Relationship between Practical Knowledge and Intentional Action*

We started to describe practical knowledge in §1 as knowledge with a certain content, and in §2 we saw why knowledge with this content seems worthy of philosophical enquiry – it seems to display certain philosophically interesting features. Here we will consider the relationship between practical knowledge and its object; intentional action. We will see that this relationship looks different to the relationship we usually expect between a piece of knowledge and its object. In (a) I will explain how the relationship between doing something intentionally and having practical knowledge (understood as in §1 and §2) seems an especially tight one, for doing something intentionally at least typically or centrally goes hand-in-hand with having practical knowledge. In (b) I will consider some options for thinking about the nature of this especially tight relationship, without aiming to settle on one. A good account of practical knowledge *should* settle on one, though, and in (c) this demand will be recorded in a second constraint on an account of practical knowledge, the *Special Relationship* constraint.

a. The Special Relationship

Consider two cases in which I express knowledge *that I am chasing the dog*. In Case 1 I am playing with my dog Mutton in the garden. Unable to see Mutton (who's hiding behind a bush), Mum shouts out of the window, "What are you doing?" and I shout back "I'm chasing the dog!" In Case 2, the dog and its temperament are unknown to me. Whilst walking in the park I am ambushed. Terrified, I flee; running in circles trying to shake it off. At some point a shift in aspect takes place; the dog's barks, I realise, are not fearsome at all, but playful; I am, and have been for a while, closer to the dog's tail than it is to mine (despite its doubtless being capable of running much faster than me). And I notice an observer - who I take to be the dog's owner - standing nearby, unflapped and smirking. At the same time, I realise my true situation and exclaim, "I'm chasing the dog!"

Ex hypothesi, at the moment of exclamation in Case 1, I am chasing the dog intentionally whereas at the moment of exclamation in Case 2 I am not. In Case 2 there are, of course, *other* things I am intentionally doing, *running* being one, but I am not intentionally *chasing the dog*. After my realisation I may *decide to continue* chasing the dog, in which case I will come to be chasing it intentionally, but I was not chasing it intentionally at the moment of realisation. Whilst the realisation was dawning on me, I was trying to *get away* from the dog and chasing something is precisely *not* a kind of trying-to-get-away-from.

Compare my *knowledge* in the two cases. Plausibly in Case 1 my knowledge has all three of practical knowledge's apparent features, whereas in Case 2 it has none. But notice further that it's hard to imagine me intentionally chasing the dog in Case 1 *without* having knowledge with these apparent features. It doesn't seem possible to change Case 1 to one in which my knowledge does *not* appear to display *FPA*, *epistemic ungroundedness* or *psychological immediacy* whilst holding fixed that I am chasing the dog intentionally: not only is it the case that the knowledge I have of what I am doing intentionally has all the apparent features attaching to practical knowledge, but also my action's being intentional seems somehow *bound up with* its being an object of knowledge of this special kind. The relationship between intentional action and practical knowledge seems especially tight.

Further considerations support the idea that the relationship between acting intentionally and knowing what one is doing is not only especially tight, but perhaps has a *conceptual* aspect. It is familiar that *denying knowledge* of an action is a way of *denying responsibility*: "How dare you upset Dan like that?" - - - "But I didn't know I was!" Denying knowledge that I was upsetting Dan may not absolve myself of *all* responsibility; Dan's defender may legitimately chastise me for being careless ("Come on, you know he's sensitive about the whole Oslo thing"), but that doesn't mean that my lack of knowledge is irrelevant to the level or kind of responsibility I bear for upsetting him. The responsibility in question seems to be of a *kind* which attaches to specifically *intentional* actions. That denying *knowledge* of one's action is absolving oneself of responsibility of this particular kind suggests some conceptual connection between knowing what one is doing and acting intentionally.

And I also seem able to absolve myself of the special kind of responsibility which goes along with acting intentionally by *admitting* knowledge but characterising it as *lacking* practical knowledge's (apparent) special features. For say I respond to "How dare you upset Dan like that?" with "I know, I feel awful - I realised I was upsetting him when I saw that look on his face, but I couldn't stop telling the story without everyone guessing

how he felt - and that would have made things even worse.” Here I characterise myself as having *noticed* that I am upsetting Dan, and my knowledge as at least in part perceptually grounded. That is, I characterise my knowledge as *not* practical in our sense. I might be in a worse moral position than had I not come to know I was upsetting him, but I am clearly in a better one than had I said “Yeah I know, I love winding him up with that story”, implicating myself in having known what I was doing all along, not on the basis of perceptual clues and not in a way that constituted a discovery.

b. Thinking about the Special Relationship

What is the correct way of understanding the special relationship between intentional action and practical knowledge? Are *all* intentional actions objects of practical knowledge? *Ought* we think of the relationship as a conceptual one? If not then how are we to make sense of the observations in the previous two paragraphs? These are questions which are answered differently by different accounts of practical knowledge and I don’t want to pre-judge things here. I do, however, want to make a few comments about the options available for thinking about these issues.

First, quite apart from the considerations in (a), it is attractive to think of practical knowledge as essential to intentional action because this would underwrite a *unified* (accommodating) explanation of all three of practical knowledge’s apparent features. On such a view, the fact that I am intentionally (say) typing would be *dependent* on the fact that I know that I am. My knowledge will display *FPA* because the fact that I am intentionally typing is dependent on *my* knowledge but not on anyone else’s, thus grounding a first-/third-person asymmetry. My knowledge that I am typing could not be epistemically grounded in any way, for the reason that the *existence* of an intentional action to be known about on any epistemic basis would presuppose my *already* knowing about it. And such a view leaves no room for me to *figure out* or *notice* that I am intentionally typing because these concepts represent my knowledge as latching on to a fact independent of it, but it is just this kind of independence between the intentional action and the practical knowledge which is denied by someone who thinks that practical knowledge is essential to intentional action.

Not only would such a view meet the *Apparent Features* constraint, but it would seem to do so in the best way possible, by giving *one* explanation of all three apparent features. An account which accommodates the features in a *unified* way seems preferable to one which does not, since it doesn’t leave us wondering *why* the three philosophically interesting features from §2 coalesce together in the phenomenon of practical knowledge.

But the idea that practical knowledge is essential to intentional action is thought by many to admit of counterexamples, and so to be untenable. The classic counterexample comes from Davidson:

[I]n writing heavily on this page I may be intending to produce ten legible carbon copies. I do not know, or believe with any confidence, that I am succeeding. But if I am producing ten legible carbon copies, I am certainly doing so intentionally.²²

Davidson’s example and ones like it have convinced many that practical knowledge *cannot* be essential to intentional action.²³ We will consider this kind of putative counterexample

²² Davidson, “Intending,” 92.

²³ See e.g. Hornsby and Stanley, “Semantic Knowledge and Practical Knowledge,” 121; Alvarez, “Acting

in detail when we consider Lucy O'Brien's views in Chapter Two, and again at the end of Chapter Five. I will suggest in those places that it is not clear that examples such as this are as damaging as they have been supposed to be, to the view that an agent's practical knowledge is essential to her action's being intentional.

But even if we do take examples like Davidson's to show that *not all* intentional actions are objects of practical knowledge, this need not entail that the relationship between the phenomena of intentional action and practical knowledge is a merely empirical one. Drawing this conclusion assumes that if the relationship between two types of phenomenon is a conceptual one, then *any* instantiation of the one type entails an instantiation of the other. But this ignores the possibility of the conceptual relationship being a *normative* one. By way of analogy, we might think there is a conceptual relationship between something's *being a car* and its *being able to get one from A to B*. But that doesn't mean that *any* car will *in fact* be able to get one A to B. Counter-examples can be found piled up on scrap-heaps. It is not plausible to conclude from these counter-examples that *being a car* is only *contingently* related to *being able to get from A to B*; it is not plausible that being a method of transportation is merely *accidental* to what it is to be a car. The correct conclusion is instead that some cars are defunct or deficient. The objects on the scrap-heap count as cars in having many of the major features of cars, but they are duds in lacking *this* one. And the possibility of thinking of them as *dud cars* rather than as non-dud non-cars seems parasitic on thinking of cars as *essentially* (*inter alia*) such that they *ought* to get one from A to B. Equally, even if Davidson's example and ones like it show that some intentional actions are not objects of practical knowledge, this needn't mean that practical knowledge is only accidentally related to intentional action. It might instead show that ideal or paradigm cases of intentional action are objects of practical knowledge; that only if it is an object of practical knowledge will an action be paradigmatically intentional.

It is important to be clear about the implications of saying that intentional action is somehow dependent on practical knowledge, for if practical knowledge *is* like this, then it appears radically different to knowledge as ordinarily understood. Knowledge, we are inclined to think, seems essentially to involve some kind of successful mental 'latching on' to facts, which would mean that the facts are 'there' to be latched on to, independently of the latching-on. It doesn't matter to the truth of the proposition *that Jess is cycling* that I know that she is. And it doesn't matter to the truth of the proposition *that I am standing on Henry's paint-tube* that I know that I am. I *might* know about these things, but their being as they are doesn't *depend* on my knowing about them. But if practical knowledge is a condition on acting intentionally; if intentional action is essentially a kind of knowing action; then the fact that *I am salting my chips*, or upsetting Dan, when these are things I do intentionally, *is dependent on my knowing that I am*. Any account which accepts that intentional action is essentially known cannot shrink back from accepting that they are providing an account of practical knowledge which is at odds with certain philosophical assumptions about what *knowledge (per se)* is like. I will end up defending just such an account in Chapter Six.

Intentionally and Acting for a Reason," 294; Peacocke, "Mental Action and Self-Awareness (I)," 373–374; Paul, "How We Know What We're Doing," 4; Setiya, "Knowledge of Intention," 172.

c. *The Special Relationship Constraint*

In this section I have tried to explain why the phenomena of intentional action and of practical knowledge seem to be in a very tight relationship with one another, and to introduce some considerations relevant to deciding the exact nature of this special relationship.

With this in mind we can now lay down a second constraint on an acceptable account of practical knowledge:

The Special Relationship Constraint

An adequate account of practical knowledge should contain a plausible conception of the nature of the special relationship between practical knowledge and intentional action.

Any account of practical knowledge which resulted in the conclusion that many of our intentional actions are *not* objects of practical knowledge would *not* contain a plausible conception of the special relationship – it won't count the relationship as special enough. But additionally, an account which construes the relationship as a conceptual one will have to have something plausible to say about examples like that of Davidson's Carbon Copier.

4. *Further Remarks on Practical Knowledge*

In this concluding section I will make a couple of further comments about the way in which some have thought about practical knowledge, and re-cap both our initial sketch of practical knowledge, and our two initial constraints.

a. *Practical Knowledge and tense*

Throughout this chapter I have been thinking of practical knowledge as knowledge of *current* or *ongoing* intentional actions. But practical knowledge is often assumed to take in both present and future intentional action. This is because knowing what one *will* intentionally do sometimes displays philosophically interesting features very similar to those displayed by one's knowledge of what one *is* intentionally doing (I will return to the 'sometimes' in (c)).

For example, I know that I will make a coffee when I have finished editing this chapter. This knowledge too seems to display *FPA*, to be *epistemically ungrounded* and to be *psychologically immediate*. Although it is possible for someone to contradict me if I say "I'm going to make a coffee in a minute" ("No you're not – the kettle's on the blink"), in many cases I do seem in a better position than someone else to know what I am going to do. The reason why seems related to the fact that my knowledge that I will make a coffee looks epistemically ungrounded: although someone else will need to ask me or recall general patterns of behaviour to know that I will make a coffee when I've finished editing this Chapter ("She's only had one coffee this morning; she usually has at least two and she tends to make one after finishing a chunk of work"), *my* knowledge doesn't seem based on facts like this - even assuming I recognise these facts. Finally, I can say *straight-off*

if prompted that I am going to make a coffee (“Do you want a coffee?” - - - “I’m going to make one myself in a minute – I’ll make you one too”), and the knowledge I express when I do so doesn’t seem to represent a discovery, or something I have noticed about myself.

Despite these similarities between one’s knowledge of one’s future intentional actions and one’s ongoing intentional actions, our starting-sketch construes practical knowledge as knowledge of what one *is* intentionally doing. There are a few reasons for this decision. First, even if practical knowledge does come in a future-directed form as well as a present-directed one, it makes for a simpler discussion to at least begin by considering practical knowledge in just one of these forms. Second, I don’t want to pre-judge the issue because one of the accounts we will consider – due to Lucy O’Brien – thinks of practical knowledge as only taking in *ongoing* actions. In the end, I will suggest that there are good reasons even beyond the similarities described in the previous paragraph, to regard practical knowledge as taking in both present and future actions; this will be captured in one of our later constraints (see Ch. 5, §2(b)). But to consider these reasons here would take us too far afield and complicate my discussion. So for now we will proceed by considering only present-directed practical knowledge.

b. Knowledge of practical reasons?

Practical knowledge is also sometimes thought to take in not only an agent’s intentional action, but also her *practical reasons* (the reasons for which she acts). My knowledge that I am typing, the thought goes, is intimately bound up with my knowledge of *why* I am typing, *viz.* in order to edit this chapter, to finish my Thesis, and so on.

Again, it has been thought that one’s knowledge of one’s intentional action and one’s knowledge of one’s reasons have similar features. I am typically in a better position to say what my reasons are than someone else might be, so my knowledge of my reasons seems to display *FPA*. My knowledge of my reasons doesn’t seem to be based on evidence (it’s not as if, asked why I’m typing, I’d say “Usually, typing means I’m working, so I must be doing it in order to work.”), and it seems like it *can’t* be based on perception, so my knowledge of my reasons seems to be *epistemically ungrounded*. Finally, I can usually give my reasons for acting straight-off, and seemingly without making a discovery about myself, so my knowledge of my reasons seems *psychologically immediate*.

Things can go wrong, of course. In some cases someone else might have a better idea of my reasons than I do, for instance if I am self-deceived (perhaps I don’t care much about getting work done; I’m just trying to avoid cleaning the house). Psychological immediacy can break down too - finding oneself in the kitchen with no idea why is something that certainly happens. In such a case one’s knowledge of one’s reasons *may* be restored by noticing that it’s 6pm, the time one usually waters the spider-plants, or pours oneself a Scotch. Still, *typically*, one’s knowledge of one’s practical reasons might be thought to have the same features as one’s knowledge of one’s intentional action.

But why think that the *very same capacity* – the very same kind of knowledge - takes in *both* one’s intentional action and one’s practical reasons? Isn’t knowledge of one’s reasons better understood as an instance of *psychological self-knowledge*, knowledge of certain kinds of mental state; beliefs, desires and – some add - intentions? Well first, the idea that reasons (including practical reasons) are mental states is not without its problems – some

prefer to think of them as (at least sometimes) *facts* in light of which we act.²⁴ But more importantly, being able to give the reason for which one is acting has seemed to some *internally* related to the fact that one's action is done intentionally.²⁵ If this is right, and if intentional action and practical knowledge are themselves internally related – an idea considered in §3(a)&(b) (and which we will consider in detail in Chapters Five and Six) – then we might expect knowledge of reasons to have a closer relationship to knowledge of intentional action than would be allowed for by a view of the former as just another kind of psychological self-knowledge.

As with the question of whether practical knowledge takes in future intentional action, I have chosen to start by ignoring the possibility that practical knowledge is also knowledge of one's reasons for acting. First because not all of our protagonists think of practical knowledge as taking in practical reasons, so I don't want to pre-judge this issue from the start. And second because deciding what to say on the issue will – again – require philosophical work which would take us too far afield to do now. We will come to the issue – again – in Chapter Five (§2(b)). I will stop short of saying that practical knowledge is knowledge *of* an agent's practical reasons, as well as being knowledge of her intentional action, but we will see that practical knowledge is importantly *sensitive to* an agent's practical reasons. The import of this distinction will be clear only in the context of the discussion of Chapter Five, so we will put the question of the practical knower's relationship to her practical reasons to one side until then.

c. *Practical Knowledge, practical reasons and intention*

In (a) I said that our knowledge of *certain* of our future intentional actions appears to have similar features to those which appear to attach to practical knowledge as described in §2. Which future intentional actions are known in this seemingly special way? The answer is those one *has decided*, or *intends* to do. Various philosophers have been impressed by the seeming fact that, as Grice puts it, “the ordinary concept of intention is such that if one intends to do A, one is logically debarred from relying on evidence that one will in fact do A.”²⁶ The reason seems to be that my intentional actions are just those actions which it is in some sense *up to me* to perform. Any intention-independent evidence that I will do something seems at the same time to be evidence that it is *not* up to me to do that thing. If a wave of nausea appears to me to be strong evidence that I will be sick, then being sick does not appear to me to be something I can choose to do or not; whereas if being sick *does* appear to me to be something I can choose to do or not, then any wave of nausea would *not* appear to me to be strong evidence that I was going to be sick.

The concepts of decision and intention also have tight relations to that of *practical reasons*: decisions and intentions are formed on the basis of, and justified by, practical reasons. So the two suggested extensions to our sketch of practical knowledge which I considered under (a) and (b) are related, in both having internal links to the concepts of decision and intention.

And it ought not go unnoticed that our starting sketch itself is of a phenomenon somehow internally linked to intention: recall that practical knowledge is knowledge of

²⁴ Hyman, “How Knowledge Works”; Dancy, *Practical Reality*, 103; Alvarez, *Kinds of Reasons*, 3.

²⁵ Anscombe, *Intention*; Vogler, “Anscombe on Practical Inference”; Newstead, “Interpreting Anscombe's Intention §§32FF”; Thompson, “Naive Action Theory.”

²⁶ Grice, “Intention and Uncertainty,” 269. See also Hampshire, *Freedom of the Individual*. Ch. 3 and Anscombe, *Intention*. esp. §§2-3.

specifically *intentional* action. These observations suggest that the concept of intention might turn out to play some special role in an account of practical knowledge. Indeed this is just the conclusion I hope to draw at the end of Part One, and the thought which will form the basis of the following discussion.

5. Conclusion

This Chapter has provided an initial sketch of our target phenomenon. §1 introduced practical knowledge as knowledge with a certain kind of content. §2 detailed three philosophically interesting features which apparently attach to knowledge with this content: *FPA*, *epistemically ungroundedness* and *psychological immediacy*. §3 explained that the relationship between practical knowledge and its object seems to be of a special kind, and considered some possibilities for thinking about its specialness. And §4 considered two ways in which some have wanted to extend the scope of practical knowledge (from present to future intentional action, and from intentional action to practical reasons), and gestured at the idea that an agent's *intentions* might play a special role in her practical knowledge.

During this Chapter we met with two initial constraints on any account of practical knowledge. It is worth re-capping these before moving on to Chapter Two:

The Apparent Features Constraint

An adequate account of practical knowledge will either accommodate, or explain away, all three of practical knowledge's apparent features.

The Special Relationship Constraint

An adequate account of practical knowledge should contain a plausible conception of the nature of the special relationship between practical knowledge and intentional action.

More constraints will be added as we go along, uncovered through the process of considering the pro's and con's of the accounts of practical knowledge we will discuss. This process starts in the next Chapter with a look at two approaches to practical knowledge, exemplified respectively by accounts due to Lucy O'Brien and Sarah Paul.

Chapter Two:

Alternatives to Intentionalism

Introduction

Lucy O'Brien's and Sarah Paul's conceptions of practical knowledge are rather different from one another. O'Brien thinks, and Paul does not, that practical knowledge is grounded in a kind of conscious awareness of acting. Paul thinks, and O'Brien does not, that knowledge of one's intention inferentially grounds practical knowledge. But despite their differences, O'Brien and Paul have it in common that they both represent *alternatives to Intentionalism about Practical Knowledge*, the view that an agent's practical knowledge is in some sense *constituted by* their intention. So I am considering O'Brien's and Paul's views in a single chapter as a means to motivating my consideration of *Intentionalism*. For in the end I will reject both views (for different reasons) and *Intentionalism* will appear likely to overcome the problems in both.

Each account is likely to seem attractive from a certain philosophical point of view. Lucy O'Brien's account may well look attractive from the point of view of someone struck by the fact that *being an agent* seems to carry with it a distinct phenomenology: *moving my body* feels different to *having my body move*. And Paul's view is likely to look attractive from a rather different starting-point; one from which it is important if possible to characterise practical knowledge as – despite not being grounded in perception – continuous with and the same in kind as our other ordinary empirical knowledge.¹

I consider O'Brien's account in §1 and Paul's in §2. I will argue that O'Brien does well on the *Apparent Features* constraint but fails to meet the *Special Relationship* constraint in two different ways. Paul does not meet the *Apparent Features* constraint and does not *clearly* meet the *Special Relationship* constraint. My discussion of Paul leads me to set out a new constraint, which her account also fails to meet: the *Why Knowledge?* constraint requires any account of practical knowledge to leave it unmysterious why the phenomenon it describes really is a kind of *knowledge*. I suggest that O'Brien's problem comes from tying practical knowledge to the phenomenology of acting, and that Paul's comes from seeing practical knowledge as inductively inferentially justified. Since these are the core elements of each view, it seems likely that the reasons for rejecting O'Brien's and Paul's accounts will carry over in some capacity to other accounts which share these core elements.

¹ Paul is also motivated by wanting to avoid identifying intentions with beliefs (Paul, "How We Know What We're Doing," 2–3.). We will consider views which make this identification in Chapter Three.

1. Lucy O'Brien's Consciousness-Based Account

A few clarificatory remarks will be useful before I introduce Lucy O'Brien's account of practical knowledge. First, O'Brien is unusual amongst our protagonists in that it is not completely clear whether the phenomenon she is interested in is quite the same as that sketched in Chapter One under the name 'practical knowledge'; different things she says pull in different directions. But what matters in this discussion is not whether as a matter of fact O'Brien is interested in our target phenomenon, but whether the account she provides is apt to do duty *as* an account of our target phenomenon. Throughout this section I will be speaking *as if* O'Brien intends her account to capture our target phenomenon, but this is only really to make discussion easier. (This practise will also gloss over the fact that O'Brien does not use the expression 'practical knowledge' but instead tends to talk of 'agent's knowledge' or 'our knowledge of our own actions'.) I will only be in a position to explain why the question of what phenomenon O'Brien is *in fact* interested in explaining is hard to call when we have considered her view in detail, so I will explain the concern more fully when I conclude this section in (f).

Second, O'Brien's consciousness-based account of practical knowledge is perhaps less well-known than that of Chris Peacocke.² I have chosen to consider O'Brien's account over Peacocke's because I think it is better worked-out in certain respects, and because it accords better with our commonsense-human methodology.³ As I suggested in my introduction to this Chapter, I suspect that the problems I bring up for O'Brien will be equally problems for Peacocke, indeed for any consciousness-based view. I won't have time to argue for this *per se* but I will make some general comments in concluding this section which explain why I think that consciousness-based accounts seem inapt to capture our target phenomenon.⁴

Finally, to see O'Brien's account in context we ought to note that it is provided as part of a broader project of understanding psychological self-knowledge.⁵ O'Brien thinks that many of our psychological goings-on are *mental actions*, which would mean that an account of practical knowledge can do duty as an account of much of our psychological self-knowledge. I won't consider the prospects for this project here; my concern is only with what O'Brien says about practical knowledge *per se*. With this in mind and in keeping with the rest of my discussion in this Thesis, I will concentrate in this section on what O'Brien has to say about our knowledge of those of our actions which are *not* mental.⁶

² Peacocke, "Action: Awareness, Ownership, and Knowledge", "Mental Action and Self-Awareness (I)" and "Mental Action and Self-Awareness (II)." See also Pickard, "Knowledge of Action Without Observation."

³ In particular, Peacocke's view makes ineliminable use of a theoretical notion of 'tryings', understood as mental events which cause both the bodily movements involved in acting, and an experiential state in which one seems to be acting in a certain way (Peacocke, "Action: Awareness, Ownership, and Knowledge?"). When one *is* acting in this way, one is entitled to take the appearance at face-value, doing which furnishes one with knowledge. But 'tryings', understood in this way, are not part of commonsense psychology; they are theoretical posits (I try to hit the stumps *by throwing the ball*, not by performing a mental action which kick-starts a bodily one). O'Brien's account does not rely on Peacocke's theoretical notion of 'trying'. For her own objection to Peacocke's 'tryings' (which are independent of mine), see O'Brien, *Self-Knowing Agents*, 146–155.

⁴ Indeed the same kinds of objection as I make against O'Brien are leveled against Peacocke and Pickard in Newstead, "Interpreting Anscombe's Intention §§32FF."

⁵ See in particular O'Brien, *Self-Knowing Agents* and "Self-Knowledge, Agency, and Force."

⁶ She talks about knowledge of non-mental actions in Chapter 9 of *Self-Knowing Agents*, and in her earlier (2003) paper, "On Knowing One's Own Actions," on which that chapter is based. Knowledge of mental actions is discussed in Chapter 6 of *Self-Knowing Agents* and in "Self-Knowledge, Agency, and Force."

The plan for this section is as follows. In (a) I describe O'Brien's account in some detail and explain how it meets the *Apparent Features* constraint. In (b) and (c) I give two different reasons why the account seems to flout *Special Relationship* constraint. The second of these reasons is that O'Brien limits the scope of practical knowledge to our more *basic* actions, and in (d) I reject O'Brien's motivation for doing so. In (e) I explain why allowing practical knowledge such a limited scope is not an optional element of O'Brien's view, but stems from understanding practical knowledge in terms of conscious awareness of action. I conclude in (f) by raising some more general worries about linking practical knowledge so tightly with the phenomenology of agency.

a. Agent's Awareness and Practical Knowledge

O'Brien accepts that practical knowledge is not epistemically grounded in any of the 'usual' ways we considered in Chapter One (§2(b)), but she denies it is *altogether* epistemically ungrounded. Her thought is that:

... our beliefs about what we are doing are given epistemic support by what we are in fact doing: we know what we are doing *because* we are doing it.⁷

The fact that the agent is doing something is made available to her as a distinctive epistemic ground for a judgment that she is doing it, *via* 'agent's awareness', a *sui generis* specifically first-personal, potentially non-conceptual awareness of acting. To be agent-aware of what one is doing is not yet to *judge* that one is doing it, but it does give the agent a *prima facie* entitlement to such a judgment.⁸ Because I am agent-aware of typing, any judgment I make to the effect that I am typing will be one I am *entitled* to make, at least assuming that I have no over-riding reasons to think that my agent's awareness is non-veridical. As long as my judgment is *true*, it will constitute knowledge. This in a nut-shell is O'Brien's account of practical knowledge.

But is it right that, on a view like this, my knowledge that I am typing grounded *in the fact that I am typing*, as O'Brien suggested in the passage quoted above? Doesn't a view like this instead involve the claim that my knowledge that I am typing is grounded in *my agent's awareness* of typing? No. To see why not we need to consider the details of O'Brien's understanding of agent's awareness.

We are agent-aware of those things we do *consciously*, something which, O'Brien points out, might be understood in one of two ways. On what we might call a 'second-order' model, I will be typing consciously only if I am typing, and if I am in some state which takes my typing as its object. On this model, my consciousness of typing would be a state separate from my typing itself, and it would be this – *rather than* my typing itself – which would entitle me to judge that I am typing. But O'Brien accepts a different – what we might call an 'adverbial' – model of acting consciously, on which acting consciously is acting *in a certain way or mode*;⁹ on which being agent-aware of typing is a matter of typing in this mode, and so on which as long as I *am* typing in this mode, my typing and my consciousness of typing are aspects of the same event or process. On the adverbial

(2005).

⁷ O'Brien, *Self-Knowing Agents*, 182; this is characterised as a "naïve thought", not something to be argued for.

⁸ See also Peacocke, "Action: Awareness, Ownership, and Knowledge," 107.

⁹ O'Brien, *Self-Knowing Agents*, 183.

model, my judgment that I am typing *can* be grounded in my typing itself, as long as I am typing in whatever way or mode constitutes it's being done consciously.

So far so good. But what *is* involved in typing in the 'conscious' mode, and how exactly does doing so entitle me to judge that I am typing?¹⁰ I act in the conscious mode, thinks O'Brien when I act "with a sense of control":

We, of course, carry out many actions as a matter of habit, or as relatively automatically and inattentively, but when we are conscious of what we are doing – when we are agent aware – we seem to act with a sense of guiding our action, with a sense of control.¹¹

Having the relevant sense of control over one's action requires that the action is done "*directly* on the basis of an evaluation of the possible ways of acting, grasped as possible actions".¹² Because the relevant possibilities are possibilities *for the agent's own* action, in realising one of these possibilities the agent "secures awareness of the possibility realized as an option *for her*",¹³ resulting in awareness which is specifically first-personal and agential, and so apt for grounding first-person judgments that *one is acting* in some way.

It is important to be clear about the epistemological role O'Brien sees agent's awareness as playing. Agent's awareness *entitles* an agent to judge that she is acting in the way given in the content of the agent's awareness. The judgment is not "justifiable with reference to reasons statable by, or comprehensible to, the knowing subject",¹⁴ but is nevertheless a *rational* transition rather than a merely *reliable* one because not only does it "[tend] (in normal circumstances) to lead to true beliefs" but because further, "its tendency to do so is explained by constitutive or relatively *a priori* features of the *nature and content* of the states or acts of the subject involved."¹⁵

The idea of being epistemically entitled to make a judgment by the presence of a likely non-conceptual state which cannot be given in justification of that judgment deserves philosophical attention, and receives it in the epistemology literature. If the idea is problematic, then so too is O'Brien's account of practical knowledge, but I am not going to assess the idea of epistemic entitlement here - I will assume that O'Brien can help herself to the notion. If the notion *is* a good one, it brings clear benefits to O'Brien's account because it enables her to reconcile the seeming tension between, on the one hand, thinking of our judgments about what we are doing as seemingly *epistemically ungrounded* and, on the other, as properly speaking *knowledge*. Thinking of the agent as *entitled* to her practical judgments by her agent's awareness enables O'Brien to think of her as *knowing* what she is doing, whilst explaining why agents won't see their judgments as epistemically grounded in any of the usual ways, for their judgments *are not* grounded in any of the usual ways, and the agent's justification is not transparent to her.

¹⁰ It is here that O'Brien offers more detail than Peacocke; see Ibid., 184–188.

¹¹ Ibid., 184.

¹² Ibid., 187.

¹³ Ibid., 185.

¹⁴ O'Brien, "Self-Knowledge, Agency, and Force," 588. For more on the relevant notion of entitlement see Sosa, "Intellectual Virtue in Perspective"; Peacocke, "Our Entitlement to Self-Knowledge" and *Being Known*; Burge, "Content Preservation" and Peacocke, "Three Principles of Rationalism." O'Brien picks out the first three as containing her favoured gloss on the notion. For more detail on the way in which *agent's awareness* might entitle an agent to judgments about what she is doing, see Peacocke, "Action: Awareness, Ownership, and Knowledge," 105–108.

¹⁵ O'Brien, "Self-Knowledge, Agency, and Force," 590. O'Brien objects to simple reliabilism in O'Brien, *Self-Knowing Agents*, 178. See also Peacocke, "Action: Awareness, Ownership, and Knowledge," 105–6.

Let's consider directly how O'Brien explains practical knowledge's apparent features. First, the fact that practical knowledge is based on agent's awareness - itself specifically awareness of *one's own* action - underwrites a first-/third-person asymmetry in the domain of intentional action. *Only I* have agent's awareness of *my typing* and so I have a route to knowledge of the fact that I am typing which is not available to you, so my knowledge that I am typing displays *FPA*. This route is, as we have just seen, not an inferential one; my agent's awareness entitles me to judge that I am typing but I do not *infer* that I am typing from my agent's awareness. Inference can only take one from judgment to judgment and my agent's awareness is not itself a judgment. My practical knowledge is not perceptually based because my agent's awareness is not a perceptual state, including a proprioceptive one, so my knowledge is *epistemically ungrounded* in any of the usual ways. Although it is nevertheless epistemically grounded - in agent's awareness - O'Brien can explain (as I explained in the previous paragraph) why it would *seem* epistemically groundless. Finally, given my agent's awareness of typing, I can judge that I am typing, on O'Brien's view, without needing to engage in consideration, or otherwise to figure out or notice that I am typing. On O'Brien's account my practical knowledge appears to be *psychologically immediate* because it *is* so.

So O'Brien's view seems to meet the *Apparent Features* constraint rather well. But it has serious problems, which I turn to now.

b. O'Brien and the Special Relationship

Part of what the *Special Relationship* constraint demands is that an account of practical knowledge should explain why *only* intentional actions are objects of practical knowledge. O'Brien accepts this.¹⁶ But is not obvious, on her view, why this should be.

Initially, she seems to have a neat explanation, which draws on the generally accepted entailment from an action's being an upshot of practical reasoning to its being intentional. We saw in (a) that we are agent-aware when we act consciously, and that conscious action is "the result of a process of evaluation of the possibilities available, grasped as possibilities".¹⁷ A process of assessing possibilities for action which results in the realisation of one of those possibilities looks like a process of practical reasoning.¹⁸ If this is right then we have practical knowledge only of our intentional actions because a) practical knowledge is grounded on agent's awareness, b) we have agent's awareness only of what we do on the basis of practical reasoning and c) actions done on the basis of practical reasoning are *intentional*.

But the suggestion faces problems. First, O'Brien seems to think of the process of assessing the possibilities for action as (at least potentially) involving operations on *non-conceptual* contents,¹⁹ which means that the 'assessment of possibilities' which give rise to agent's awareness of the ensuing action doesn't look like practical reasoning as ordinarily understood; this involves operations on *propositional attitudes*.

¹⁶ O'Brien, *Self-Knowing Agents*, 169–170.

¹⁷ *Ibid.*, 184.

¹⁸ Whether or not this is O'Brien's view is unclear. For some relevant but inconclusive remarks see (*Ibid.*, 185–6.).

¹⁹ *Ibid.*, 167; *ibid.*, 186.

But even putting that issue to one side, a problem remains. Recall that acting *consciously* was identified with acting “with a sense of control”,²⁰ and opposed to acting “as a matter of habit, or [...] relatively automatically or inattentively”.²¹ This would be another reason to link practical knowledge with *intentional* action, if the relevant kind of ‘control’ was that we typically have over our specifically intentional actions. But it is very unclear why doing something on the basis of an assessment of the possibilities for action would guarantee that one acts consciously, or with sense of control in *this* sense. For assume that the agent engages *in the process of assessment* consciously (‘attentively’); I cannot see what rules out that the chosen action is *done* (‘realised’) inattentively and/or automatically. I might for instance be fully consciously engaged in choosing between making tea or coffee, and might consciously and attentively decide in favour of the latter, but then go on to *make coffee* on autopilot, my attention taken up with something else entirely; working out a crossword clue, talking on the phone, or whatever.

It does not help to stress, as O’Brien does, that agent’s awareness requires realising a possibility for action *directly* in response to one’s assessment.²² For whatever ‘direct’ might mean here, the problem does not depend on imagining a temporally extended or causally complex ‘gap’ between one’s assessment of the possibilities for action and one’s realisation of one of them. The problem does not rest on *any* particular conception of the *relation* between the assessment and the realisation, only on the idea that the assessment and the realisation are distinct events or processes. And this is hard to deny given that the assessment might not have resulted in the realisation (I might have got distracted; there might have been no coffee left etc.). As long as they are conceived as distinct, any amount of conscious attention to the assessment does not seem to entail conscious attention to the realisation.

An option here would be to drop the idea that the *assessment* is what underwrites agent’s awareness, and to think of it instead as a property of the *realisation* itself. On this view, consciously making coffee, and so being agent-aware of doing so, would just be consciously attending to *making coffee* whilst doing it. But then we lose the link between *φ-ing consciously* and *φ-ing intentionally*. For I can consciously attend to those things I do non-intentionally too - like sneezing, for example, or falling over.

I am not denying that sneezing and falling over on the one hand, and making coffee on the other, involve different kinds of awareness. The problem for O’Brien is how *she* is able to distinguish these kinds. She needs to do so in order to explain the epistemological and psychological differences between my practical knowledge that I am making coffee, and my non-practical knowledge that I am sneezing or falling over. The worry is that if agent’s awareness is merely a matter of consciously attending to the realisation of some possible action, then it’s hard to see what the difference is between agent’s awareness, which is supposed to be particular to our *intentional* actions, and the kind of awareness we have of those things – like sneezing or falling over – which we do *non-intentionally*. And this means that we lose the link between *practical knowledge* and *intentional action* which the *Special Relationship* constraint requires us to keep intact.

The dilemma is between an account which ties agent’s awareness to a process we can identify or link with practical reasoning, but which has difficulty explaining why agent’s awareness should in fact attach to the resulting actions, and an account which

²⁰ Ibid., 183 *et passim*.

²¹ Ibid., 184.

²² Ibid., 185 *et passim*; O’Brien italicizes ‘directly’ but does not explain in what sense it is being stressed.

understands agent's awareness as a property of the *realisation* of a possibility for action, but which seems to undermine the link required by the *Special Relationship* constraint, between *practical knowledge* and *intentional action*.

c. O'Brien on Practical Knowledge's scope

We have considered why it is not obvious that O'Brien can accommodate the fact that *only* intentional actions are objects of practical knowledge, but she also has a difficulty, I will suggest, in allowing that *enough* intentional actions are objects of practical knowledge.

In Chapter One, we saw that even if φ -ing intentionally does not *entail* knowing that one is φ -ing, it is at least *typically* or *paradigmatically* an object of practical knowledge. But according to O'Brien's account only a very restricted class of our intentional actions are objects of practical knowledge. O'Brien thinks that we only have practical knowledge of our actions under their more *basic* descriptions. We will see here why she thinks this, and I will explain why her motivation is mistaken in (d). In (e) we will see that the nature of her account of practical knowledge as grounded in conscious awareness of acting means that she is barred from adopting a more realistic, *i.e.* broader, conception of practical knowledge's scope.

O'Brien distinguishes between one's knowledge *that one is acting* (rather than not) and one's knowledge of *what* one is doing; knowledge *that one is* – specifically - *φ -ing*.²³ She thinks that whilst one is straightforwardly first-person authoritative *that* one is acting (when one is), one has only limited authority about *what* one is doing: we have *FPA* over our actions only under their *basic* descriptions.²⁴

We saw in Chapter One (§1) that an action can be intentional under more than one description. One of these descriptions characterises one's action as basic, as something one can 'simply do', that is, do without needing to do anything else. All other descriptions of the action will be non-basic. If I illuminate the room by turning on the light, which I do by flipping the switch, which I do by moving my finger, my action is *basic* under the description *moving my finger* and non-basic under the others.²⁵

We have *FPA* over our actions only under their more basic descriptions, thinks O'Brien, because less basic descriptions "seem to advert to the subject's main purpose or motive in doing what she is doing", and because:

... it seems [...] that while the subject will [...] be authoritative about what her purpose in acting is, and authoritative about the fact that she is acting for the purpose of doing one thing rather than another, she may not be authoritative about whether her purpose came off. ²⁶

O'Brien illustrates the idea with an example of someone playing snooker. Call this person Bev. Bev intends to get the pink ball in the corner pocket; she hits the cue-ball which strikes the pink, sending it along the top cushion and neatly into the corner pocket. It is implausible, thinks O'Brien, that Bev is authoritative about whether "she is

²³ Ibid., 160–161.

²⁴ Ibid., 160–3.

²⁵ See Ibid., 163. The example is, of course, from Davidson ("Actions, Reasons, and Causes," 4).

²⁶ O'Brien, *Self-Knowing Agents*, 162.

getting the ball in the corner pocket”²⁷ because Bev, no less than her opponent, has to wait and see whether the ball goes in. (Note the similarity of this example to Davidson’s Carbon Copier example introduced in Chapter One (§3(b)).)

O’Brien concludes that:

The claim of first-person authority with respect to our actions is to be understood as relative to certain descriptions which could be regarded as basic.²⁸

And she thinks that restricting the scope of practical knowledge to *basic* actions “draws the boundaries between those actions over which we are authoritative and those which we are not, in intuitively the right place”:

Those actions which we take ourselves to *know* how to carry out ‘just like that’ [...] do seem to be those actions we can know that we are doing immediately, without any needing [*sic*] to know that we have done anything else, or needing to check by monitoring feedback from bodily awareness of other perceptual faculties.²⁹

Because practical knowledge is first-person authoritative (as O’Brien accepts³⁰), restricting the scope of authority to basic actions is equally restricting the scope of *practical knowledge* to basic actions.

But O’Brien is *not* right to think that the basic/non-basic distinction is co-located with the divide between those actions of which we have (authoritative) practical knowledge and those of which we do not. For our knowledge of many of our *non-basic* intentional actions is not only first-person authoritative, but has all the other features of practical knowledge too. The point does not depend on assuming that basic action-descriptions are purely bodily descriptions of actions; O’Brien accepts that some of the things we can ‘simply do’ have higher-level descriptions than this. *e.g.* “picking up the cup just before us, writing familiar words...”³¹ The objection is that much of our practical knowledge is of actions which are non-basic even when ‘basic’ actions are described at a level higher than mere bodily movement.

Consider writing a thank-you letter; rearranging your bookshelf; replacing the inner-tube on your bike; making a cup of coffee; planning a holiday; making a cake. All of these can be done only by doing other things, so none are basic actions. It is very hard to imagine them being done unknowingly, and the knowledge in question has the apparent features of practical knowledge. Let’s consider one example in detail.

I am writing a thank-you letter. I can see my hand holding the pen, passing over the page with delicate movements and leaving writing in its wake, I am proprioceptively aware of the movements I am making and can feel the pressure between the pen and the page. I am unlikely to be specifically *attending* to any of this, but clearly I *could*. But my knowledge *that I am writing a thank-you letter* is independent of all this for familiar reasons (see again Ch. 1, §2(b)): I don’t need to *wait and see* what I am writing, unlike a nosy neighbour would have to. *She* must wait until I have written “Dear Molly, ...” to know that it is a *letter* I am writing, and wait a little longer – until I have written “I’m writing to thank you

²⁷ Ibid.

²⁸ Ibid., 189.

²⁹ Ibid., 167.

³⁰ Ibid., 158–160.

³¹ Ibid., 167.

for...” - in order to know it’s a *thank-you* letter. My knowledge is not grounded in this way.

Partly for this reason, my epistemic position is different to that of my nosy neighbour. My neighbour’s observations are consistent with a different hypothesis, *e.g.* that I am making a prop for a play. She can choose the most likely hypothesis on the basis of further considerations (perhaps she knows I’m in a play which requires just this kind of prop; or perhaps she knows that my grandmother Molly has recently sent me a gift). But it would be absurd to imagine *me* choosing between hypotheses in this way. So there is an epistemic asymmetry between me and my neighbour: my knowledge is *first-person authoritative*. And because I don’t hypothesise, my judgment doesn’t seem based on *evidence* either.

If asked, I can say straight-off what I am doing. I can say that I am *writing a thank-you letter* and not just that I am *moving my hand* or *writing*. And my capacity to say straight-off doesn’t seem to depend on my having *noticed* or *discovered* that I am: my knowledge that I am writing a thank-you letter is *psychologically immediate*.

As well as having all three of practical knowledge’s apparent features, my knowledge that I am writing a thank-you letter seems tightly bound up with the fact I am doing it intentionally. It is hard to imagine my writing a thank-you letter without knowing that I am; if I don’t know I am writing a thank-you letter, then I won’t know what, or even whether, to write as I approach the page. I might end up with picture of a cat instead. A stretching of the imagination apt to deliver a case of *unknowingly* writing a thank-you letter seems also to deliver a case of doing so *non-intentionally*; doing so in my sleep, for example. A different stretch which delivers a letter-writing which is known only by observation or on the basis of evidence, also delivers a case of non-intentional letter-writing; writing in the manner of someone speaking in tongues. Both stretchings of the imagination represent me as too cognitively cut-off from what I am doing to count as doing it intentionally. And so my practical knowledge of this *non-basic* action of writing a thank-you letter seems very tightly bound up with the fact that I am doing it intentionally.

Restricting practical knowledge to basic action misdescribes the facts. Even worse, not only does practical knowledge *in fact* take in non-basic actions, it seems to *need* to do so. We have already mentioned that an action’s being an object of practical knowledge has moral and legal relevance (*Introduction (d)*; Ch. 1, §3(a)). If I poison you knowingly I am on different footing, morally and legally speaking, than if I poison you unwittingly; and assuming I *know* I am poisoning you, it makes a difference to my moral and legal status whether my knowledge is *practical* or not. *Poisoning* is not something I can simply do, and so is not a basic action. But very generally, we care the most, from a moral and legal perspective, about people’s actions under their *non-basic* descriptions. The only reason it matters morally and legally speaking that I moved my hand *thus* is that it was an act of poisoning. And very many morally significant actions are too complex and extended in time to be easily thought of under single basic descriptions at all. Consider *saving the baby*, *going vegetarian*, *taking up an aggressive tax-avoidance scheme*, *spying for the Russians*. To maintain the link between practical knowledge and the relevant kind of moral and legal responsibility, we must think of practical knowledge as taking in actions under their non-basic descriptions.

So any account of practical knowledge which restricts it to basic action flouts the *Special Relationship* constraint, for the reason that it classifies a vast number of our intentional

actions as *not* objects of practical knowledge, when in fact they not only *are* but *need to be* in order to be subject to the kinds of moral and legal classification we know them to be subject to.

d. The Snooker Example reconsidered

We must resist restricting practical knowledge's scope in the way O'Brien does. But O'Brien did not make the restriction for no reason. She thought that it followed from the possibility of being "ignorant of the consequences of [one's] actions"³², illustrating the idea with the snooker example. So I want to explain here why I don't think that attention to the snooker example *does* warrant the scope restriction O'Brien makes. I will distinguish two ways in which O'Brien *generalises* from the snooker example, and suggest that neither generalisation is warranted. Bev might lack practical knowledge that she is potting the pink (although I will suggest later (Ch. 5, §3(a)) that even this is not obvious), but this doesn't mean that practical knowledge does not *in general* take in intentional actions under their non-basic descriptions.

Note first that *even if* knowing (at *t*) that one is potting the pink always requires knowing (at *t*) that the pink will go in the pocket, this is consistent with knowing non-perceptually and authoritatively (at *t*) that one is potting the pink. For one can know non-perceptually and authoritatively that the pink will go in the pocket. Let's agree that this is not possible for someone like Bev, who is not very good at snooker. Still, it *is* possible for someone like Ronnie O'Sullivan, who is.

If the pink in question isn't particularly tricky and/or if he is on good form, Ronnie might know *in* (and *whilst*) *striking the cue-ball* that *the pink is going (to go) in the pocket*. We can see in his behaviour that Ronnie knows that the pink is going in whilst striking the cue-ball – after doing so he immediately gets up, walks round the table to where the white ball will end up and gets down to take his next shot. When he's in his best form he *doesn't even bother looking* to see if the pink goes in the pocket – he already knows that it will.

We can stipulate that that the movement of the cue and the passage of the balls are identical for Ronnie and for Bev. And let's agree that neither player will know authoritatively and non-perceptually that they are *potting the pink* unless they know authoritatively and non-perceptually that *the pink will end up in the pocket* (although I will come back to this assumption). Being skilled and well-practised at something can mean that *because* one knows one *is doing* it, one also knows that one *will succeed in having done it*. We should resist generalising from the fact that someone *relatively unskilled* at a certain kind of non-basic action needs to 'wait and see' that they are pulling it off, to the conclusion that non-basic actions of this kind could *only ever* be known by their agents on the basis of observation, and so not authoritatively, and so not practically.

We imagined that both Bev and Ronnie needed to know, whilst taking their shot, that they *will have potted* the pink, in order to know whilst taking their shot that they *were potting* it. But not all cases look like this. In certain cases, one can *be doing* something without later *having done* it, a phenomenon often labelled 'the openness of the progressive' (the 'progressive' being the aspect attaching to propositions of the form " x is φ -ing"). And if someone can *be φ -ing* without ever *having φ -d*, then she can *know that she is φ -ing* without knowing that she ever *will have φ -d*. Writing a thank-you letter seems to work like this: I

³² Ibid., 160.

might be *writing* a thank-you letter – and know that I am - and never end up *having written* it – if for instance I suffer a stroke half-way through and never regain my capacity to write (and at any rate have other thing to worry about).

This gives us a second reason for not generalising in the way O'Brien does from the snooker example: *that* was an example in which knowing that one *is doing* something requires knowing that one *will have done* it. Bev (although not Ronnie) needed to rely on perception for the latter, and so the former too. But where doing something does not require having done it, *knowing* that one is doing it does not require *knowing* that one has done it. Perception has no role to play in this kind of case, analogous to the role it played in the Bev example.

The cases don't carve up neatly. We can even describe a snooker case which looks more like the letter-writing case: Ronnie has just struck the cue ball exactly as before. The shot proceeds in just the same way, when BOOM!!! – the arena is blown sky-high by a bomb. Later, on the news, we hear what seems perfectly true, "Earlier this evening the Crucible Theatre was blown up *just as Ronnie O'Sullivan was potting* a crucial pink." The pink never ended up in the pocket but *Ronnie was potting it* when the bomb went off, and if the explosion does not undermine the fact that he was potting it, nor does it undermine the fact that he knew that he was.

We needn't consider which features of different kinds of case are relevant to them falling into one or the other category. And I have not tried to exhaust all the ways in which one can, and cannot generalise from the Bev case. I have just highlighted two ways in which O'Brien's original snooker case is *unlike* other kinds of case, by way of explaining why generalising from the Bev case in the way O'Brien does is unmotivated. Even if *Bev* fails to have practical knowledge that she is potting the pink, this doesn't mean that *no* non-basic actions are objects of practical knowledge. It doesn't even mean that no *pink-pottings* are objects of practical knowledge.

It is possible that I have been reading O'Brien's scope restriction too strongly, for what she *says* is that attention to the snooker example shows that "drawing a description of an action from an intention is *not sufficient* to give one authority over the action".³³ But if the thought were only that we are not authoritative over *all* of our intentional actions, then restricting the scope of practical knowledge would not be a solution. For O'Brien accepts that we can fail to be authoritative even about some of our *basic* actions.³⁴ We needn't decide here whether or not I have been reading the scope restriction – as stated - overly strongly, for as we are about to see, O'Brien's *explanation* of practical knowledge seems to commit her to a view on which practical knowledge takes in only our basic actions; perhaps even only our *bodily* ones.

e. Agent's Awareness and Basic Action

Assuming that O'Brien can give a good account of what is involved in knowing our basic actions (that is, putting to one side the worries from (b)), I want to explain here why I don't think she can broaden the scope of her account of practical knowledge, even if she might want to in light of the considerations of (c) and (d).

³³ O'Brien, *Self-Knowing Agents*, 162; my emphasis.

³⁴ *Ibid.*, 189. See also Roessler and Eilan, "Agency and Self-Awareness: Mechanisms and Epistemology," 23–24.

O'Brien gives the following gloss on the way in which acting on the basis of an assessment of the possibilities for action (grasped as such) is supposed to secure agent's awareness:

Agent's awareness of bodily action would [...] involve the agent having a grasp of the possible ways that she could move her body as a basic action, and carrying out one action rather than another on the basis of an assessment of the possibilities available to her. The agent's grasp of which actions are available as basic actions will be based on a general grasp of the ways in which she can move her body without doing anything else. This grasp will itself be based on ways that she has moved her body in the past. It will also be based on a particular grasp of the position of her body at the time of action, which itself will be based on the ways she has recently moved it.³⁵

O'Brien here explicitly commits herself to the idea that the possibilities which are grasped prior to action are possibilities for *moving one's body*. Recall also that this grasp is supposed to be at least potentially non-conceptual. Agent's awareness is supposed to attach to an action which is realised on the basis of one's assessment of the possibilities, and somehow to do so in virtue of being realised on this basis (again, we are ignoring the problems with this idea discussed in (b)). If the process of assessing possibilities is supposed to furnish one with agent's awareness, and the grasp of these possibilities is a grasp of possibilities for *bodily movement* then it seems able only to deliver awareness of whichever basic bodily movement results from the assessment. And if practical knowledge is essentially grounded on agent's awareness, then practical knowledge could only be of one's basic bodily movements.

The worry is not that (say) *making a coffee* is not something an agent can consider as a possibility for action, but that possibilities *like this* don't seem to be of the right kind to play the role O'Brien needs them to play. This is at least in part because by 'conscious', O'Brien seems to have in mind some kind of 'feeling of doing', and this seems to at least involve a bodily component. It is something we can recreate in imagination when we imagine ourselves moving our bodies; the kind of thing which misfires in Alien Hand Syndrome, when the body moves but the 'feeling of doing' is absent or in certain kinds of phantom-limb hallucination, in which it seems to an (e.g.) armless person that she is moving her arm.³⁶

The experience O'Brien has in mind is more than the experience of one's body's moving; it is the experience of *moving one's body* – of *doing* something with one's moveable parts. And for this reason it is something properly called '*agent's awareness*', and conceived as distinct from proprioception. But whilst awareness of this kind might entitle me to judge that I am *moving my arm* (rather than just that my arm *is moving*), it doesn't seem able to entitle me to judge that I am *making coffee* or rather generally to judge that I am doing anything described in terms which *go beyond* bodily movements: the awareness in question seems *inherently* unable to stretch much, or at all, beyond my bodily actions.

Understanding agent's awareness in such a way that it could take in actions conceived independently of their bodily components would enable O'Brien to broaden the scope of practical knowledge so that it took in actions under descriptions at a much higher level than the bodily. But doing so would also be moving away from an account which

³⁵ O'Brien, *Self-Knowing Agents*, 184.

³⁶ That O'Brien is keen to accommodate phenomena of these kinds is made clear in *Ibid.*, 189.

bases practical knowledge on the phenomenon O'Brien describes under the label 'agent's awareness'; on a kind of conscious 'feeling of doing'.

f. Conclusion

O'Brien's account happily met the *Apparent Features* constraint but had two separate kinds of difficulty with the *Special Relationship* constraint. It seemed unclear why agent's awareness would be awareness of *only* intentional action, and it was not possible to have agent's awareness, in O'Brien's sense, of those of our intentional actions which were *non-basic*, which meant that O'Brien could not account for the fact that we *do* – at least typically – have practical knowledge (as described in Chapter One (§§1&2)) of our actions under their non-basic descriptions. So we should reject the idea that practical knowledge is grounded in agent's awareness.

I began my discussion of O'Brien by saying that it was not clear to me whether she is interested in quite the same phenomenon as we are. I can now explain why this is hard to call. On the one hand, she wants to account for a kind of knowledge which has just the content and special (apparent) features we set out in Chapter One (§1 & §2); this is shown by how well she meets the *Apparent Features* constraint. But on the other hand, an account of practical knowledge as grounded in agent's awareness seems badly placed to do so, not least because the *scope* of the kind of knowledge with this content, and these features, seems much broader than O'Brien's account allows for. There is another problem, too, with thinking of practical knowledge – of *our* target phenomenon – as grounded in the phenomenology of conscious action, which I will outline by way of concluding.

Words like 'conscious' and 'aware' can mean all sorts of things, but for O'Brien acting consciously is understood in opposition to acting absent-mindedly, inattentively or as a matter of habit. I am agent-aware of what I am doing when I am *attentively involved* in it, focussed on it, when it is central to my experiential state. But very often what one is intentionally doing is *not* central to one's experiential state, or an object of one's attention, as when I am making a coffee whilst trying to solve a crossword clue, or like just now when I was typing 'coffee' but concentrating on the point I wanted to make. Recall Falvey reading his book whilst waiting for his dough to rise. What he's *occupied with* is what's going on in his book, but he's still *making bread*. Very often too, we do things both intentionally and habitually, as when I lock the door on leaving the house. All these things are not only done intentionally, but are at least typically objects of practical knowledge – if asked one can say 'straight-off' that one is doing any of these things, for example, and one's knowledge doesn't appear to be epistemically grounded. With this in mind, tying practical knowledge very tightly to *conscious awareness* or the *phenomenology* of acting begins to look like going in the wrong direction.

This attention-independence of practical knowledge doesn't make it unusual as a kind of knowledge. On the contrary, knowledge *in general* is usually thought of as a dispositional kind of thing. My knowledge that the pangolin is an endangered species does not require me to be constantly consciously thinking about it, or attending to the fact. And just as well, for I know an *indefinite* number of facts. This isn't a delusion of omniscience (indeed it is consistent with another truth; that I am ignorant of an indefinite number too); it is a fact about what knowledge must be like. If I needed to consciously attend to facts in order to know them I would be overwhelmed with conscious thoughts and unable to get on with things; unable to *use* any of my knowledge.

The fact that I needn't be consciously engaged with a fact in order to know it is as true of my knowledge that I am making a coffee, that I am typing and that I am writing a PhD Thesis as it is of my knowledge that the pangolin is endangered. And the reasons are the same too; I am currently engaged in a number of ongoing projects, all of which I know that I am doing, but I can't attend to all of them at the same time.

I have explained how, because O'Brien's account of practical knowledge sees it as grounded in the conscious experience of acting, it is ill-placed to capture *our* notion of practical knowledge. And I think we have good reason to think that *any* account which takes this approach will be similarly ill-placed. Consciousness-based accounts of practical knowledge seem bound to misdescribe the phenomenology of practical knowledge. I don't mean by this that practical knowledge has a distinctive phenomenology which such accounts misdescribe, but that such accounts go wrong by tying practical knowledge to a particular phenomenology, when – like propositional knowledge more generally – it doesn't seem to have any.

2. Sarah Paul's Inferentialist Account

Like Lucy O'Brien, Sarah Paul accepts that practical knowledge is ungrounded in perception, but she denies that it is wholly *epistemically ungrounded* in any of the usual ways. On her view, my practical knowledge that I am typing is constituted by a belief that I am typing which is inferentially grounded on my knowledge *that I intend to be typing*, plus some supporting premises.³⁷

The discussion is structured as follows. In (a) I explain how Paul's account fares on our two initial constraints. We will see that unlike O'Brien, she has some difficulty meeting the *Apparent Features* constraint, and that it is at least not obvious how well she meets the *Special Relationship* constraint, although her difficulty is very different from O'Brien's (indeed Paul's account does very well in accommodating practical knowledge's scope). In (b) and (c) I consider Paul's epistemology in more detail and bring out a further problem for her account, which takes the form of a dilemma: at first sight it is hard to see how the inference Paul sees as underwriting practical knowledge can justify my judgment *that I am typing* rather than my judgment *that I can or might be typing* (b). And although there is a way of bolstering Paul's inference to get over this problem, doing so comes at the cost of introducing a problematic regress. Because of this she flouts a new constraint, which I introduce in (d): the *Why Knowledge?* constraint demands that an account of practical knowledge ought not leave it mysterious why the phenomenon it describes *is* a kind of knowledge.

a. Inferentialism and our two Constraints

Paul aims to “explain agential knowledge with a simple, familiar epistemic structure”,³⁸ that of inductive inference from evidence. We already know that inductive inferences can furnish us with knowledge, and Paul's strategy is to find a structure of this kind underlying practical knowledge. She accepts that practical knowledge can't be inferred

³⁷ Paul develops the account given in Grice, “Intention and Uncertainty.” See also Donnellan, “Knowing What I Am Doing.” Paul gives an account of future-directed practical knowledge very similar to her account of present-directed practical knowledge, but I will concentrate here only on the latter (see again Ch. 1, §4(a)).

³⁸ Paul, “How We Know What We're Doing,” 19.

simply from perceptual clues or from a combination of these and general knowledge of one's tendencies (see again Ch. 1, §2(b)). Instead, she thinks that we infer our practical knowledge from, *inter alia*, knowledge of our intentions. I know that I am typing, she thinks, because I intend to be doing so, I know this of myself, and I have inferred that I *am* typing on this basis. We will consider the inference, including its subsidiary premises, in more detail in (b). First, though, let's see how Paul's inferentialist account fares on our two initial constraints.

For Paul, the *particularities* of practical knowledge – its special (apparent) features – are to be explained by the particularities of one's knowledge of the crucial premise about one's intention. Paul seeks to accommodate practical knowledge's apparent *FPA*, explaining this as a product of the agent's special epistemic relationship to her own *intentions*. There is a first-/third-person asymmetry between my and your knowledge that I am typing because there is a first-/third-person asymmetry between my and your knowledge of my *intentions*, and because my knowledge of what I am doing is inferred from this. I won't here consider Paul's account of our knowledge of our own intentions.³⁹ The important point is that assuming that our knowledge of our intentions itself displays *FPA* – thinks Paul – so will the practical knowledge which is inferred from it. For this reason Paul would take herself to meet the requirement from Chapter One that a good account of practical knowledge should have something plausible to say about exactly what practical knowledge's *FPA* amounts to.

Whether or not she does is unclear. One might worry that her account locates *FPA* in the wrong place. John Schwenkler puts the point like this:

[O]n an account like Paul's while there will clearly be fundamental differences between your and my respective ways of knowing about my *intentions*, it appears that once that knowledge is in place, you may be in every bit as good a position as I am to know what I am doing on the basis of the knowledge of those intentions plus my general tendencies and the evident favourability of my circumstances [...] [Yet it] seems as if it is knowledge of one's *actions*, and not just of the intentions that underlie them, that ought to be characterized by first-person privilege [...].⁴⁰

Paul might respond that another person *de facto* won't ordinarily have as detailed a grasp of my tendencies, capacities and circumstances as I have, and that given this, *FPA* really does transmit from one's knowledge of one's intention to the practical knowledge inferred from it. Further discussion would be needed to consider whether this response is adequate. But I will pass over this here, because there is a more serious problem with Paul's capacity to meet the *Apparent Features* constraint.

The problem is with her capacity to explain practical knowledge's apparent *epistemic ungroundedness*. Recall that in order to meet the *Apparent Features* constraint, an account must *either* accommodate *or* adequately explain away each of practical knowledge's apparent features. Because she thinks that practical knowledge is inferred from evidence, Paul must explain away the seeming *non-evidential* nature of practical knowledge. But it's not clear that she can.

In Chapter One (§2(b)) I said that one of the manifestations of practical knowledge's seeming *epistemic ungroundedness* – which includes its *evidential* ungroundedness – was that

³⁹ But see Paul, "How We Know What We're Doing" and "How We Know What We Intend."

⁴⁰ Schwenkler, "Perception and Practical Knowledge," 143.

if asked how one knows that one is intentionally doing something, an agent will typically have nothing to say. How do I know that I am typing? Well, I just *do*. It certainly doesn't seem to me that my knowledge that I am typing is based on the fact that I *intend to be* doing so. But why should this be if my practical knowledge that I am typing *is* inferred from knowledge that I intend to be doing so?

Paul considers this kind of worry but dismisses it as “unconvincing”⁴¹ once we recognise that inferences needn't occur in or be tracked by consciousness. But although Paul is right to think that a judgment can be inferentially grounded without coming at the end of a conscious ‘premise-premise-conclusion’ train of thought, this isn't an adequate response to the worry about apparent *epistemic ungroundedness*. (It might be a good explanation of practical knowledge's apparent *psychological immediacy*, but I won't consider this here.)

For part of the reason we might think of a given belief as inferentially based is that its subject *could*, if asked, give the premises of the supposed inference as her reasons for holding the belief. I'm dreading getting soaked later. How do I know I'll get soaked? I haven't gone through any ‘premise-premise-conclusion’ thought-process in coming to know this, but if you ask me why I believe I'll get soaked, or how I know I will, I can *reconstruct* my grounds: “It's pouring down with rain and looks unlikely to stop. I have to go out later, I don't have a mac or a brolly, so I'll (probably) get soaked...”

Paul owes us an account of why, *if* asked how one knows what one is doing (or will do), an agent won't represent her knowledge of her intention as her grounds in this way. The apparent *epistemic ungroundedness* of practical knowledge means that not only would Paul's inference take place non-consciously, but it would actually be *opaque* to the knower; unreconstructible on questioning. And not only is the inference unreconstructible, but it is likely to be *actively disavowed* by the practical knower. If someone were to suggest to me that my evidence for thinking that I am typing is my knowledge that I intend to be doing so, I wouldn't find the suggestion particularly compelling. Usually, the criterion for thinking of a bit of knowledge as inferentially based is exactly that its premises are reconstructible and would be avowed by the knower if asked for them. But these criteria are not met by my knowledge that I am typing, and in light of this we lack positive reason to think that it *is* inferentially based on my knowledge of my intention.⁴² Paul does not meet the *Apparent Features* constraint because she fails to *explain away* the apparent epistemic – specifically *evidential* - ungroundedness of practical knowledge.

What about the *Special Relationship* constraint? Thinking of practical knowledge as inferred from intention certainly explains why only intentional actions should be objects of practical knowledge, and because intentions can concern actions under their non-basic descriptions, Paul looks likely to get practical knowledge's scope right too. On these points Paul's account is clearly an improvement on O'Brien's. But it's less clear whether Paul accurately characterises the nature of the special relationship between practical knowledge and intentional action, for the view is consistent with the brute empirical possibility that someone might simply fail to make the relevant inference, indeed that someone could *systematically* fail to do so. The result would be a person who had the normal capacity to act intentionally but whose knowledge of these actions – if any - didn't appear to display *FPA*, *epistemic ungroundedness* or *psychological immediacy*, and this

⁴¹ Paul, “How We Know What We're Doing,” 10.

⁴² For a similar thought see Hampshire and Hart, “Decision, Intention and Certainty,” 9–10.

might be thought implausible. Clearly there is more to say here, but we should at least be suspicious of Paul's capacity to meet the *Special Relationship* constraint.

One more point is worth making before we move on. One assumption of Paul's inferentialist view is that the capacity to know about one's intentions is independent of the capacity to know about one's own intentional actions. But this might be doubted. Kieran Setiya, for example, argues that in the same way that knowledge of our perceptual states is transparent to knowledge of the world, our knowledge of our intentions is transparent to knowledge of what we are *doing*.⁴³ In both cases, the result is that the capacity for self-knowledge of one's own psychological state (perception; intention) is dependent on the capacity for knowledge of what those states represent (perceptible goings-on; intentional actions). And if this is right, Paul's account must be rejected, for it rests on a false assumption. I won't consider Setiya's argument here, but I will come to a related conclusion in (c); that Paul's inference fails to reduce the capacity to know what one is doing to the capacity to know what one intends.

Summing up, then. Paul fails to meet the *Apparent Features* constraint because she fails to explain away the appearance of epistemic ungroundedness which attaches to practical knowledge. The appearance of epistemic ungroundedness means that any inference would not only need to *occur non-consciously*, but would also be unreconstructible by the agent. And second, she doesn't *obviously* meet the *Special Relationship* constraint because her view seems to allow for a wide-spread dislocation between *acting intentionally* and having *practical knowledge* of one's action. In the next section I will consider Paul's epistemology of practical knowledge more closely, and will describe what I think is an even more serious problem than the ones gestured at here.

b. Practical Knowledge's evidence-base

On Paul's view, my knowledge of what I intend gives me the crucial premise from which to infer knowledge of what I am actually doing, but clearly this premise is not enough on its own to ground my practical knowledge. In order to infer that I *am φ -ing* from the fact *that I intend to be φ -ing* Paul says that I must also know (or at least justifiedly believe⁴⁴) *first* that I *can φ* if I intend to, and second that my circumstances are conducive to φ -ing.⁴⁵

To facilitate discussion, let's set out Paul's inference-schema for present-directed practical knowledge:

Inference-Schema for Practical Knowledge

- | | |
|--|--|
| P1) | I intend to be φ -ing |
| P2) | I can φ if I intend to φ |
| P3) | My circumstances are conducive to φ -ing |
| C) I am φ-ing | |

Paul needs it to be the case that premises of the form (P1)-(P3) justify an agent in making a judgment of the form (C); *that I am φ -ing*, but it is unclear at first glance how they could do so. The problem is not that the premises do not *entail* the conclusion; the

⁴³ Setiya, "Knowledge of Intention."

⁴⁴ From now on I will talk simply of one's *knowledge* of one's intention, leaving the parenthesis implicit.

⁴⁵ Paul, "How We Know What We're Doing," 15.

premises of an inductive inference can *never* be expected to do this. The problem is that it is hard to see how they even make the conclusion *likely* true, and so reasonable to believe. For instance, although my knowledge that I currently intend to be typing, that I *can* type if I intend to, and that my circumstances are conducive to typing clearly justify the conclusion that I *could* or *might* be typing, they don't seem to justify the conclusion that I *am in fact* typing. There seems to be a justificatory gap between the premises and the conclusion of Paul's inference-schema.⁴⁶

Paul appears cognisant of this kind of worry; she recognises that she needs to explain the sense in which “what the agent knows without observation extend[s] beyond his own mind to *what actually happens*”.⁴⁷ But I don't think that her explanation of how this works is successful. Briefly, the idea is that because (as we commented in §1(d) above) *doing* something need not entail that one *will have done* it (the ‘openness of the progressive’), someone can know that they *are doing* something without knowing that they *will have done it*.

Paul considers how Jen knows that she is walking to the faculty, independently of perception. Her suggestion is that because Jen knows that her intention in walking is to get to the faculty, so as long as she can be confident that walking to the faculty is something she can do if she intends, and that her circumstances are conducive to doing it ((P2) and (P3)), she can know that in walking she *is walking to the faculty*. And because she can *be walking there* even if she doesn't end up *getting there* (the openness of the progressive), she can *know* that she is walking there on the basis of knowing that this is her intention, and in particular without having to rely on perception to give her clues about whether she *is getting* there.⁴⁸

It ought to be clear from the comments I made in §1(d) that I think that these kinds of consideration about the progressive are interesting, and of relevance to understanding practical knowledge. But I don't think they help Paul bridge the justificatory gap between the premises and conclusion of her suggested inference. For even if we accept that Jen can know that she is walking to the faculty on the basis of knowing that her intention in walking is to walk to the faculty, and that the mechanism behind this has something to do with the openness of the progressive, Paul tells us nothing whatsoever about how Jen knows *that she is walking*. Paul might be able to explain, with reference to features of the progressive tense, how practical knowledge *spreads* amongst action-descriptions (in this case, from *walking* to *walking to the faculty*) but she doesn't tell us how it gets started.

Jen is supposed to know that she *is walking* by inference from knowledge that (1) she intends to be walking, (2) that she *can* walk if she intends to and (3) that her circumstances are conducive to walking. But the justificatory gap remains unbridged – these premises seem to justify Jen only in believing only that *might* or *could* be walking; they fall short of grounding her practical knowledge that she *is* (in fact) walking (so too that she is walking to the faculty).

⁴⁶ For a similar objection to the inferentialist model (although pre-dating Paul's particular presentation), see O'Brien, *Self-Knowing Agents*, 177.

⁴⁷ Paul, “How We Know What We're Doing,” 16; my italics.

⁴⁸ Ibid. §5.

c. Infinite Regress

The foregoing assumes a certain reading of what I will call the ‘capacity premise’ (P2) - that “I can φ if I intend to φ ” - on which knowing that one can φ if one intends to means knowing that φ -ing is the kind of thing one is able to do intentionally (in execution of one’s intention). But there is a stronger reading of the capacity premise, on which “I can φ if I intend to φ ” means something more like “Intending to φ now is generally sufficient for φ -ing now”.⁴⁹ If Jen believes that intending to now be walking is generally sufficient for now *being* walking, then she would seem to be justified in believing that she is walking, on the basis of knowing that she intends to be doing so (as long as she knows that her circumstances are conducive). If I know that A is generally sufficient for B , and I know that A obtains then I can thereby know – by inference - that B obtains.

For this to work, Paul would need to say something about what ‘generally sufficient’ is supposed to mean here. Clearly, the sufficiency in question cannot be *logical* sufficiency since it is – and is recognised by agents to be - brutally empirically possible to intend to be doing something yet not actually be doing it, as when I intend to be saying something funny, but am not. The sense of ‘sufficient’ would need to *strong* enough to justify an inference from one’s intending (now) to φ to one’s (now) φ -ing, but *weak* enough *not* to generate logical entailments between intending and doing. The capacity premise, however it is to be understood in detail, needs to embody a recognition that intentions to be φ -ing *tend* to be accompanied by φ -ings.

We can bypass the question of how exactly to understand the stronger interpretation of the capacity premise, by simply adding in a premise whose content captures this tendency of intentions to result in the intended actions. The result is the following, bolstered inference-schema:

Bolstered Inference-Schema for Practical Knowledge

P1)	I intend to be φ -ing
P2')	I can φ if I intend to φ
P2'')	I tend to be φ -ing when I intend to be φ -ing
P3)	My circumstances are conducive to φ -ing
<hr/>	
C)	I am φ-ing

For clarity I have split the capacity premise into two, one interpreted along the original lines (P2') and one (P2'') in which the dispositional link between intention and performance is made explicit.

Although the bolstered premise bridges the justificatory gap we considered in (c), its prospects are not good. For both versions of the capacity premise presuppose prior practical knowledge, and this, as I will explain, leaves Paul without an account of how practical knowledge is possible.⁵⁰

⁴⁹ Thanks to Jane Heal for this suggestion.

⁵⁰ Strictly, this means that the current objection is as much of a problem for the original inference-schema as for this bolstered version.

According to the bolstered inference-schema, Jen's judgment that she is walking is justified in part by her knowledge that she tends to (be) do(ing) what she intends to (be) do(ing). But how does she know *this*? Paul suggests that past experience of exercising one's agency will play some role in an agent's knowledge of the supporting premises, although does go into detail.⁵¹ Still, it is easy to see how past experience of agency might furnish one with knowledge of (P2"). Jen will have had past experiences of intentionally doing those things she intended to be doing, and I guess Paul takes it that the knowledge embodied in a premise of the form of (P2") is a generalisation from these past instances.

The generalisation will be based on many past instances of intending to be doing something which she was then. Jen's knowledge that she tends to be doing what she intends to be doing (P2") is an inductive generalisation from various past bits of knowledge of conjunctions of the form "I intend to be φ -ing, and I *am* φ -ing". Knowing the conjunctions will have required knowing both conjuncts, and herein lies the problem, for *this* will have required having had various bits of prior knowledge of the form "I am φ -ing." But this seems to mean that her current practical knowledge essentially presupposes past pieces of practical knowledge, since these past pieces are required for her to know premises of the form (P2"), which in turn are needed to ground judgments of the form "I am φ -ing" rather than merely judgments of the form "I might be φ -ing".

If this is right, Paul's account of the epistemology is incomplete: she seems able to explain how a given bit of practical knowledge is justified by reference to her inference. But when we ask what entitles the agent to her knowledge of *her premises* (or why she is reasonable in believing them), we find ourselves needing to cite previous bits of practical knowledge. If Paul is right, these must be justified by further inferences of the same kind, which depend on yet further previous bits of practical knowledge, and so on. Any bit of practical knowledge depends on further (prior) bits of practical knowledge and so on *ad infinitum*. The regress which ensues is vicious because it leaves Paul unable to explain why judgments inferred in the way Paul suggests would count as *knowledge* at any level.

It might be suggested, in response, that the regress is not infinite; and so not problematic. One might say that although a great deal of our practical judgments do depend for their justification on prior practical knowledge (for the reasons just given), *at bottom* there is a kind of knowledge of one's own intentional action which is *not* practical, and so which *does not* itself rely on the same kind of inference, but which, given a critical mass of instances, can license the generalisation to (P2"). Once this generalisation has been made, an agent is in a position to derive – *via* Paul's bolstered inference – *practical* knowledge. Positively, the most plausible version of this suggestion claims that the regress bottoms-out in observation-based judgments of the form "I am (intentionally) φ -ing".

But the view this suggestion leaves us with is not attractive, for many will find it deeply implausible that as a matter of fact there is a period of time during someone's development in which she is able to act intentionally but can come to know what she is doing only by observation. This would also be a stretch of time during which one's knowledge of what one was intentionally doing lacked *FPA* and *psychological immediacy*: if a little tot knows that she is intentionally (e.g.) eating banana only by observing herself doing this, then intentionally eating banana is something she could *fail to notice*, and is known, if at all, as a discovery about herself. But it doesn't seem particularly plausible –

⁵¹ Paul, "How We Know What We're Doing," 15.

for reasons we considered in Chapter One (§2(c)) – to think of someone as able to act intentionally but capable of *noticing*, or failing to notice, what they are doing.

A more sophisticated kind of response involves denying that the bits of knowledge of the form “I am φ -ing” are knowledge of one’s actions *as intentional*.⁵² One can only know that one is *intentionally* φ -ing, on this view, *after* one has made the generalisation to (P2’): all self-knowledge of *intentional* action is *practical* knowledge and is based on Paul’s bolstered inference. Prior to generalising to (P2’) our tot is both incapable of practical knowledge, *and* incapable of intentional action. These capacities go together. This avoids the unpalatable consequence that during an initial stage our tot can act intentionally but know that she is doing so only observationally and psychologically mediately.

But this response has a different consequence which I think is equally implausible. On this picture our tot would have to go through a period in which she could intend to do things, and know this about herself, but could not act intentionally. The idea that the capacity to form intentions and to know that one has them, could float free of the capacity to *act intentionally* and to know that one is doing so, is hard to take seriously. It is hard to think of a creature forming intentions and recognising themselves as doing so, but not being able to execute these intentions. Sure, sometimes we fail to execute our intentions, and a small child is likely to be more incompetent than most. But it is hard to imagine a period in a creature’s life in which intentions were commonly and self-consciously formed but *never* executed. It is hard to know how such a creature at such a stage could be motivated to form intentions they recognised did not tend to come to fruition. Yet this is just what the present suggestion entails.⁵³

d. Conclusion

I have described various problems with Paul’s account. It was not obvious how well she was able to meet the *Special Relationship* constraint and she seemed to fail to meet the *Apparent Features* constraint. In the previous two sections I have described a dilemma which seems to undermine her ability to meet a fourth and very important constraint, which we can now state as follows:

The Why Knowledge? Constraint

Any account of practical knowledge ought not to leave it mysterious why the phenomenon it describes is a kind of knowledge.

On one horn of the dilemma, described in (b), Paul flouts this constraint because a judgment of the form “I am φ -ing” is not justified by the inference’s premises; at best they justify the agent in judging “I might (or could) be φ -ing”. Yet practical knowledge is knowledge that one *is in fact* φ -ing. Bolstering the premise-set to bridge this justificatory gap comes at the expense of inviting a problematic regress, for the extra premise needed to do so itself depends on prior bits of practical knowledge. The regress cannot be stopped without generating implausible consequences. For these reasons we ought to reject Paul’s account, and because the problems are problems with the very idea that practical knowledge might be inferred from knowledge of one’s intention, it seems likely that any view of this broad kind will be subject to the same objections.

⁵² Thanks to Alison Fernandes for this version of the response.

⁵³ Some empirical work relevant to these thoughts is discussed in Roessler, “Intentional Action and Self-Awareness,” 385–388.

Concluding Part One

In Part One I have tried to do three things; to clarify our target phenomenon - *practical knowledge*; to consider and reject two broad approaches to understanding it by considering and rejecting exemplars, and to introduce and explain three constraints on a good account of practical knowledge. It is useful to re-cap these three constraints:

The Apparent Features Constraint

An adequate account of practical knowledge will either accommodate, or explain away, all three of practical knowledge's apparent features.

The Special Relationship Constraint

An adequate account of practical knowledge should contain a plausible conception of the nature of the special relationship between practical knowledge and intentional action.

The Why Knowledge? Constraint

Any account of practical knowledge ought not leave it mysterious why the phenomenon it describes is a kind of knowledge.

Neither of the accounts I considered in Chapter Two met all of these constraints. O'Brien met the *Apparent Features* constraint but, because she tied practical knowledge to an agent's conscious awareness of acting, the account was unable to meet the *Special Relationship* constraint. It was hard to see why *only* intentional action should be an object of practical knowledge, and to see how very many of our intentional actions – the non-basic ones - could be. In addition, the view conflicted with the fact that we can know what we are doing even when we are not consciously attending to doing those things; with the fact that practical knowledge, like all propositional knowledge, appears to be *dispositional*. Practical knowledge as we described it in Chapter One seems much more pervasive than O'Brien's picture allowed, and the account was rejected.

Paul's account didn't have the same problems as O'Brien's, but it still had problems. It was not able to meet the *Apparent Features* constraint and not obviously able to meet the *Special Relationship* constraint, although Paul's worries here were different from O'Brien's. Most seriously, I suggested, a dilemma infected Paul's capacity to meet the new *Why Knowledge?* constraint. The problem seemed to be in particular with the idea that judgments about *what one is doing* are – on Paul's view – inferred from judgments about one's intention and abilities. I suggested that the inference would go through only if one of the premises linked intending with acting, but introducing such a linking premise led to the regress problems we considered at the end of §2. I suggested that Paul's account ought to be rejected.

Both O'Brien's and Paul's accounts did have benefits, though, which are best seen by comparing the two, since they mirror one another. Because Paul's account *does not* tie practical knowledge to the phenomenology of acting, and because it linked practical

knowledge's *content* to that of intention, it was easy to see how Paul avoided the kind of problems which O'Brien had with the *Special Relationship* constraint (*viz.* of accounting for the breadth of practical knowledge's scope, and its dispositional character). And because O'Brien's account *does not* think of practical knowledge as *inferred* from evidence of any kind, she was able to avoid the kinds of problem Paul had with meeting the *Apparent Features* constraint (*viz.* of explaining why practical knowledge should appear epistemically ungrounded).

So perhaps what we need is a view which combines the helpful features of both accounts, and dispenses with the problematic ones. I start to consider such a view in Part Two.

Part Two

Cognitivist Intentionalism

In this Part we will start to consider *Intentionalism* – the idea that an agent’s practical knowledge is in some sense constituted by her intention. The considerations in Part One gave us reason to take this idea seriously, by suggesting that we dismiss two major alternative ways of thinking about practical knowledge. But how is the idea supposed to work? The idea is apt to seem strange from the get-go given the common assumption that knowledge is constituted by *belief* (when true and held in epistemically favourable circumstances). How could practical knowledge be constituted by *intention*?

One move is to think of intention as itself a kind of belief. If that is right, then intentions could constitute knowledge in the way beliefs do, and practical knowledge’s special features – those which differentiate it from other kinds of knowledge – could be explained by reference to the differences between intentions and ‘ordinary’ beliefs. I will refer to the idea that intentions are a kind of belief as ‘Cognitivism about Intention’ – *Cognitivism* for short – and I will call a view of practical knowledge which includes this idea *Cognitivist Intentionalism*, or *CI*.

David Velleman and Kieran Setiya both defend versions of *CI* and I will consider both versions in Chapter Three. Both views are found problematic, although Velleman’s rather more so. In Chapter Four I argue that *Cognitivism* itself is untenable, which leads me to reject *any* version of *CI*. This sets the scene for a consideration of the view I will defend in Part Three; one on which practical knowledge is constituted by intention *rather than* by belief.

Chapter Three:

Cognitivist Intentionalism

Introduction

David Velleman and Kieran Setiya both accept that practical knowledge is constituted by intention, and that intentions are a kind of belief. In the language I will be using, they are both *Intentionalists* about practical knowledge, and *Cognitivists* about intention: they both accept *Cognitivist Intentionalism* or *CI*.

Accepting *Cognitivism* means that Velleman and Setiya are able to adopt *Intentionalism* in a way which fits with what I will call the *Standard Metaphysics of Knowledge*, the familiar picture of *knowledge* as constituted by a belief which is true and held in epistemically favourable circumstances. Velleman and Setiya both agree that *practical* knowledge is constituted by a belief of a special kind – an intention-belief (as I will call them) – which is true and held in epistemically favourable circumstances. But they disagree on the nature of these circumstances. Velleman's conception of them is *internalist* whereas Setiya's is *externalist*.

This difference, it turns out, is important. The main point of §1 will be that Velleman inherits the major problems faced by Sarah Paul, for the reason that on point of epistemology (and despite first appearances), the two epistemological pictures are more or less the same. We will see how Setiya's externalism immunises him against these problems in §2, but we will also see that Setiya has problems of his own.

1. Velleman's Internalist Cognitivist Intentionalism

Velleman's account of practical knowledge¹ must be understood from the perspective of a *Causal Theory of Action* (henceforth 'CTA'). CTA comprises an account of the *constitution* and of the *explanation* of intentional action.² Intentional actions are *constituted* by bodily movements which are the causal upshots of the agent's reasons or motivations (understood as mental states), and they are to be *explained* in terms of these reasons, where this explanation is understood as a species of efficient-causal explanation.³

¹ I concentrate on Velleman's view as presented in his *Practical Reflection*, but later work (*The Possibility of Practical Reason*) will become relevant in (d). Velleman is clear that the later work does not represent a substantial change of mind, but a development of the earlier account ("Precis of The Possibility of Practical Reason," 223).

² Velleman thinks of the objects of practical knowledge as *autonomous* or *full-blooded* actions, not *intentional* ones (Velleman, "Precis of The Possibility of Practical Reason," n. 6.), but for reasons noted earlier (Ch. 1, §1), I will proceed as if he gives an account of our knowledge of the latter.

³ Velleman's version of the causal theory is atypical. He describes his argument with standard CTA in his Introduction to *The Possibility of Practical Reason*, 5–12. This argument is not relevant to my discussion in this Chapter.

Velleman's account of practical knowledge is an attempt to respond to two problems. First, he wants to reconcile practical knowledge's *psychological immediacy* – what he calls its 'spontaneity' – with *CTA*, in light of the following seeming tension between the two:

We usually know what we're doing, and we seem to know it quite spontaneously, without having to discover it. We feel as if we're inventing what we do and hence that we don't have to find out. But if our actions are caused by our motives, then they must follow a predetermined course that isn't really ours to invent [...] How, then, can we have spontaneous knowledge of them?⁴

I will adopt Velleman's terminology throughout this section, speaking of practical knowledge as 'spontaneous' rather than 'psychologically immediate', but ask the reader to recall that these labels name the same feature of practical knowledge.

The second problem is to do with the *internal* coherence of the idea of spontaneous knowledge. The worry is that if practical knowledge is spontaneous – not based on a discovery or episode of noticing (etc.) – then it cannot be based on *evidence*. But if an attitude with an empirical content is *not evidentially grounded*, it doesn't deserve the name 'knowledge'. Velleman notes that certain philosophers – he has in mind in particular Elizabeth Anscombe and Stuart Hampshire⁵ – have suggested that practical knowledge is "a unique kind of knowledge that simply doesn't require grounds".⁶ But he thinks that this position is deeply problematic in that it "[makes] epistemological exceptions" in practical knowledge's favour,⁷ by exempting it from "the usual requirements of evidential support".⁸

Any attitude, in order to be knowledge, must "be able to vouch for its truth by means of evidence",⁹ and so knowledge (at least empirical knowledge) *must* be based on evidence for thinking its content true. Anything called 'knowledge' which does not meet this evidential criterion is not "worthy of the name"¹⁰ 'knowledge'. Velleman, then, is mindful of the importance of meeting the *Why Knowledge?* constraint, and of potential conflicts between doing so and adequately explaining practical knowledge's *psychological immediacy* (and so meeting the *Apparent Features* constraint).

His response to both problems looks like this: far from being in conflict with the causal aetiology of intentional action, practical knowledge itself plays an ineliminable role *within* it. Practical knowledge is *spontaneous* because it *causes* the actions it concerns. And despite its spontaneity, practical knowledge is nonetheless *evidentially grounded* because the beliefs which constitute it stand as evidence for their own truth. Velleman's response to both problems rests on his identification of *intentions* as a kind of belief. Practical knowledge plays a role in the causal aetiology of intentional action in virtue of the fact that its constituting attitudes – intentions-beliefs – do so. And the beliefs which constitute

⁴ Velleman, *Practical Reflection*, 4. For a similar worry see Roessler and Eilan, "Agency and Self-Awareness: Mechanisms and Epistemology," 20–21.

⁵ Velleman, *Practical Reflection*, 25. Velleman specifically cites Hampshire and Hart, "Decision, Intention and Certainty" here, adding that he "would list Anscombe, Grice, Kenny, and Melden as agreeing, or likely to agree with Hampshire and Hart." (Ibid, n.14)

⁶ Velleman, *Practical Reflection*, 25.

⁷ Ibid. A similar worry is found in O'Brien, *Self-Knowing Agents*, 178.

⁸ Velleman, *Practical Reflection*, 105.

⁹ Ibid.

¹⁰ Ibid., 25.

practical knowledge are *self-evidencing* because they are intention-beliefs, which unlike ordinary beliefs, tend to *cause their own truth*.

I will explain these ideas in more detail in (a), and then explain how the account meets two of our constraints in (b). In (c) and (d), though, I will argue that Velleman's account is no real improvement on Paul's in relation to the *Apparent Features* and *Why Knowledge?* constraints. I will conclude in (d) that Velleman's account has significant internal difficulties and ought to be rejected.

a. Velleman on Practical Knowledge

In outline, Velleman's view is as follows: Intentions are a kind of belief, and practical knowledge is constituted by these intention-beliefs. Specifically, intentions are beliefs of the form "I will φ " or "I am φ -ing",¹¹ but they differ from ordinary beliefs in two important ways. First, they are formed in response to the agent's *motivations* for acting as intended (practical reasons), rather than in response to sufficient prior *evidence* for thinking that one will act as intended (theoretical reasons). My intention-belief that I will go to the bar, for example, is not formed in response to reasons which make my going to the bar look *likely* (if my bar-going already looked likely it's hard to see what point there would be in forming the intention to go there), but in response to reasons I have *for going there* (I want a pint). It is because the attitudes constituting practical knowledge are not formed in response to evidence for their truth that practical knowledge does not represent a discovery; that it is *spontaneous*.

Second, intention-beliefs are *self-fulfilling*. Since intentions (this is where CTA comes in) tend to cause the actions they concern, they tend to cause their own truth. When my intention-belief that I will go to the bar causes me to go to the bar, my intention-belief turns out true. Not all intention-beliefs cause their own truth (sometimes we don't do what we intend), but as a category they have a tendency to do so. This is a *causal* property of intention-beliefs, but it underwrites an *epistemic* property: because intention-beliefs are *self-fulfilling*, they are also *self-evidencing* or *self-justifying* (I will use these interchangeably). If my intention-belief that I will go to the bar is likely to *cause me* to go to the bar, then my having this intention-belief is evidence that I *will go to the bar*, so is evidence for its own truth. It is because the attitudes constituting practical knowledge are *self-evidencing* that they – if true – meet "the usual requirements of evidential support"¹² and so count a kind of *knowledge*.

But why *are* intention-beliefs self-fulfilling; what is the *mechanism* here? The answer is that agents capable of intentional action come pre-packaged with a desire to know what they are doing.¹³ This desire *inhibits* an agent from doing things she does not believe she will do, and *prompts* her to do things she *does* believe she will do. Having an intention-belief that I will go to the bar thus *adds to my antecedent motivations* for going there and, *ceteris paribus*, causes me to do so. Let's consider the idea in more detail by going through another example.

¹¹ In fact Velleman thinks that intention-beliefs have more complex contents than this. The details don't matter at the moment but we will come to them in (d).

¹² Velleman, *Practical Reflection*, 105.

¹³ *Ibid.*, 47, *et passim*. In later work Velleman alters the suggestion so as to allow that one's aim to know what one is doing is determined not by a desire, but by some sub-personal mechanism. His reasons are given in *The Possibility of Practical Reason*, 12–20. Throughout, Velleman thinks of it as constitutive of being an agent that one aims to know what one is doing.

I am working with the radio on quietly in the background. My concentration is flagging and I am likely to get distracted if I leave it on, because an interesting programme is about to start which would likely monopolise my attention. I had wanted to finish this section before I took a break, so perhaps I should turn the radio off. On the other hand, perhaps listening to something interesting for half an hour might set me up for better concentration later; so perhaps I should take my break now, and turn the radio up so I can hear it better. I have some motivation for turning the radio *off*, and some for turning it *up*; neither set of motivations are overwhelming in strength. What to do? In the end I decide to turn the radio *off*, and go on to do so.

Here is Velleman's *causal-motivational* story: because my motivations to turn the radio *off* and to turn it *up* were roughly equal in strength I was not compelled to do either, and prior to deciding what to do I was (more or less) *just as likely* to do either. When I decide to turn the radio *off* I do so on the basis of my *motivations* for turning it *off* (had I decided to turn it *up* it would have been on the basis of my motivations for doing this instead). Deciding is forming an intention-belief, in this case with the content *that I will turn the radio off*. Because I now have a belief (an intention-belief) that I will turn the radio off, my desire to know what I am doing kicks in, inhibiting me from doing anything I *don't* believe I will do (including turning the radio *up*) and prompting me to do what I *do* believe I will do, *viz.* to turn the radio *off*. The balance of my motivations has been shifted; I *now* have stronger motivation for turning the radio off than for turning it up. I am caused, by these motivations, to turn the radio off.

Here is Velleman's *epistemological* story: because I recognised that my motivations to turn the radio *off* and to turn it *up* were roughly equal in strength, my evidence that I would turn the radio off was, and appeared to me to be, no stronger than my evidence that I would turn it up (and *vice versa*). So I didn't have sufficient reason to believe I would either turn the radio up, or turn it off. Nevertheless I can form an intention-belief on the basis of one set of motivations, and I in fact decide to turn the radio *off*, forming an intention-belief that I will do so. Intention-beliefs are self-justifying (because they are self-fulfilling), so I now have a *justified* intention-belief, one which 'meets the usual requirements of evidential support'. All that needs to be the case for me to thereby have *knowledge* is for the intention-belief to turn out *true*. And it does (*ex hypothesi*): I turn off the radio (for reasons explained in the previous paragraph).

What exactly do I have practical knowledge of in this example? That I *will turn off* the radio, or – a bit later – that I *am turning* it off? Jonathan Dancy¹⁴ worries that Velleman confuses *present-directed* with *future-directed* knowledge of intentional action. Velleman clarifies things as follows:

In the first instance, [...] an agent aims at the former. But [...] I also think that an agent's primary means of attaining the former is by attaining a short-term version of the latter.¹⁵

Intention-beliefs are typically formed prior to performing the action in question, but because an agent knows, in virtue of her intention-belief, that she *will* act in a certain way, this means that when she comes to *acting* in that way, her action comes as no surprise and she knows that she *is* performing it. I simply assume here that Velleman's strategy of

¹⁴ Dancy, "On Knowing What One Is Doing," 239–240.

¹⁵ Velleman, "Replies to Discussion on the Possibility of Practical Reason," 277. Velleman is responding directly to Dancy here.

understanding present-directed practical knowledge in terms of a ‘short-term version’ of future-directed practical knowledge is unproblematic.

Let’s sum up by clarifying how Velleman explains practical knowledge’s spontaneity consistently with meeting the *Why Knowledge?* constraint. The formation of intention-beliefs is not constrained by prior evidence for their truth in the way that the formation of ordinary beliefs is; this is why practical knowledge is *spontaneous* and does not represent a discovery on the basis of evidence. But *once formed*, intention-beliefs are evidentially grounded; they constitute evidence for themselves. And so the practical knowledge they constitute is – like empirical knowledge in general – evidentially grounded. This is how, thinks Velleman, we can maintain that practical knowledge is ‘worthy of the name’. Velleman’s trick has been to separate the *psychology* of practical knowledge from its *epistemology*. Its spontaneity is a feature of its psychology; the intention-beliefs which constitute it are not *formed in response to* or *occasioned by* sufficient evidence for their truth. But this does not mean that on point of *epistemology* they are not evidentially grounded. Intention-beliefs themselves give an agent sufficient evidence for thinking them true. In the terminology from Chapter One, we can think of Velleman as pulling practical knowledge’s *apparent epistemic ungroundedness* apart from its *apparent psychological immediacy*, seeking to *explain away* the former appearance, and to *accommodate* the latter, motivated by a recognition of the importance of meeting the *Why Knowledge?* constraint.

Velleman’s *Cognitivism* plays a central role in his account. It is because intentions are a *kind of belief* that they are capable of constituting knowledge, which is understood according to Standard Metaphysics of Knowledge as constituted by belief which is true and held in epistemically favourable circumstances. And it is because they are a *special* kind of belief that first, they can be formed without sufficient prior evidence (they are ‘spontaneous’ and practical knowledge is *psychologically immediate*) and second, they are nevertheless justified (they are self-justifying, and so if true will constitute *knowledge*). Because intention-beliefs are self-justifying they need not be *inferred* from evidence in order to be *grounded in* evidence.

b. The Special Relationship Constraint

Let’s start to assess Velleman’s view. Like Paul, he ties the content of practical knowledge to the content of intention, and so like Paul would seem to avoid O’Brien’s scope difficulty. Indeed, we can assume that *any* version of *Intentionalism* is bound to do as well as Paul did on this score.

Paul had a different problem with the *Special Relationship* constraint, *viz.* that it seemed brutally empirically possible on her view for an agent to be capable of intentional action yet to lack the capacity for practical knowledge (see again Ch. 2, §2(a)). Velleman seems to improve on Paul’s account here. For a case in which someone acts intentionally will be a case in which their action is caused in part by their intention-belief, and this means that anyone acting intentionally has a true intention-belief that they are acting. And since intention-beliefs are *self-justifying*, anyone acting intentionally will have a true *and justified* intention-belief, which means that they’ll have knowledge that they are acting as intended. So on Velleman’s view, acting intentionally either entails or comes very close to entailing practical knowledge, and the *Special Relationship* constraint seems to be met.

But he doesn’t do so well on our two other constraints, as we will see.

c. The Apparent Features Constraint

Velleman initially appears to do well on the *Apparent Features* constraint, since he happily accommodates both practical knowledge's *psychological immediacy* ('spontaneity') and its *FPA*. We have already seen (in (a)) how he explained practical knowledge's *psychological immediacy* and I won't bring up any problems for this explanation. Let's consider why, on Velleman's view, practical knowledge should display *FPA*.

A new bit of terminology is useful here. For Velleman, practical knowledge is "the cause of what it understands".¹⁶ This characterisation picks up on the fact that, as described above, practical knowledge's constituting attitudes (intention-beliefs) play a role in the causal aetiology of the *objects* of that knowledge (the intended actions): my intention-belief that I will turn off the radio causes me to turn it off.¹⁷ The idea that practical knowledge is 'the cause of what it understands' underwrites a first-/third-person asymmetry in relation to propositions about what someone is doing intentionally, because only a person's *own* beliefs (their *intention-beliefs*) have the relevant causal role. So Velleman seems to accommodate practical knowledge's *FPA* and to explain exactly what it amounts to, which was a requirement of meeting the *Apparent Features* constraint (see again Ch. 1, §2(d)).

It is with practical knowledge's apparent *epistemic ungroundedness* that Velleman has a problem. Indeed, it is the same problem I suggested Paul faced. Like Paul, Velleman must explain away this apparent feature of practical knowledge because like Paul, Velleman *denies* that practical knowledge is epistemically ungrounded: it is, he thinks, grounded on the evidence it provides for itself and *only* because it is, is it 'worthy of the name' "knowledge".

On the face of it, Velleman has a story to tell which was not available to Paul. For on his view, practical knowledge is not *inferred from* evidence, even though it is epistemically grounded in evidence. So perhaps I am at a loss to answer the question "How do you know you are turning off the radio?" because I hear it at a request for the evidence *in response to which* I formed my intention-belief, and I recognise that I didn't form my intention-belief in response to any evidence, but in response to my motivations.

But this response is problematic. For there is no reason to think that an agent can *only* hear "How do you know you are φ -ing?" as a request for information about the aetiology of her intention-belief or knowledge. If Velleman is right that knowledge *per se* requires evidential grounds, then in seeing myself as *knowing* that I am turning off the radio, I must seemingly either recognise my judgment as epistemically grounded, *or* flout some rationality requirement. We can stipulate that I can think of myself as knowing that I am turning off the radio, consistently with being rational, and ask: why *wouldn't* I hear "How do you know you are turning off the radio?" as a request for my evidential grounds?

A pragmatic response is available: I won't hear the question in this way because *everybody knows* that intention-beliefs are *self-evidencing*. Because this is common knowledge and if it is also common knowledge that I am *intentionally* turning off the radio, then there is no point in my interlocutor asking for my evidence; the involvement of my intention ensures that the *answer* to their question is already mutually known. This would explain

¹⁶ Velleman, *Practical Reflection*, 102.

¹⁷ We shall see later (Ch.5; §1(e)) that this is not the only way to understand the claim that practical knowledge is 'the cause of what it understands'.

both why I don't hear the question as a request for my evidence *and* why it seems like a stupid or strange question to ask.

One worry about this response is that it seems to depend on thinking of ordinary agents as at least implicitly accepting Velleman's unusual metaphysics of practical knowledge; as accepting that practical knowledge is constituted by a kind of belief which unlike ordinary beliefs is self-justifying and can be formed in the absence of reasons for thinking it true. I will complain about this conception of intention directly in (d), but I will put this worry to one side here in order to concentrate on a bigger problem with the pragmatic strategy.

The pragmatic response it isn't a satisfactory answer to our question, for my interlocutor can *force* me to hear the question as request for my knowledge's epistemic grounds, rather than for its psychological aetiology. She can ask me directly, "*What is your justification* for believing that you are turning off the radio?" or "*On what evidence* do you believe that you are turning off the radio?" But now Velleman is in just the same boat as was Paul (see again Ch. 2, §2(a)); it just doesn't seem to me as if my knowledge that I am turning off the radio *is* evidentially grounded. From my perspective, my intention to be turning it off just doesn't seem to constitute *evidence* that I am, *even if* I accept that my intending to do things usually leads to my doing them.

Interestingly, Velleman accepts this datum:

To the question "How do you know what you're doing?" the most obvious reply is "Well, I'm the one who's doing it!" – which suggests that doing something is a special way of knowing about it.¹⁸

I find what Velleman says in this passage very plausible (and will provide my own account of how doing something is a special way of knowing about it in Chapter Six). The problem for Velleman is that his account doesn't do his own observation justice.¹⁹

Like Paul, Velleman thinks that practical knowledge is *evidentially based* on the agent's intention. And like Paul, Velleman fails to explain why this evidential basis should be opaque to the agent. The fact that Velleman denies that practical knowledge is *inferred from* evidence might explain why it should appear to be *psychologically immediate* ('spontaneous') but it does not adequately explain away practical knowledge's apparent *epistemic* – in particular, *evidential - ungroundedness*. Velleman does no better than Paul on the *Apparent Features* constraint.

d. The Why Knowledge? Constraint

I have already remarked that Velleman is mindful of the need to meet the *Why Knowledge?* constraint; to give an account of practical knowledge on which it is not mysterious why it really is a kind of *knowledge*. Velleman's worry about meeting this constraint is a worry about how practical knowledge could be *spontaneous*, consistently with its being *evidentially grounded*. And despite the fact that (as I argued in (c)) he doesn't sufficiently explain away the *appearance* of epistemic ungroundedness, he does have a story to tell about practical knowledge's evidential basis.

¹⁸ Velleman, *Practical Reflection*, 18.

¹⁹ A similar objection can be found in Langton, "Intention as Faith," 258.

But in order to work, his story about practical knowledge's grounds must be *complete*. My complaint against Paul's account was that *her* story about the agent's epistemic grounds was *not* complete, because either her premises failed to justify a conclusion about what the person is *actually doing*, or – on a bolstered version of the Pauline inference – because her premises depended essentially on prior practical knowledge, leaving Paul's account of practical knowledge's grounds ultimately unexplained. For this reason, I argued, Paul flouted the *Why Knowledge?* constraint. Again, I don't think Velleman improves on Paul's account in relation to this issue. For although their accounts look very different, this difference is down to a different conception of the *psychology* of practical knowledge, and not a different conception of its *epistemology*. On point of epistemology Velleman's account pretty much coincides with Paul's bolstered account. But the problem for Paul's bolstered account was a problem with its *epistemology*.

In (a) we saw how thinking of intention-beliefs as *self-fulfilling* meant seeing them as *self-justifying*, but Velleman's epistemology is more complicated than I suggested there. Agents' intention-beliefs are not justified merely because they *are* self-fulfilling, but because, in addition, agents *believe* that they are:

... the agent isn't entitled to expect himself to do something unless he has reason to believe that he'll do it if he expects to.²⁰

And such a person *does* have a reason to believe that he'll do what he expects (intend-believes): a feature of Velleman's account I have so far suppressed is that intention-beliefs not only *are* self-fulfilling, they also *represent themselves* as such. This, for Velleman, is part of the *content* of intentions; in the abstract, intentions have the content "I'll φ because I hereby intend to."²¹ The 'because' represents one's intention-belief as likely to cause the intended action, whereas the 'hereby' indicates that the intention which will do the causing is *this very one*. Both the 'because' clause and the 'hereby' clause are within the scope of one's belief. My intending *to turn off the radio* thus involves my believing:

- a) That I intend to turn off the radio
- b) That this intention will or is likely to cause me to turn off the radio
- c) That I will turn off the radio

For Velleman, these 'three' beliefs are aspects of the same attitude, so unlike Paul, Velleman does not think I *infer* from (a) and (b) to (c). Still, on point of *epistemology* the similarity with the bolstered Pauline account is glaring. Both Paul and Velleman think that a belief of the form "I will φ " is *epistemically grounded* on an evidence-set which includes a belief of the form "I intend to φ " and a belief about how intention typically links to action. The belief in (b) is about a *particular* (singular) causal relationship *this* intention-belief is likely to enter into, but it must, presumably, be justified by a generalisation linking intentions-to- φ with φ -ings. With this in mind, Velleman's and Paul's *epistemologies* turn out to look pretty much identical: for both, part of what evidentially grounds an agent's practical knowledge is her knowledge that she *tends to do*, or to be doing, what she *intends to do*, or to be doing.

²⁰ Velleman, *Practical Reflection*, 97. For 'expectation' read 'intention-belief': In *Practical Reflection* Velleman first characterizes practical knowledge's constituting attitudes as 'self-fulfilling expectations'; later in the book he identifies these with intentions (Velleman, *Practical Reflection*, Ch. 3 & 4).

²¹ Ibid., 140; my square brackets.

Does Velleman fall into the same regress as did Paul (see again Ch. 2, §2(c))? Yes, I think so, although it takes a little longer to see why, for an option is open to Velleman which was not open to Paul. Paul was explicit that agents come to know the link between intention and action ‘from experience as an agent’, and this committed her to thinking of this knowledge as inductively generalised from past instances of conjunctive knowledge of intention and of action. It was the thought that the relevant knowledge was justified by an inductive generalisation which caused the regress problem. But Velleman need not accept that the knowledge of the link between intention and action is justified in this way. There is another option available to Velleman; he can ‘go functionalist’.

What intentions *are*, says the functionalist, is completely determined by their functional role, part of which is to cause actions which match their contents. An isomorphic account of concept-possession has it that someone possessing the concept of *intention* just is someone who recognises intention’s distinct functional role as characterising certain psychological attitudes. On this view, an attitude’s functional role is – as I will say – *transparent* to someone possessing the concept of the attitude. Because it is part of intention’s functional role that they tend to cause the actions they represent, this will be transparent to anyone possessing the concept of *intention*. This would explain how an agent could know that she tends to do what she intends, *without* having inductively generalised from past conjunctive beliefs of the form “I intend to be φ -ing and I am φ -ing”, and so without presupposing prior practical knowledge in the way Paul did. On the functionalist view, the agent knows this in just virtue of possessing the concept of *intention*.

I’ll accept for the sake of argument that Paul’s regress is avoided by the functionalist strategy.²² The problem is that the strategy is not available to Velleman. The self-fulfilling nature of intentions is only one aspect of intention’s functional role, and it is implausible that its other aspects, as Velleman understands them, are transparent to anyone possessing the concept of *intention*. Because the functionalist cannot pick and choose which aspects of an attitude’s functional role will be transparent to someone possessing the concept of the attitude, Velleman must *choose between* his complete conception of intention’s functional role, and the functionalist account of concept-possession. Either way his account of practical knowledge falls down, as we shall see.

What are the other aspects of intention’s functional role, and why is it implausible that they are transparent to someone possessing the concept of *intention*? We have already seen that for Velleman, intentions are a special kind of belief. In more detail, he thinks of ‘ordinary’ beliefs as having three jointly defining features: a particular *direction of guidance*, *constitutive aim*, and *direction of fit*. An attitude’s *direction of guidance* is a matter of whether it tends to cause, or to be caused by what it represents. Ordinary beliefs *tend to be caused by* what they represent, but intentions, recall, are *self-fulfilling* which means that they tend to *cause* what they represent. So intention’s direction of guidance differs from that of ordinary belief. But intentions are nevertheless a *kind of belief* because they share a *constitutive aim* and a *direction of fit* with ordinary beliefs. I want to suggest that it doesn’t

²² Could Paul herself not step back from her commitment to agents knowing the linking premise ‘I tend to do what I intend’, and adopt the functionalist response? Not, I don’t think, without stepping back from her inferentialist strategy itself. For the Functionalist response requires thinking of a belief of the form “I intend to be φ -ing” and a belief of the form “I am φ -ing” as somehow internally related to one another. But this would collapse the logical gap between premises and conclusion which is a feature of *inductive* inference. Would the resulting view nevertheless be a good one? I can’t address the altered Pauline view in detail here, but it seems unlikely to do any better than her view (as considered in Chapter Two) in relation to the *Apparent Features* constraint.

seem plausible to think that anyone possessing the concept of *intention* recognises them as having these features, even implicitly.

Belief's *constitutive aim* is *truth*: an attitude is a belief only insofar as it aims to be true. Intentions are supposed by Velleman to share belief's constitutive aim because:

When I form [an intention] I aim to settle a question in my mind only insofar as I can thereby settle it in fact; I aim, that is, to avoid representing an arrangement that I am not thereby managing to make.²³

My intention to turn off the radio, thinks Velleman, 'aims at being true' in the sense that its *point* is to cause me to turn off the radio, a condition which would render its content – *that I will turn off the radio* – true. Any attitude which did not have this aim would not be an intention, thinks Velleman, or a belief. (Intentions are unlike ordinary beliefs in that they tend to meet the aim off their own bat; they are self-fulfilling.)

If intentions constitutively aim at truth, they must be *capable* of being true. But it's not obvious that commonsense psychology thinks of intentions as ever being true (or false). We certainly don't *say* that people's intentions are – like their beliefs – true (or false); we say that they are (competently) carried out, or not. In common parlance, talk of 'true *beliefs*' is usually a semantic notion, whereas talk of 'true *intentions*' is usually a moral one. So it is not obvious that this feature of intention's functional role – as Velleman conceives it – is transparent to people possessing the content of *intention*.

It might be responded that the fact that we don't speak of intentions as (semantically) true or false doesn't entail that they are *not*. Perhaps we don't think of intentions as having true (false) contents because they don't *represent* their propositional contents *as* true, as do beliefs. Instead they represent their propositional contents *as to be made* true. But this doesn't mean, the response continues, that intentions' contents *are not* true or false, so it doesn't mean that intentions can't aim at truth.

But Velleman can't give this response, for it requires *distinguishing* intention's *direction of fit* from that of belief, and Velleman thinks that intention and belief *share* a direction of fit:

Because choosing [i.e. *forming an intention*] entails settling a question in one's mind, it requires more than representing an answer *to be* arranged. If it were still *to be* arranged that I was going to act, then it would not yet be settled that I was going to act [...]. Settling on a future action thus requires representing the action as arranged: my choice makes it true that I am going to act, by representing it as true that I'm going to act. It therefore has the same direction of fit as belief.²⁴

This point is not *ad hominem*, but goes to the heart of Velleman's epistemology of practical knowledge, for we have already seen that Velleman's commitment to *Cognitivism* is central to his claim that intentions can constitute practical knowledge. And his commitment to *Cognitivism* just is his commitment to the claim that intentions share two out of the three defining features with 'ordinary' beliefs; that intentions are unlike ordinary beliefs *only* in relation to their direction of guidance.

²³ Velleman, *The Possibility of Practical Reason*, 25.

²⁴ Ibid. See also Ibid., 10.

There is another problem for the claim that intentions share a constitutive aim with beliefs, too; at least on Velleman's understanding. Velleman glosses the aim-claim in terms of correctness conditions:

The norm of correctness for belief is [...] internal to the nature of belief itself. The concept of belief just is the concept of an attitude for which there is such a thing as correctness and incorrectness, consisting in truth or falsity. [...] Philosophers have traditionally accounted for this feature of belief by saying that beliefs constitutively aim at the truth.²⁵

Assuming intentions have propositional, truth-valued, contents, it is still not obvious that commonsense psychology represents them as incorrect insofar as they fail to cause facts matching these contents. Assume I have an intention whose content is *that I am typing 'causal'*, but that I am actually writing 'casual', so that my intention's content is false. What has gone wrong? On Velleman's view, my intention is incorrect because it is false, in just the same way that Bev's belief that the pangolin is prospering is incorrect because it is false (the pangolin is endangered).

But we don't think of things in this way. If I recognise that I am typing 'casual' and not 'causal' I will – as Anscombe pointed out²⁶ – correct *my action* and not *my intention*. Whereas if Bev recognises that the weight of evidence shows that pangolins are being hunted to extinction, she will *correct her belief*. If she doesn't she's irrational. But I am not irrational for maintaining my intention to type 'causal' when I realise I am not currently doing so. It is not rational to *give up on* what one is writing just because of making a typo. Nothing would ever get written. So again, possessing the concept of *intention* doesn't seem to entail thinking of intentions as sharing a *constitutive aim* with beliefs.

I take no stand here on whether Velleman is right that intentions share a constitutive aim and direction of fit with – and so are a kind of – belief (though I will argue against *Cognitivism* from a different direction in Chapter Four). But *if they do*, then this does not seem to be transparent to people in possession of the concept of *intention*. So intention's functional role is not transparent to possessors of the concept. So the functionalist way out of Paul's regress problem is not available to Velleman, and he has to fall back on understanding the "I tend to do what I intend" premise as constituting an inductive generalisation from past practical knowledge. And so – like Paul – his account of practical knowledge's grounds is incomplete and he flouts the *Why Knowledge?* constraint.

And I take no stand on whether we *should* be functionalists about concept-possession. But *if we should* then we have to reject Velleman's *Cognitivism*, since – again – the functional role which intention *appears to us* to have is not the one Velleman describes. But Velleman cannot give up his *Cognitivism*, because as I explained above (in (a) and (b)) it is central to his explanation of how practical knowledge's spontaneity is consistent with its status as a kind of knowledge.²⁷

²⁵ Velleman, *The Possibility of Practical Reason*, 16.

²⁶ E.g. in Anscombe, *Intention*. These issues will be explored in detail in Ch. 5.

²⁷ And to his explanation of how practical knowledge's spontaneity is consistent with *CTA*, accepting which Velleman seems to think non-negotiable (see e.g. *Practical Reflection*, 3).

e. Conclusion

At the end of the last Chapter I suggested we might expect a view on which practical knowledge is *constituted by* intention to fare better than one on which it is *inferred from* knowledge of one's intention, and Velleman's *internalist CI* offered just this. But his account ended up being more or less the same as Paul's on point of epistemology: on both accounts knowing what one is doing depends on knowing what one intends and knowing that (roughly) intending to do something makes doing it likely. Because his account of the epistemology of practical knowledge is similar to Paul's, Velleman faces the two main problems faced by Paul's account, the first of which is about practical knowledge's *epistemic appearance* (as ungrounded) and the second of which is about its *epistemic grounds*.

So it is worth considering a version of *CI* which involves a different understanding of practical knowledge's epistemological underpinnings. Kieran Setiya offers an *externalist* version of *CI*, to which I now turn.

2. Setiya's Externalist Cognitivist Intentionalism

The main problem for Velleman and Paul – that of a potential regress within practical knowledge's justification – depended on the presupposition that for an agent's tendency to do what she intends to justify her belief that she was φ -ing (her intention to be doing so), the agent must *believe* that she tends to do what she intends. It was not enough that she *has* a disposition to φ when she intends to, she must in addition *recognise* this about herself.

Kieran Setiya dispenses with this presupposition. On his view merely *being* disposed to do what one intends justifies an agent in believing that she will φ when she intends to. On this *externalist* version of *CI*, the agent doesn't need to know that she has this disposition, and so questions about *how* she knows it just don't come up, which means that the completeness of Setiya's account of practical knowledge's epistemology is not threatened by regress.

Setiya characterises himself as a Cognitivist, but unlike Velleman he offers no clear or settled view about how intention and belief relate, in different places suggesting that one's belief that one is φ -ing is *constituted by* one's intention to φ ;²⁸ that forming an intention is *forming* a belief;²⁹ that beliefs *figure in* our intentions;³⁰ that intentions *have corresponding* beliefs;³¹ that intention is *a species of* belief;³² and that intending is a *mode of* believing.³³ At least some of these characterisations of the relationship conflict with one another, and at least some of his works contain more than one characterisation. In light of this interpretative difficulty I will treat Setiya as taking over Velleman's conception of the relationship between intention and belief. This seems fair given that Setiya introduces

²⁸ Setiya, "Practical Knowledge," 396.

²⁹ *Ibid.*, 397.

³⁰ *Ibid.*, 402.

³¹ *Ibid.*, 407.

³² Setiya, "Practical Knowledge Revisited," 131.

³³ Setiya, "Knowing How," 303.

his conception of practical knowledge as a response to Velleman's, from within the Cognitivist viewpoint, nowhere criticising Velleman's version of Cognitivism.³⁴

I will spell out Setiya's externalist version of *CI* in (a) and consider how well it meets our constraints in (b). I will conclude in (c) that although Setiya's account represents an improvement on Velleman's, his commitment to *Cognitivism* is not obviously defensible.

a. An Externalist Cognitivist Intentionalism

How do I know that I am currently (intentionally) making Turkish coffee, according to Setiya? First, because I am making it intentionally, I am doing it in execution of my *intention* to be making it. And my intention to be making it is a kind of belief that I am making it (Cognitivism). So I have a belief that I am making it, which is true, because I *am making* it. And second, my belief is – externalistically – *justified*, because I am disposed to do what I intend, which means that not only is my belief *true* but it is *non-accidentally* true.

Rather generally, *dispositions* underwrite non-accidental relations between conditions and outcomes. This is reflected in the fact that citing a disposition can be a way of *explaining* some fact or event. “Why did that liquid start to boil?” - - “Because it is water, and its temperature reached 100°C, and water boils (is disposed to boil) at 100°C”. When we explain a phenomenon by saying “Because...”, we represent what follows the ‘because’ as non-accidentally related to the phenomenon to be explained. This is a conceptual point linking the concepts of *explanation* and *non-accidentality*.

Thought of in one way, giving the *justification* for a belief is *explaining why* that belief constitutes knowledge (or: would constitute knowledge if it were true). So we can see why a belief's being non-accidentally true is relevant to its constituting *knowledge*. Setiya can be thought of as trying to explain how my intention-belief that I am making Turkish coffee counts as knowledge, if true, by describing the feature of it which makes its truth *non-accidental* to the fact that I *have* it. I am disposed to do what I intend-believe that I will do. And this means that when I *do in fact* do what I intend-believe I will do, my intention-belief is non-accidentally true, and so knowledge. Setiya dispenses with Velleman's and Paul's assumption that in order for my disposition to do what I intend to play a role in the justification of my intention-belief, it must be something I *recognise*. Instead, the disposition just needs to be one I *have*.

Two clarifications are in order before we turn to assessing Setiya's account. First, his view is *not* that because I am disposed to do what I intend, my intention-belief that I am making Turkish coffee is *grounded on evidence*. Setiya *disagrees* with Paul and Velleman that practical knowledge is evidentially grounded in any sense. On this score, Setiya's view has more in common with O'Brien's: they both agree that agents are *entitled* to make judgments about what they are intentionally doing without this meaning that their judgments are “justifiable with reference to reasons statable by, or comprehensible to, the knowing subject”.³⁵ The disagreement between O'Brien and Setiya is over what condition of the agent entitles her to judgments about her intentional actions. For O'Brien it is the agent's agent-awareness; for Setiya it is their disposition to do what they intend.

³⁴ Setiya also accepts some version of *CTA* (see e.g. Setiya, “Explaining Action,” 380), but it does not play an important role in his account of practical knowledge as it did in Velleman's.

³⁵ O'Brien, “Self-Knowledge, Agency, and Force,” 588.

Second, to avoid confusion I want to explain why I have chosen not to characterise Setiya's view in the way that he himself characterises it; as one on which,

... it is *knowledge how* that explains the otherwise troubling possibility of knowledge in intentional action and of the knowledge embodied in prospective intention.³⁶

Setiya gives at least two glosses on *knowing how to φ* in his early papers on practical knowledge. He starts by thinking of knowing how to φ as (roughly) the *capacity (ability)* to φ in execution of one's intention to φ ,³⁷ later moving to think of it as the *disposition* to φ when one intends to φ .³⁸ These two understandings cannot be identified, since - rather generally - having a disposition to do something presupposes having a capacity to do it, but not *vice versa*. So Setiya must choose between the disposition-characterisation and the capacity-characterisation of *knowing how to φ* .

And it is clear that Setiya should choose the *disposition*-understanding. For whereas being *disposed* to do what one intends ensures that *if* one acts as intended, one's intention-beliefs is non-accidentally true, we cannot say the same for the agent's *capacity* to do what she intends. I am *capable* of executing my intention to do the washing up but this doesn't bear on whether I will, given that I intend to, in the same way that the fact that I am *disposed* to wash up when I intend to bears on whether I will, given that I intend to. (The choice between a disposition- and a capacity-reading of Setiya's epistemology mirrors the choice between Sarah Paul's initial inference-schema and the bolstered version in Chapter Two (§2(c)).)

To introduce more confusion, 'knowing how' to do something is usually understood as referring to what Ryle had in mind (see again Ch. 1, §1), and Setiya himself encourages this use.³⁹ But Ryle seemed to be talking about a *capacity*, and not a disposition! The upshot is that it is confusing to talk of *knowing how to do something* as underwriting one's knowledge that one is doing it. So this is why I avoid representing Setiya's view as one on which it is *know-how* which underwrites practical knowledge, and describe it as one on which a disposition to do what one intends acts as external justification for the agent's intention-belief. Let's see how the view fares.

b. Assessing Setiya's Externalist CI

I have already explained why Setiya's externalism means that he overcomes the problem which led both Paul and Velleman to flout the *Why Knowledge?* constraint. Like Paul and Velleman he also avoids O'Brien's scope-related difficulty with meeting the *Special Relationship* constraint.⁴⁰

What about the problem Paul and Velleman had with meeting the *Apparent Features* constraint? Neither could explain why, given that (on their views) practical knowledge *is* epistemically grounded on evidence, it does not *appear* to be. Again, Setiya's externalism is helpful here. On his view - as on O'Brien's - there is no reason why my justification

³⁶ Setiya, "Practical Knowledge," 388; my italics.

³⁷ Ibid., 407.

³⁸ Setiya, "Practical Knowledge Revisited," 135.

³⁹ Setiya, "Practical Knowledge," 403.

⁴⁰ Setiya is mindful of the importance of getting the scope of practical knowledge right (ibid., 394), and also of needing to accommodate what I have called the *dispositionality* of practical knowledge (Setiya, "Knowledge of Intention," 171).

should be transparent to me, because in general being entitled to a belief by some external condition does not involve being *aware* of one's being so entitled.

Setiya can give a similar account of practical knowledge's *FPA* as did Velleman: if one is acting intentionally it is because one is executing one's intention-belief. This intention-belief also represents it as being the case *that one is acting* as intended, which means that it is *true*. And because by executing one's intention-belief, one is thereby manifesting one's disposition to do what one intends, one's intention-belief is *non-accidentally* true. Only *one's own* intention-belief could be non-accidentally true of one's own action in this special way, and this explains the first-/third-person asymmetry in relation to facts about what one is intentionally doing.

Setiya's account is a clear improvement on Velleman's (and Paul's), but is not without its problems. I will mention two. First, Setiya does not say enough about how the epistemology of practical knowledge is supposed to work and second, he fails to improve on a seemingly problematic aspect of Velleman's account which he himself brings up, and which I passed over in the previous section.

As I explained in (a), Setiya thinks that the fact that when an intention-belief is true, it is non-accidentally so, is relevant to the question of the agent's *knowledge* of her intentional action. But he holds back from endorsing "an epistemology on which its being no accident that a belief is true *suffices* for it to count as knowledge".⁴¹

Setiya does not explain why he is wary of endorsing such an epistemology, and I won't second-guess his wariness here. I just want to remark that *if* he rejects such an epistemology, then his account of practical knowledge is incomplete since it renders the requirement that the agent's belief be non-accidentally true merely *necessary* for practical knowledge. So his explanation of how this condition is met – by the agent's being disposed to do what she intends – is only a *partial* explanation of what is involved in having practical knowledge. *Unless* Setiya accepts that non-accidentally true belief is sufficient for knowledge, then he flouts the *Why Knowledge?* constraint, at least pending a description of what more is needed for knowledge. My complaint here is just that Setiya is unclear about the status of his account.

The second worry is less *ad hominem* and concerns whether Setiya's account is consistent with our commonsense-human approach. Start by considering an objection Setiya makes to Velleman's account. Setiya thinks that even if Velleman can explain why an intention-belief will be justified *once formed*, he cannot explain why the agent would be justified *in forming* it, given that they do not form it in response to prior evidence for its truth:

Like many others, I find the forming of beliefs without prior evidence epistemically suspect, even when the beliefs in question are believed, and even known, to provide sufficient evidence for themselves, once formed. To take an example from another context, consider "the situation you would be in if you believed yourself (perhaps with good reason) to be watched over by a benevolent spirit, who sees to it that whenever you form a belief on a certain subject-matter (say, the winners of horse-races), it is true." If I come to believe that Malabar will win, and I know that I have done so, my belief will be supported by sufficient evidence, and it may well count as knowledge. I know my own belief and I know that the benevolent spirit will make it true. But it still seems to me

⁴¹ Setiya, "Practical Knowledge Revisited," 136; my emphasis.

unreasonable to *form* that belief. And this is just what I am doing when I make a decision [or: form an intention] on Velleman's account.⁴²

Setiya thinks that the fix is to think of the agent's justification as pre-dating their forming their intention-belief.⁴³ And on Setiya's view, this is indeed the case, for the agent's justification is embodied in a *standing* condition of theirs; their disposition to do what they intend.⁴⁴

But I am not sure about Setiya's description of the problem in the Malabar case, and if I am right, Setiya's 'fix' is no such thing. Setiya thinks that the problem is with the idea that an agent could be *reasonable* in forming a belief that Malabar will win without prior evidence. But what would count as reasonable here? If we assume – as the example does – that agents can form beliefs *at will*, then it becomes appropriate to assess their belief-formation by the dictates of *practical* rationality. If beliefs *can* be formed at will, then performing the action of *forming the belief that Malabar will win* on the basis of the desire for Malabar to win and the belief that if one *believes that Malabar will win*, then he will, is no less reasonable than performing the action of *feeding Malabar magic oats* on the basis of a desire that Malabar will win and a belief that if one *feeds Malabar magic oats* then he will. Both 'actions' – forming the belief that Malabar will win, and feeding Malabar magic oats – seem perfectly reasonable ways of satisfying one's desire for Malabar to win. By analogy, *if* Velleman's agent can form intention-beliefs at will, she seems perfectly reasonable in doing so.

But there *is* something wrong with the Malabar case, I think, and analogously something wrong with Velleman's account, because it is *not* obvious that we *can* form beliefs at will. The problem is with the *psychology* of the case quite independently of what we think about its *epistemology*.⁴⁵ Obviously Velleman won't accept this – on his view it is a feature of intentions that they *are* formed at will and not in response to prior evidence (see again §1(a)&(d)). *Qua* beliefs they constitutively aim at truth, but unlike ordinary beliefs they meet their aim *not* by being formed in response to antecedent evidence, but by causing the actions they represent; by being *self-fulfilling*. But it is far from clear that Velleman is doing anything other than *claiming* that intentions embody a kind of belief that can be formed at will.

The problem *for Setiya* is this: first, he misdescribes the Malabar case as showing that there is an *epistemological* problem with forming beliefs at will, and infers that Velleman's account of practical knowledge is subject to an analogous problem. But quite apart from whether there is an epistemological problem with forming the Malabar belief and analogously in Velleman's picture with forming an intention-belief, there is a *psychological* block to forming beliefs without sufficient prior evidence. This is the real lesson of the Malabar case, I suggest. And if this is right, Setiya's account is no better off than Velleman's. They both hold that intentions are a kind of belief which can be formed at

⁴² Setiya, "Practical Knowledge," 400–401. The quotation is from Cian Dorr Dorr, "Non-Cognitivism and Wishful Thinking," 99–100; the square brackets are mine.

⁴³ Setiya, "Practical Knowledge," 405.

⁴⁴ Jane Heal has pointed out to me that it's not clear that Setiya's account *is* an improvement on Velleman's on this score: isn't Velleman's agent's *belief* that she is disposed to do what she intends, a standing state just like her disposition itself? This seems right to me but I don't have space to consider the point in detail.

⁴⁵ Rae Langton agrees with me here, and goes even further, suggesting that being self-fulfilling actually *undermines* a belief's capacity to constitute knowledge (Langton cites Humberstone, "Direction of Fit," 62). I don't consider these thoughts here, but see Langton, "Intention as Faith."

will, and not in response to prior evidence for their truth. But Setiya does not, any more than does Velleman, give us an explanation of why forming an intention-belief at will is any more psychologically plausible than forming the analogous Malabar belief at will.

Velleman and Setiya both *bold* that intentions are a special kind of belief which can be formed at will, but for many the observation that intentions *can* be formed in the absence of sufficient prior evidence for their truth will be a reason to *deny* that intentions are a kind of belief. Setiya brings up a problem for Velleman which – if I am right about how it is properly understood - he cannot answer.

c. Conclusion

Setiya's externalism enabled him to avoid the two major problems which led me to reject Paul's and Velleman's accounts. But Setiya's account is not without problems of its own. First, it is not obvious whether Setiya thinks that non-accidentally true belief is sufficient for knowledge. If it is not then his account of what is required for practical knowledge is incomplete and he flouts the *Why Knowledge?* constraint. And second, like Velleman's, his account depends on thinking of intentions as a kind of belief which can be formed at will. But the idea of a kind of belief which can be formed at will seems to conflict with what appears to be a commonsense-psychological truism about belief: you can't *decide to believe*. Within the methodological context of this Thesis, this is particularly problematic.

The idea that intentions are a kind of belief has, then, caused problems for both Velleman and Setiya. And in the next chapter I will give a more sustained argument against Cognitivism which will lead us to rule out *any* account of practical knowledge which depends on it.

Chapter Four:

The Content of Intention

*Introduction*¹

It is often assumed,² although rarely argued,³ that intentions are propositional attitudes. I will call the claim that intentions are in all cases propositional attitudes ‘Propositionalism about Intention’, or *Propositionalism* for short,⁴ and I will argue here that it is false. This matters for now because it undermines *any* version of *CI* - including but not restricted to those of Velleman and Setiya. And it will matter later on (especially in Chapter Six) because it will constrain any version of *Intentionalism*.

In §1 I introduce the key ideas needed to make my argument; those of a *propositional attitude* and of what is involved in an intention’s being *executed*. I give the argument in §2: accepting for the sake of argument that *certain* intentions *are* propositional attitudes, I go on to show that certain others *cannot* be; no *propositional* intention could be executed unless some *non-propositional* intention is executed. We do execute our intentions, so some intentions are not propositional attitudes. I consider in §3 how the Propositionalist might respond, but no response is found compelling. So in §4 I conclude that Propositionalism ought to be rejected, and a further constraint on an account of practical knowledge - the *Intention-Content* constraint - is introduced. This constraint requires that an account of practical knowledge does not entail Propositionalism, and *CI* is rejected as flouting this constraint.

Two notes before we start. First, I restrict my discussion in this Chapter to *singular* propositional intentions; intentions whose contents are propositions which make singular reference to an individual. I do this in order to simplify my discussion, but it is justified by a feature of the argument I will give in §2. For the intentions which I am going to argue *are not* propositional are intentions which the Propositionalist will hold have *singular*

¹ The ideas in this section have benefited greatly from discussions with, in particular, Chris Cowie, Ryan Cox, Tim Crane, Alison Fernandes, Alexander Greenberg, Jane Heal, Richard Holton, John Maier, Peter Sutton, and Jen Wright; from members of the ‘Serious Metaphysics Group’ in Cambridge, and of the ‘Seriously Summery Summer Surrogate Seminar’ at King’s College London.

² For representative examples of assumptions, see Harman, “Practical Reasoning,” 441; Fodor, “Propositional Attitudes,” 505–6; Searle, “The Intentionality of Intention and Action,” 259; Vermazen, “Objects of Intention”; Williamson, *Knowledge and Its Limits*, 21; Shah and Silverstein, “Reasoning in Stages,” 101, and various works of John Broome’s, e.g. “Normative Practical Reasoning,” 176.

³ I know of only two arguments for the view. Wane Davis argues that we “achieve greatest generality” if we translate all intention-attributions into propositional form (Davis, “A Causal Theory of Intending,” 43). This argument would work only if doing so would not misrepresent the facts; this Chapter argues that it does. John McDowell assumes that intentions are able to constitute practical knowledge and infers from the propositionality of practical knowledge to the propositionality of intention (McDowell, “What Is the Content of an Intention in Action?” 424). I will not discuss his argument here, but will make it clear in Chapter Six why I think the principle it rests on – that the content of an attitude must match the content of any piece of knowledge it constitutes – is false.

⁴ Thus my ‘Propositionalism’ is distinct from the ‘Propositionalism’ considered in e.g. Montague, “Against Propositionalism” and in Crane, *The Objects of Thought*, 108; the view that all *intentional* attitudes are *propositional* ones.

propositional contents. From now on, then, ‘proposition’, ‘propositional’ etc. will be qualified by ‘singular’, although this qualification will be left implicit.

Second, my conclusion is, I think the same as Annette Baier’s conclusion in “Act and Intent”,⁵ and some of the considerations used in my argument are also used by Baier. But since our arguments themselves are distinct, and since Baier is responding to a particular account of intentions which is *not* strictly a Propositionalist account,⁶ I will not discuss Baier’s argument in this Chapter, although I will indicate some points of convergence between Baier’s discussion and mine in footnotes.

1. Preliminaries

The idea in §2 will be that attention to what is required for a *propositional* intention to be *executed* will show that not all intentions are propositional. So we need to know what it means for an attitude to be *propositional*, and what it means for an intention to be *executed*. This is our task in this section.

a. Propositional Attitudes

I do not aim to thoroughly investigate the notion of a propositional attitude, just to say enough for our purposes. Propositional attitudes are attitudes with propositional contents. What are *propositional contents* like? I will note three uncontroversial features of propositional contents: they are *truth-apt*, they can be picked out (in English) using indicative *‘that’-clauses*, and they are *about* something with propositional shape. We can’t use the first two features to find out whether a given intention is propositional or not, for reasons I will explain. But I think we can use the third.

First, propositional contents are truth-apt. So if we can decide whether an intention has a truth-apt content, we can decide whether it has a propositional content.

We observed in our discussion of Velleman that it sounds odd to call intentions true or false (Ch. 3; §1(d)), but we also saw there that this need not entail that they *are not* true or false. Although the strategy was not available to Velleman (for reasons there explained), a Propositionalist of a different stripe might explain the oddness of *calling* intentions ‘true’ or ‘false’ by reference to their *direction of fit*. To recap the thought, imagine that Sharon has an intention whose propositional content is *that she will deceive Phil*. Now assume the orthodox view of intentions as (*contra* Velleman) representing their contents *as to be made true* rather than *as being true*. The propositional content of Sharon’s intention is *that she will deceive Phil* and this content is true, but it is represented by Sharon’s intention *as to be made true* and not *as true*. Contrast Linda’s *belief that Sharon will deceive Phil*. Linda’s belief has the same – true – propositional content as Sharon’s intention (*ex hypothesi*).⁷ But Linda’s

⁵ For related discussions see also Müller, “Radical Subjectivity: Morality Versus Utilitarianism”; Clark, “The Action as Conclusion,” 501–502; Thompson, “Naïve Action Theory”; Alvarez, *Kinds of Reasons*, 65–73; Dancy, “Action, Content and Inference,” 285–286; Boyle and Lavin, “Goodness and Desire.”

⁶ That found in Chisholm, “The Structure of Intention.”

⁷ I am simply assuming here that *beliefs* are propositional attitudes, an assumption which has been rejected by David Lewis, who thinks of believing not as a relation to a proposition but as the self-ascription of a property (Lewis, “Attitudes de Dicto and de Se”). I think that the discussion of this chapter could be restated from the Lewisian standpoint, but cannot argue for this here. I will come back to Lewis’ idea in concluding the Thesis.

belief represents the propositional content *as being* true. Both women's attitudes *are* true, but Linda's attitude is committal about its truth-value whereas Sharon's is non-committal.

So it is consistent with it *seeming* odd to *call* intentions true (or false), that they do in fact have true (false) contents. For this reason, the fact they don't *seem true* (false) doesn't tell us whether or not they *are in fact* propositional attitudes.

Second, propositions can be expressed (in English) using an indicative 'that'-clause, and propositional attitudes can be ascribed using a psychological verb which embeds such a clause: Bev hopes *that her sister will call later*; Lauren fears *that she's being followed*. So if we can decide whether an intention can be ascribed in this way, we can decide whether it has a propositional content.

Is often observed that natural-language intention-attributions rarely contain an indicative 'that'-clause, but instead some kind of infinitival 'to'- construction. "Sharon intends *to deceive Phil*"; "Judge Julie intends Bev *to go to prison*". But the Propositionalist will (and does⁸) resist the inference from this fact about the grammar of common natural language intention-attributions, which they accept, to the conclusion that the attitudes thereby ascribed have non-propositional contents, which, of course, they don't. They are likely to point out that it is not in general a failsafe procedure to read off the semantic structure of what is expressed from the surface linguistic structure of what expresses it. For some reason, the thought will continue, intention-attributions are unlike belief-attributions (and attributions of other propositional attitudes) in that their content-clauses are not typically transparent to the attitudes' contents.

The observation that natural-language intention-attributions *in fact tend not to* contain 'that'-clauses does not render false the starting assumption that propositional attitudes can be attributed with a psychological verb embedding a 'that'-clause. Nor does it mean that intentions have to be understood as an exception to that claim. For 'to'-clauses can be transformed into 'that'-clauses quite easily.⁹ Intentions *can* be attributed with a psychological verb embedding a 'that'-clause and for the Propositionalist, this is the *proper* way to attribute them, since it is more revealing of their structure than the natural-language 'to'-attributions.

Where the 'to'-clause contains a subject-expression and a predicate-expression, these can be put together in constructing a 'that' clause: "Judge Julie intends Bev *to go to prison*" becomes "Judge Julie intends *that* Bev *will* go to prison". (We cannot simply carry over "go to prison" because "Bev go to prison" is not an indicative sentence. By inserting "will" we secure something clearly truth-apt.) We can do the same for any ascription of the form "*S* intends *a to φ*", turning it into "*S* intends *that a will φ*". Transformation is only slightly more complicated for intention-attributions with subjectless content-clauses, such as Sharon's intention *to deceive Phil*. Here we must come up with a subject-expression in constructing our 'that'-clause. But a suitable candidate is close at hand, since intentions whose contents are attributed with "*to φ*" *only ever* concern the actions or properties of *the intender*. So we can non-arbitrarily plug in an expression referring to the intender, and this gives us a propositional subject. More concretely, "Sharon intends *to deceive Phil*" becomes "Sharon intends *that she [Sharon] will deceive Phil*".¹⁰

⁸ The tactic is commonplace (see again n. 2), but for a more thorough discussion than is usual, see Davis, "A Causal Theory of Intending," 43.

⁹ Ibid.; Wallace, "Propositional Attitudes and Identity," 149.

¹⁰ This suggestion maps on to the standard syntax for infinitives, which construes infinitival clauses as

Of course the fact that there are rules for transforming ‘to’-clauses to ‘that’-clauses does not show that intentions attributed using “to φ ” in the content clause were all along really propositional intentions. This conclusion would follow only if transforming the “to”-form into the “that”-form introduced no semantic difference, and the mere possibility of transformation does not speak to this, just as the possibility of transforming any statement about *you* into a statement about *me* (rule: replace any term referring to me with one referring to you) does not prove that we are identical. The point for now is that we *can’t* infer from the lack of ‘that’-clauses in typical natural-language intention-attributions to the conclusion that the intentions thereby attributed are not propositional. The fact that intentions *are not* typically attributed using ‘that’-clauses does not entail that they *cannot* be, nor that their contents might be more perspicuous if attributed using a ‘that’-clause.

Finally, intentional states are *about* intentional objects. Tim Crane describes the idea, in a different context, nicely:

[O]ne cannot adequately understand what [a given intentional] state is without knowing what its [intentional] object or objects are. [...] Intentional objects are just whatever one’s intentional states are directed on. My thought [that my doctor smokes] is directed on my doctor, on smoking, and on the fact that he smokes: intentional objects can be ordinary objects, properties, events or states of affairs (so ‘object’ here does not mean *particular*).¹¹

A *propositional* intentional state will be about a propositionally shaped intentional object. So if we can decide whether an intention is *about* a propositionally shaped intentional object, then we can decide whether the intention is a propositional attitude.

For an attitude to be *about a propositionally shaped object* it must minimally be *about* a thing, a property, and the-instantiation-of-the-property-by-the-thing.¹² For short, we can say that a propositional attitude is about some *x*’s Φ -ing, where ‘ Φ -ing’ is instantiating some property or behaving in some way (including but not limited to *acting* in some way). Tim’s thought is about some *x*’s Φ -ing; it is about *his doctor’s smoking*. This entails that it is also about *his doctor* and about *smoking*. If Sharon’s intention to deceive Phil is propositional, it also must be about some *x*’s Φ -ing. It is clear from our attribution that her intention is about *deceiving Phil*. But if the Propositionalist is right, then it is also about *Sharon’s deceiving Phil*, and so about *Sharon* too.

The problem with the first two features of propositional contents was that it didn’t look possible to see whether an intention’s content had these features without already having decided whether it was a propositional attitude. So although propositional intentions *will*

containing hidden anaphoric reference, labeling this with ‘PRO’. In the standard syntax, “Sharon intends to deceive Phil” is rendered as “Sharon intends [*PRO to deceive Phil*]” (Radford, *Transformational Grammar: A First Course*, 313–317). I won’t be taking issue with the standard syntax: my argument is that the contents of intentions such as Sharon’s don’t make *semantic* reference to their subjects.

¹¹ Crane, “The Intentional Structure of Consciousness,” 37–38; see also Crane, *The Objects of Thought*, Ch. 4.

¹² I say ‘minimally’ because it might *also* be about *two* distinct objects and the relation between them. My belief that Henry likes Izzy can be thought of as being about *Henry*, about *liking Izzy* and about *Henry’s liking Izzy*. But it can also be thought of as about *Henry*, *Izzy*, about *being liked by* and about *Henry’s liking Izzy*. I will stick to considering the simplest available decompositions of the propositional contents we will be considering, which will always distinguish a single object and a single predicate within the proposition.

have these features, we can't look for these features in deciding whether the intention *is* propositional. But we can, I think, use the final feature in this way, because there is an important relationship between what an intention is *about* and what is required for the intention to be *executed*, as I shall now explain.

b. Executing an Intention

Central to our commonsense understanding of agency is a distinction between an intention's being *executed* and its content's merely being *satisfied*; a distinction recognised (although not necessarily in these terms) by anyone with the concepts of *intention* and *intentional action*. Let's consider the distinction.

I intend to water the tomatoes. Remaining neutral on whether my intention's content is propositional (as I will throughout this section), we can distinguish two kinds of case in which I go on to water them. In one, I *carry out* my intention, watering them *intentionally*; in the other, I water them without carrying out my intention and so not intentionally - even though watering them was what I intended to do. In this second case my watering is related too accidentally to my intention to count as *carrying it out*. I slosh water all over the grow-bags after tripping over the step, for instance.¹³

Intentions, I am assuming, can also concern the behaviour and properties of objects other than oneself: Judge Julie's intention for Bev to go to prison is like this. The distinction also attaches to intentions like this, for there are two kinds of case in which what Julie intends comes to pass. In the first, Bev goes to prison in virtue of Julie *carrying out* her intention (there will be a story to be told about *how* she carries it out - by *sentencing Bev*, let's imagine). In the second, Bev goes to prison alright, but independently of anything Julie does. Here Julie's intention is not carried out, even though what Julie intended comes about.

The distinction in the second pair of cases is slightly different to that in the first. One of the differences between the cases in the first pair was whether what was intended was *done intentionally*, whereas what was intended was not done intentionally in *either* case in the second pair. This was because whilst *watering the tomatoes* was something I could do, *Bev's going to prison* was not something *Julie* could do; poor old Bev would have to do that.

But there is something common to the distinctions in both pairs. Intuitively, in the first case of each pair, *what is intended* comes to pass *because* the intention comes to fruition in the way that it should, whereas in the second case of each pair, *what is intended* comes to pass *despite* the fact that the intention *fails to* come to fruition in the way it should. It is this common distinction that I want to capture by talking about the *execution* and the mere *satisfaction* of an intention. We can be more precise about the distinction as follows:

- i. An intention's content is *satisfied* when what is intended is (somehow) brought about
- ii. An intention is - further - *executed* when what is intended is brought about,
 - a. *by the intender herself* and
 - b. *intentionally*

¹³ This distinction is familiar from the literature on causal deviance and *CTA*, but the distinction itself is a commonsense datum, not something particular to *CTA*.

These conditions are supposed to be plausible independently of whether or not Propositionalism is true. *What is intended* can be *brought about* if it is a proposition or state of affairs, but also if it is *an action*, and so *not* propositionally shaped. Note also that (i) and (ii) are not intended as *definitions* of ‘satisfaction’ and ‘execution’: they are merely necessary conditions. Further, they are not supposed to comprise a *theory* about intention-execution and satisfaction or a reductive explanation of the notions. Positively, they articulate the *commonsense-action-theoretic* distinction illustrated above, with some precision but on the same explanatory level. These points will become important in §3.

Let’s see how (i) and (ii) classify the second case in each of our pairs. The case in which Julie intends Bev to go to prison, and Bev does, but with no input from Julie, meets the *satisfaction* condition (i) since *what Julie intended* (for Bev to go to prison) is brought about. But it flouts the execution condition (ii) because this is not brought about *by Julie* (it flouts (ii.a)). The case in which I intend to water the tomatoes, and I do water them, but only because I trip also meets (i), because *what I intended* (to water the tomatoes) is brought about. And it meets (ii.a) because this is brought about *by me*. But it flouts the execution condition because I do not bring it about intentionally, but by accident (it flouts (ii.b)).

I have stuck to an intuitive conception of whether *what* is intended is brought about, because I have wanted to avoid making assumptions about what *is* intended, and so about whether the intentions in question are propositional or not. In the next section we will consider what is involved in bringing about what is intended when an intention is propositional.

2. An Argument against Propositionalism

I start by conceding, for the sake of argument, that *some* intentions *are* propositional attitudes. Since it is plausible that an intention is about *at least* those objects and properties mentioned in the content-clause of its natural-language ascription, I will start by assuming that when an ascription’s content-clause contains an explicit subject and predicate, it attributes a propositional intention. Julie’s intention for Bev to go to prison, for instance, will be assumed to have the content *that Bev will go to prison*. All ascriptions of intentions for something *other* than the intender to do something (or become or remain some way) will contain explicit subject- and predicate-expressions, so I will assume that whenever an intention is naturally so-ascribed it is a propositional intention. And I will refer to any intention whose content is propositional as an ‘I(Φx)’, where ‘I’ stands for ‘intention’ and ‘Φx’ makes the structure of the attitude’s content explicit.

Not all natural-language intention-attributions’ content-clauses contain explicit reference to a subject, as we saw in §1(a); some contain only a predicative expression. My intention *to water the plants* is like this, as is Sharon’s intention *to deceive Phil*. The Propositionalist holds that even these intentions are – *really* – I(Φx)’s. I am going to argue that at least some of them are not, although I will *not* be arguing from facts about the linguistic properties of these attributions – I explained why such a strategy wouldn’t get us far in §1(a).

a. The Problem

Let's start by considering what seems to be involved in executing an *agreed* $I(\Phi x)$. Bev has been found guilty of murder and Judge Julie intends her to go to prison. We will assume that Julie's intention has the propositional content *that Bev will go to prison* – although I will sometimes use more colloquial language to indicate it.

If Julie's intention is to be *executed*, she needs to meet conditions (ii.a) and (ii.b) from §1(b). Meeting (ii.a) requires her to *bring about* what she intends. She can't do this by magic; she needs to employ her causal powers. In particular, she needs to employ them in bringing about *Bev's going to prison*. She needs to *cause Bev to go to prison*. Let's imagine she does this in the usual way – by sentencing Bev to a prison term (and doing whatever this involves – uttering certain words in court, signing various forms *etc.*). But she also needs to meet (ii.b), which requires her to *intentionally* bring it about that Bev goes to prison. She must *intentionally* sentence Bev. And this requires her to execute a further intention, *to sentence Bev*.¹⁴

For the Propositionalist, this new intention of Julie's is another $I(\Phi x)$. This means that it is not only about *sentencing Bev*, but is also about *Julie*; its content is *that Julie will sentence Bev* (from Julie's perspective, *that I will sentence Bev*). In order to execute this intention Julie must (ii.a) bring it about *that she sentences Bev*, and (ii.b) do this intentionally. How does she do this? I will argue that the Propositionalist can give no acceptable answer to this question.

b. Somehow Self-Causation: Regress

The first suggestion is that all $I(\Phi x)$'s are executed in the same way, *viz.* by their subjects intentionally *causing the x to Φ* . Julie's intention for Bev to go to prison – whose content was the proposition *that Bev will go to prison* – was executed by Julie's *causing Bev* to go to prison. Just so, on the current suggestion, Julie's intention to sentence Bev – whose content is the proposition *that she will sentence Bev* – will be executed by Julie's *causing herself* to sentence Bev. But this suggestion leads first to odd results, and second to a vicious regress.

Prima facie, we can ask *how* Julie will cause herself to sentence Bev. This 'How?' question intelligibly comes up whenever something or someone causes something to happen.¹⁵ "The brick broke the window." - - - "How?" - - - "By smashing into it"; "Julie will cause Bev to go to prison." - - - "How?" - - - "By sentencing her." If Julie causes herself to sentence Bev, it seems we can ask how she does *this* too.

Here is *one* way in which she might do so: knowing herself prone to nerves when passing custodial sentences, Julie downs a few shots of tequila before walking into court. She knows tequila makes her vindictive and prone to power-trips, and the tequila technique has overcome her reluctance to send defendants down in the past. If it works this time, then by employing it she will have caused herself to sentence Bev.

¹⁴ I assume here that doing something intentionally is executing an intention to do it. Michael Bratman has argued against this 'Simple View' of the relation between intentional action and intending (in *Intention, Plans, and Practical Reason*, Ch. 8). It would complicate my argument to consider Bratman's objections to the Simple View here, but if he is right I think a version of my argument will survive. An interesting reason why he may not be right is given by Falvey in "Knowledge in Intention," 41.

¹⁵ See also Baier, "Act and Intent," 653.

But although this describes *a* possibility, it is not the one I had in mind for this example, which was imagined as one in which Julie was *simply able to* sentence Bev. Let me be clear about what ‘simply’ means here. The idea is not that Julie is able to sentence Bev *as a basic action* in the sense described in Chapter Two (§1(c)): Julie cannot sentence Bev except by uttering certain words, signing certain forms etc., so under the description ‘sentencing Bev’ her action will be *non-basic*. Still, Julie seems ‘simply’ able to sentence Bev in the sense that she seems able to do so without doing anything – like employing the tequila technique – to *cause herself* to do so.

If this appearance is taken at face value then (assuming Propositionalism) we have a counterexample to the claim that executing *any* $I(\Phi x)$ requires the intender to cause the x to Φ . But perhaps there is some reason why we ought not take this appearance at face-value. Perhaps there is a more plausible way of thinking about how Julie might cause herself to sentence Bev, than that offered in the previous paragraph; a way which is consistent with the *appearance* that Julie can sentence Bev without causing herself to do so. Without wanting to trawl through suggestions for this ‘more plausible way’, let’s simply *posit* some action α , which fits the bill; which is *ex hypothesi* an action by doing which Julie causes herself to sentence Bev, and which for whatever reason does not conflict with the *appearance* that Julie need not act on herself in so doing.

Even helping ourselves to α , there is trouble. For if doing α is to count as the *execution* of her intention to sentence Bev, then it must be something Julie does *intentionally* (it must meet (ii.b)), and so must be done in execution of an *intention to α* . Like all intentions, this is, for the Propositionalist, another $I(\Phi x)$; its content is *that Julie will α* . Our current assumption is that an $I(\Phi x)$ is executed only if the intender causes x to Φ . So executing this new intention requires Julie to cause herself to α . How? Again it seems Julie will need to *do something* to cause herself to α – call it β . If doing β is to count as the execution of her intention to α , she must β *intentionally* (ii.b), and so must execute an intention to β . But again her intention to β is – for the Propositionalist – another $I(\Phi x)$, whose content is *that Julie will β* , which the current proposal tells us will be executed only if she *causes herself* to β . How does she do this? By doing something to cause herself to β – call it γ ... And so on *ad infinitum*. The regress is vicious because the conditions on executing any $I(\Phi x)$ generate further distinct $I(\Phi x)$ ’s, which must be executed in order for the initial $I(\Phi x)$ to be executed. Executing *any* intention becomes impossible.¹⁶

But we *can* execute our intentions, so either they are not all $I(\Phi x)$ ’s – and Propositionalism is false – or not all $I(\Phi x)$ ’s intentions are executed by somehow causing the x to Φ . Let’s consider another way in which an $I(\Phi x)$ might be executed.

c. No-How Self-Causation: Incoherence

I said above that *whenever* a person causes something to happen, there is something to be said about *how* – *by doing what* – the person did so. This was one of the assumptions which generated our regress. But perhaps this assumption is false; perhaps at least sometimes we can cause things to happen without *doing anything* which causes them to happen; perhaps there are *some* cases in which people *simply* (in the sense of ‘simply’ described above) cause things to happen; in which they cause things to happen *no-how*. If Julie can cause herself to sentence Bev no-how, this would stop the regress from (b); it would mean that although executing her intention for Bev to go to prison requires executing the

¹⁶ Two distinct but closely related regresses are described in Ryle, *The Concept of Mind*, 65–6 and in Davidson, “Problems in the Explanation of Action,” 102.

further intention *to sentence Bev*, this further intention can be executed without itself generating any more intentions.

But apart from its capacity to block our regress, what reason do we have to think that things ever cause other things to happen *without doing anything* to so cause them? The idea is very hard to understand, perhaps incoherent.¹⁷ One way to bring out the problem is to consider the question of *when* the effect happens and to ask for an explanation of this fact.¹⁸ Say the brick broke the window at *t* - why did the window break *then*? Simply mentioning *the brick* won't explain why the window broke *at t* rather than some time before or after *t*, at which the brick also existed. We need to mention something the brick *did at t* – *viz.* smashing into the window - to explain why it broke the window *at t*.

It might be accepted that the idea of an object's causing something to happen without *doing something* to cause it to happen is *in general* incoherent, but suggested that when that object is an agent, and the something that happens an action of theirs, there is an exception to the general rule that objects can cause things to happen only by doing things. After all, agents are in an especially metaphysically intimate relationship to their own actions, so perhaps it is true that the brick can cause the window to break only *somehow*, but still maybe Julie can cause herself to sentence Bev *no-how*.

But any appearance of a solution in this suggestion is chimerical. For it is unclear that we can say anything about this 'especially metaphysically intimate relationship', except that it constitutes or underwrites an agent's capacity to cause their own actions *no-how*, where she can cause other kinds of happenings only *somehow*. What such a suggestion would need to provide is an explanation of *how and why* no-how self-causation is possible, but instead it only claims that it *is* possible for agents, albeit now in the language of 'metaphysical intimacy'. The re-phrasing makes the suggestion itself no more coherent or less mysterious.

To be clear, my point is not that agents cannot *cause themselves to act*.¹⁹ I accepted that Julie *could* cause herself to sentence Bev by downing tequila shots, although she *does not* in the example we are working through. There are everyday examples too. Because I am not a morning person, I usually have to *get myself* to get up early: I do so by setting an alarm the night before. Because I am shy I sometimes have to *make myself* approach people in social situations: I do so by giving myself a pep-talk. These are intelligible, colloquial, examples in which agents cause themselves to act; in which they are not able *simply* (in the sense of 'simply' explained above) to act. But they are intelligible precisely because there is something to be said in such cases about *how* – *by doing what* – the agents cause (more colloquially: *get*) themselves to act.²⁰ In contrast, it is hard to understand the difference between my causing myself *no-how* to (say) get up early, and my being able to do so *without causing myself* to at all.

It is not an attractive option for the Propositionalist to suggest that Julie can cause herself to sentence Bev *no-how*. But there is *prima facie* another way in which the regress might be stopped. To this I now turn.

¹⁷ Adrian Haddock worries about this kind of suggestion too; see Haddock, "At One with Our Actions, but at Two with Our Bodies," 159.

¹⁸ Thanks to Jane Heal for this way of putting things.

¹⁹ See also Baier, "Act and Intent," 658.

²⁰ See also Ibid., 652–655.

d. Execution without Self-Causation: Collapse

The Propositionalist thinks that Julie's intention to sentence Bev has the propositional content *that Julie will sentence Bev*. The problem is to explain what is involved in executing this intention. In (b) and (c) we considered two ways of understanding the claim that Julie needs to cause herself to sentence Bev in order to execute her intention. But both were problematic. So the Propositionalist may be tempted to deny that in order to execute her intention *that she will sentence Bev*, Julie needs to *cause herself* to do so.

On this suggestion, executing certain $I(\Phi x)$'s requires the intender to cause x to Φ , but executing others does not; Julie's intention *that Bev will go to prison* is of the first kind (Julie *can't* execute this *except by* causing Bev to go to prison), but her intention *that she will sentence Bev* is of the second kind (she can execute *this* without causing herself to sentence Bev). Like the suggestion in (c), this suggestion promises to stop the regress in (b).

There are three closely related explanatory challenges for this suggestion. First, *which* $I(\Phi x)$'s are the privileged ones; those which can be executed without the agent causing x to Φ ?²¹ Second, *what is it about* these privileged $I(\Phi x)$'s which explains why they can be executed in this special way, given that most $I(\Phi x)$'s can be executed only by causing x to Φ ? And third, *how are* these intentions executed?

An initial response to the first challenge suggests that the privileged set of $I(\Phi x)$'s contains just those intentions which concern the intender's *own* actions. I will refer to these as $I(\Phi i)$'s ('i' for *me*). But the suggestion does not work. As we saw above, executing *some* $I(\Phi i)$'s, *does* involve causing oneself to Φ . For the Propositionalist, my intention to get up early is an $I(\Phi i)$; its content is *that I will get up early*, but I may not be able to execute it other than by causing myself to get up early (by setting my alarm). A second suggestion is the privileged set of intentions is the set of *basic* $I(\Phi i)$'s; intentions to do something as a basic action. But this won't work either, for many *non-basic* actions can be done 'simply' in the current sense, *i.e.* without causing oneself to do them. Julie – we have now accepted – doesn't need to *cause herself* to sentence Bev, but *sentencing Bev* is not a *basic* action of hers; she does it by uttering certain words and signing various forms etc.

So the Propositionalist owes us an account of how she intends to delimit the class of $I(\Phi x)$'s the execution of which *does not* involve self-causation. We can assume that it will contain *only* $I(\Phi i)$'s, but it does not look easy to specify *which* $I(\Phi i)$'s it will contain. The second explanatory challenge can be met only by meeting the first, since saying what it is about $I(\Phi i)$'s in the relevant class that accounts for the fact that executing them does not require self-causation would require saying *which* $I(\Phi i)$'s they are.

Rejecting Propositionalism, by the way, affords a very easy answer to the first two explanatory challenges. The non-propositionalist identifies the class of intentions which can be executed without self-causation with the class of non-propositional intentions; those whose contents lack a propositional subject. And it is exactly this feature of these intentions which explains why they can be executed without self-causation: because they are not *about* oneself, we should not expect bringing about *what they are about* to involve exerting one's causal influence *over oneself*. Julie's intention for Bev to go to prison is about Bev's going to prison, and so Julie needs to *cause* Bev's going to prison. If Julie's intention

²¹ Someone taking the route considered in (c) would have to meet this challenge too; I consider it here because I think that the Propositionalist is more likely to pursue the current option.

to sentence Bev is *not* about *Julie's sentencing Bev*, but is just about *sentencing Bev*, then Julie need not cause herself to sentence Bev. She need cause only *Bev's sentencing*.

Let's assume for the sake of argument that the Propositionalist can meet the first two explanatory challenges (although I think it very unlikely). Let's consider the third. This means assuming for the sake of argument that Julie's intention to sentence Bev is in the class of intentions which can be executed *without* self-causation. What *is* involved in Julie's executing her intention to sentence Bev?

Note first that we cannot altogether dispense with the idea that executing an intention involves exerting one's causal influence in bringing about what the intention is about. That idea was essential to the distinction between *execution* and *satisfaction* in §1(b), and this distinction is itself essential to the commonsense concept of intentional action.

So the idea that certain $I(\Phi)$'s can be executed without self-causation has to be the idea that in executing it, one exerts one's causal influence over something *other than oneself*, in such a way as to *result in* or *constitute one's Φ -ing*, but *without* this itself amounting to *causing oneself* to Φ . For we are trying to explain how an $I(\Phi)$ might be executed *without* the intender causing herself to Φ (the self-causation idea led to regress (b) or incoherence (c)).

This doesn't leave much theoretical wiggle-room for describing what is involved in Julie's executing her intention to sentence Bev. Here is the only thing – I think – left to say: Julie must *exert her causal influence over Bev in such a way as to cause Bev to get sentenced*. We have already seen how she might do this – by uttering certain words and signing some relevant forms. In thinking of her as exerting her causal influence *over Bev* we need not think of her as physically interacting with Bev in the first instance; the idea is just that Bev is affected by what Julie does. Importantly, *Julie's* role is wholly on the *cause* side of the cause-effect relation instantiated when she executes her intention, whereas on the *effect* side we find poor old Bev, and facts relating to her new status.

As a description of how, in executing her intention to sentence Bev, Julie's causal powers are directed, the above seems plausibly true and certainly a way of stopping the regress in (b). But the Propositionalist cannot accept the above account, because accepting it undermines Propositionalism. For the suggestion involves thinking of Julie's intention as *executed* by Julie's directing her causal powers over Bev, in a sentencing-kind-of-way. And if *executing* it involves directing *her causal powers* over only things other than herself, *having* it seems plausibly to involve a directing of *her mind* over only these same non-Julie objects. So what her intention to sentence Bev is *concerned with*, or *about*, seems to be simply *sentencing Bev*, and not also about *herself*, and so not about *her* sentencing Bev. But this would mean that its intentional object is not *Julie's sentencing Bev* but simply *sentencing Bev*. Because its intentional object is not propositionally structured, her intention is not a propositional attitude.

e. Summing up

Let's re-cap the argument of this section. I started by assuming for the sake of argument that some intentions *do* have propositional contents, and took as an example Julie's intention for Bev to go to prison. We said that its content was *that Bev will go to prison*.

We then asked what is involved in executing an $I(\Phi x)$. In (b) we imagined that *all* $I(\Phi x)$'s are executed by *somehow causing the x to Φ* , but this led to a vicious regress. Two ways of stopping the regress were considered in (c) and (d). In (c) the suggestion was that some $I(\Phi x)$'s can be executed by causing the x to Φ *no-how*, but this seemed at best mysterious and at worst incoherent. So in (d) we considered the suggestion that some $I(\Phi i)$'s can be executed *without any* self-causation. But despite being plausible as a description of what Julie might do in executing her intention, the suggestion seemed to undermine Propositionalism.

Before moving on I would like to return briefly to some of the considerations in §1(a). There I said that it seemed odd to call intentions 'true' or 'false', and that in natural language we don't tend to attribute intentions using 'that'-clauses. I said that we could not *argue from* these observations to the falsity of Propositionalism, and I have not done so - my argument proceeded from metaphysical considerations about the requirements on intention-execution, derived from commonsense-action-theory. But if my argument works, it explains and vindicates our natural language attributional practises and semantic intuitions. And it means that the relatively commonplace practise within philosophy of mind of forcing that-clauses into intention-attributions is *not* a colloquially unusual but philosophically proper way of making the content of these attitudes perspicuous; it is a misleading practise which obscures the true content of the attitudes thereby attributed.

Let me turn to considering how the Propositionalist might respond to my argument.

3. Three Responses and Conclusion

I will consider three responses. The first is that my argument presupposes a problematic view about the role of causation in agency. The second claims that there are reasons for accepting Propositionalism which are independent of the considerations in §2. And the third suggests an account on which an agent's being represented in the content of her intention is consistent with her *not* acting on herself in executing it. Taking each in turn I explain why I find these responses unconvincing.

a. Causation and the Execution of Intention

I said in §1 that executing an intention requires the agent to bring about – cause - what the intention is about, and then argued in §2 that there was no satisfactory way in which the intender could do this if Propositionalism were true. The current objection is that along the way I have presupposed an unattractive conception of agency and that if we accept a more attractive one, the problems I have described disappear.²²

The unattractive view is the 'Agent-Causal' view that agents are a special, *sui generis* kind of cause. It is thought unattractive because being *sui generis*, the kind of causation it posits seems *mysterious* from a naturalistic perspective. As Davidson put it, we want the notion of "cause" implicated in agency to be the "ordinary notion of cause which enters into scientific or commonsense accounts of non-psychological affairs".²³ The attractive view

²² Hugh Mellor, for instance, has made this suggestion to me.

²³ Davidson, "Actions, Reasons, and Causes," (Introduction, xv). Examples of the Agent-Causal view can be found in Chisholm, "Freedom and Action" and in Taylor, *Action and Purpose*; some arguments between Agent-Causality and CTA are discussed in Bishop, "Is Agent-Causality a Conceptual Primitive?"

is *CTA*, and is designed *inter alia* to avoid just these kinds of difficulty by assimilating causation by agents to causation in the “ordinary” sense, viewing agents as causing things in virtue of their mental states and events causing things - where *intentional* agency is a special case in which the mental events in question are (manifestations of) her beliefs, desires and/or intentions. The current objection has two parts. First, that I am assuming Agent-Causalism in saying that executing an intention requires *one’s bringing about* what it is about, and second, that if we adopt *CTA* instead, my argument will not work.

But it is not true that I assume Agent-Causalism. To understand why, we need to return to what I said in §1(b) about the *status* of the requirements on intention-execution; of (ii.a) and (ii.b). I characterised these conditions as articulating merely necessary commonsense-action-theoretic conditions on intention-execution. The argument between *CTA* and Agent-Causalism is an argument about whether an agent’s *causing something to happen* can be explained in terms of causal relations of the ordinary kind; between facts or events. *CTA* says it can, Agent-Causalism says it can’t. My argument *does* require accepting (ii.a) and (ii.b), but it does not require accepting that (ii.a) and (ii.b) are the last word about how to understand intention-execution; that they cannot be re-cast in terms of the causation of bodily movements by mental states or events, or by facts about these.

If my argument works (if *any* argument works), it works because the concepts it employs sit in a certain implicational structure. If *CTA*’s proposed reduction works (if *any* reduction works) it works because the reducing concepts sit in an implicational structure which can underwrite the implicational structure of the reduced concepts. It follows that my argument will work just as well at a lower level of description as it does at the commonsense-human level. I don’t assume Agent-Causalism, and my argument would not be undermined by the truth of *CTA*.

b. Execution-Independent Arguments for Propositionalism

The second objection targets my argument’s penultimate move; the inference from the assumption that Julie *does not act on herself* in executing her intention to the conclusion that her intention is therefore not *about* her. (The ultimate move takes us from here to denying that her intention is propositional; this move seems safe given our understanding of content in terms of ‘aboutness’ – see again §1(a)). The move is legitimate only on the assumption that the *only* consideration relevant to what an intention is about is how an agent’s causal powers are directed in executing it. But, the objection continues, there are independent considerations which suggest that Julie’s intention is about herself. I will consider two.

First, it looks like a datum that Julie’s intention can be executed only *by Julie*. And this might be taken to suggest that her intention is *about* herself; is about *her* sentencing Bev and not just about *sentencing Bev*. If her intention is about *her* sentencing Bev, then it looks like a propositional intention after all; an intention *that she will sentence Bev*.

I agree that the datum needs accommodating: intentions can *only ever* be executed by their subjects.²⁴ But it is not best accommodated by building reference to the intender into the content of the intention. This would mean that *all* intentions are about the intender, and this seems implausible. Julie’s intention for Bev to go to prison doesn’t seem to be about

²⁴ Don’t I execute your intention if I do what you tell me? No. I obey your command. We do make this distinction and ought not to ignore it.

Julie in the way that her intention to dye her own hair is about her.²⁵ If we can accommodate the datum without making what seems like an implausible claim, then we should.

And we can – very easily in fact. For the datum falls out of something we are already committed to, *viz.* the commonsense requirements on intention-execution from §1(b). According to (ii.a), an intention would be *executed* (and not merely *satisfied*) only if what it is about is brought about *by the intender herself*. That entails that only Julie can execute her intention for Bev to go to prison, even though this intention’s content makes no reference to Julie, and the same goes for her intention to sentence Julie. Since our commonsense conditions on intention-execution explain our datum independently of any assumptions about the intention’s content, the datum gives us no reason to think of Julie’s intention to sentence Bev as being about Julie, and so no independent reason for Propositionalism.

Another argument is given by Hector-Neri Castañeda.²⁶ The thought is that intentions must contain reference to the subject because their canonical expressions do. Julie might express her intention to sentence Bev by asserting “I am going to sentence Bev”. But Castañeda’s argument ought to be rejected, for it is not in general the case that what someone says when they express an attitude matches the content of that attitude. This principle would get us the contradictory result that my (first-order) desire *for a coffee* (or: *to have a coffee* or *that I will have a coffee*), when expressed in my (plausibly canonical) assertion “I want a coffee” is identical with the (second-order) desire *for wanting a coffee* (or: *to want a coffee* or *that I want a coffee*). So the principle that we can *in general* read off the content of attitudes from the contents of their canonical expressions leads to contradiction in the desire case, and must be rejected.²⁷

Of course is open to the Propositionalist to suggest other independent reasons for thinking Propositionalism true, even though in executing certain intentions we do not act on ourselves. I have considered and rejected only the suggestions I think are most likely to be made. But it is important to recall that the current strategy for defending Propositionalism has *two* parts, not just one. For not only does the Propositionalist need to find a plausible reason, independent of the considerations in §2(d), to think that the intentions I say are about only *acting in some way* are in fact about *the intender’s own acting* in that way, she would also need to meet the first two explanatory challenges I described there. And it looked hard to do this from a Propositionalist perspective. It looks hard, that is, to give membership conditions for the privileged set of $I(\Phi_i)$ ’s; those which can be executed without self-causation, and therefore to explain what it is about the intentions in this privileged class which *explains why* they are executable without self-causation, given that $I(\Phi_x)$ ’s in general – like Julie’s intention *for Bev to go to prison* – will be executed only by causing the x to Φ .

²⁵ The implausible claim that all intentions are about *the intender* ought to be distinguished from the plausible one that all intentions are about *doing something*, about *acting* in some way. My argument has required assuming that some intentions are *not* about *doing something* but are about *a state of affairs* (Bev’s going to prison, for example). For arguments that intentions are *always* about acting in some way, see Baier, “Act and Intent,” 649 and Clark, “The Action as Conclusion,” 501–2. Cf. Vermazen, “Objects of Intention.”

²⁶ Castañeda, “Intentions and Intending,” 141.

²⁷ The principle does seem to work in the case of belief, but for reasons just explained, we should see this as a special case, and not as a paradigm, of the semantic relationship between an attitude’s canonical linguistic expression and its content. I will come back to the idea of linguistic expression in Chapter Six (esp. §4(c)).

In light of these issues it seems hugely preferable to accept that at least some intentions are not propositional attitudes, but attitudes towards doing things, or towards things (actions) which are to be done.²⁸

c. *An Irreducibly Agential "I"*

One thing my argument picks up on is that by intending something someone is in the capacity of *agent*, whereas *what* she intends is in the capacity of *patient*. One way of putting my conclusion would be to say that executing any intention requires executing at least one intention in relation to which the subject appears *only* as agent. I *can* act on myself, and my intentions *can* be executed by my so doing – my intention to get up early is like this if executing it requires setting an alarm the night before. If my argument works, it shows that it is impossible for *all* intentions to be like this: some intentions must include me only as agent (i.e. as *intender*), not also as patient (as part of *what is intended*).²⁹

The final objection is that I have ignored a theoretical possibility, *viz.* that not all I(Φ i)'s are made the same; some contain a special "I" (understood pronominally) which represents the intender in an 'irreducibly agential' way.³⁰ My intention to get up early tomorrow contains the usual "I", which represents the self as a patient in just the same way that Bev is represented as a patient in Julie's intention *to sentence Bev*. But the "I" in Julie's intention to sentence Bev – understood as her intention *that she will sentence Bev* – is different. This "irreducibly agential" "I" picks out Julie *qua agent*. The execution of intentions containing this special "I" is incompatible with the intender's being acted on – and so acting *on themselves* – as patient.

If this idea works, it affords the Propositionalist responses to *all three* explanatory challenges of §2(d). The class of privileged I(Φ i)'s which are executable without self-causation is just the class containing the irreducibly agential "I". And it is *because* the self-reference they involve is irreducibly agential that they are executable without self-causation; indeed *what* is intended will not be brought about if the intender acts *on herself*, because the state of affairs thereby instantiated is *not* one which contains the intender only *qua agent*, but one in which she is just as much *patient*.

But the proposal is hard to accept because it is hard to make sense of as a stable position. It is hard to get a grasp on the view without it's collapsing either into the view I think we should adopt – one on which there is *no* reference to Julie in the content of her intention to sentence Bev – or into the version of Propositionalism which I have been arguing against, which sees Julie's intention to sentence Bev as concerning Julie, *qua patient*.

It is not easy to explain the worry here, in part because I find it hard to understand the proposal. But perhaps we can (literally) use our imaginations to a grasp on what the 'irreducibly agential "I" is supposed to be. I suppose that however it is supposed to work, this "I" is part of the representational content of an episode of imagining oneself *doing*

²⁸ This positive suggestion is developed in Clark, "The Action as Conclusion."

²⁹ This is close to how Baier puts things (Baier, "Act and Intent," 658).

³⁰ Ryan Cox, Alison Fernandes and Chris Cowie have all (independently) suggested this to me. Castañeda (Castañeda, "Intentions and Intending," 141–144) has a similar view in mind, except that his special "I" is not irreducibly *agential*, but irreducibly *first-personal*. I don't respond to Castañeda's suggestion because I think that positing an irreducibly *agential* "I" seems better-placed as a response to the present worry.

something. This is not intended as a definition; just a way-in to understanding what we are talking about. I imagine myself feeding Egbert (the hamster) a mealworm. Because I am imagining *myself* feeding him a mealworm, what I am imagining must – I suppose the thought goes – contain reference to me. But since I am not imagining myself as in any way *acted on*, it must contain irreducibly *agential* reference to me.

I find it hard to see how thinking of my episode of imagination as being about myself, under an irreducibly agential mode of presentation, is supposed to be any different from thinking of it as not being *about me* at all, but about only Egbert and mealworms and feeding. It is true that what I am imagining is different to what I would imagine if I imagined, say, *Alexander* feeding Egbert a mealworm, but we can account for this difference without thinking of the imaginative episode as being *about me qua* mealworm feeder; we just need to say that it's *not* about Alexander - or anyone else - *qua* mealworm feeder. Additionally, when I explicitly direct my imagination *at* myself – at *my feeding Egbert a mealworm* – what comes to mind is an imagined scenario from a perspective which is outside myself *qua mealworm feeder*; I imagine myself feeding Egbert the mealworm *from an outside perspective*, as if during an out-of-body experience.³¹ My perspective *qua* imaginer is distinct from my perspective *qua* imagined mealworm feeder; it is the perspective of someone perhaps in the top left-hand corner of the room.

The current objector is bound to complain that I am not taking her suggestion seriously. The irreducibly agential “I” is supposed to be something we are all familiar with. It is the “myself” I imagine when I imagine *myself* feeding Egbert and when my perspectives *qua* imaginer and *qua* imagined coincide. But I can't see *why we are to think of* this episode of imagination as being *about me* at all and not just *about feeding Egbert a mealworm*. Why is it not enough for me to *imagine myself* feeding Egbert that my perspective *qua* imaginer coincides with my perspective *qua* imagined? And we are back to square one, since the current objector will presumably say that this *is* enough, but that *by* imagining things from the mealworm feeder's perspective, one *is employing* the irreducibly agential “I”. And I will complain that I don't understand the difference between our positions, and we'll be going round in circles.

I don't know how to bring this argument to an end. I suspect the suggestion of an ‘irreducibly agential “I”’ is incoherent *if* it means that one is at the same time treating oneself as an object *of thought* but not as an object *of action*. But I don't know how to argue this except by gesturing, as I have done in the previous two paragraphs, at why I – personally - find the idea hard to grasp. I can only hope the reader finds it hard too, and that she also shares the sense that it is to be explained by the incoherence of the idea of an ‘irreducibly agential “I”’ which shares a syntax with the ordinary “I”, rather than by personal imaginative limitations.

Luckily, I think I can bypass these loose ends. For we would only *need* to make sense of an irreducibly agential “I” if we were antecedently convinced of the need to think propositional, those intentions which I have argued are not. It is likely that we will have such a reason only given independent reason to think Propositionalism true. In the

³¹ Some related issues – albeit from a more empirically-minded point of view than the one relevant to our discussions – are discussed in Marcel, “The Sense of Agency: Awareness and Ownership of Action.” Marcel argues that empirical data from abnormal psychology support the idea that the content of the experience of agency does not include reference to the subject. For discussion see Roessler and Eilan, “Agency and Self-Awareness: Mechanisms and Epistemology,” 43–46.

absence of such a reason, and given the argument against Propositionalism in §2, I think we can safely ignore the suggestion about the irreducibly agential “I”.

d. Conclusion

This Chapter has argued that Propositional is false; that if any intention stands a chance of execution, it does so only *via* the execution of some non-propositional intention. Although I assumed for the sake of argument that some intentions *are* propositional attitudes, it is consistent with my conclusion that *none* of them are, and it is a consequence of my argument that in their most fundamental form, intentions are non-propositional attitudes. The possibility of executing intentions depends on the existence of non-propositional intentions. The argument of this chapter will lead to a new constraint on an account of practical knowledge, which I will introduce in concluding this Part of the Thesis.

Concluding Part Two

Velleman's and Setiya's accounts of practical knowledge depended centrally on Cognitivism, the idea that intentions are a kind of belief. Holding that intentions are a kind of belief looks *prima facie* a clever move, one which enables a theorist to explain both why practical knowledge has *unusual* features compared to knowledge of others' intentional actions or one's own non-intentional ones but also why it is nonetheless still a kind of *knowledge*. On such a view, practical knowledge is *unusual* because it is constituted by beliefs of a funny kind, but it is *knowledge* because it is constituted by *beliefs* nonetheless. We saw some internal problems with Velleman's and Setiya's accounts in Chapter Three (Velleman's view came out worse than Setiya's), but it might be thought that the general strategy of understanding practical knowledge as knowledge in intention, where intentions are a funny kind of belief, is a good one and worth pursuing in different guises.

But if the argument of Chapter Four is correct, the general strategy is unworkable. If intentions are not *propositional* attitudes – at least not fundamentally so – then intentions *are not* a kind of belief, since beliefs *must be* propositional attitudes because they must be truth-apt. Any version of *CI* will flout what I will call the *Intention-Content* constraint:

The Intention-Content Constraint

An adequate account of practical knowledge must not entail that intentions are propositional attitudes.

Part One motivated taking some version of *Intentionalism* seriously; the idea that practical knowledge is *constituted by* intention. And the current Part has motivated *rejecting* the idea that intentions are a kind of belief; it has rejected *Cognitivist Intentionalism*. Part Three will consider a version of *Intentionalism* which rejects *Cognitivism*, first *via* an interpretation of Anscombe's view in *Intention* in Chapter Five, and then as developed in my own way in Chapter Six.

Part Three

A Better Intentionalism

This Part introduces, develops and defends two versions of the idea that an agent's practical knowledge is constituted by her intention, where intentions are conceived as distinct attitudes from beliefs.

Chapter Five provides an interpretation of Anscombe's account of practical knowledge in *Intention*, which I call *Non-Cognitivist Intentionalism* (NCI). I argue that NCI happily meets the *Apparent Features* and *Special Relationship* constraints, and does so in a unified way. NCI also meets a further constraint which is introduced and explained: the *Structure-Sensitivity* constraint demands of an account of practical knowledge that it should explain how practical knowledge takes in both present and future intentional actions, and that it is sensitive to the agent's practical reasons. But it is not clear how NCI meets the *Why Knowledge?* constraint, nor how it meets the *Intention-Content* constraint introduced in my Conclusion to Part Two, consistently with a final new constraint, the *Knowledge-Content* constraint, which demands that an account explain why practical knowledge should be a kind of propositional knowledge.

Chapter Six introduces and defends *Non-Propositionalist Intentionalism* (NPI) which is a stronger view than NCI, and entails it. NPI takes over the benefits of NCI, and – I argue – in addition meets the *Why Knowledge?*, *Intention-Content* and *Knowledge-Content* constraints. I conclude that NPI deserves to be taken seriously as an account of practical knowledge.

Chapter Five:

Anscombe's Non-Cognitivist Intentionalism

Introduction

I have chosen to consider Anscombe's account last for a couple of reasons. First, I think Anscombe provides a *unified* way of meeting various of our constraints, the benefits of doing which are best seen against the background of the discussion we have had up until now. And second, because I think the account is roughly right, at least if developed in a certain way. I will develop it in this way in Chapter Six, so it makes sense to consider the view I find in *Intention* immediately prior to doing that.

Intention is famously opaque, and although I do think the view I describe here as 'Anscombe's' stands a fair chance of actually *being* her view – at least in relation to its major elements – the focus of this chapter is not exegetical but philosophical. My aim is to describe the view I have found in *Intention* which is also a view which I think is broadly correct. So I will not be defending my interpretation against alternatives, although I will try to make it clear throughout how I get to my interpretation from what Anscombe actually says in *Intention*.¹

The main features of Anscombe's view, as I understand it, are set out in §1. This involves explaining how two different characterisations Anscombe gives of intentional action converge. On one of these, intentional action is action to which a special sense of the question 'Why?' has application; on the other it is action which is the object of practical knowledge. By tracing the interrelations between these three concepts – intentional action, the special sense of 'Why?' and practical knowledge – we will gain a deeper understanding of all three. For a central element of Anscombe's view about practical knowledge is that it is conceptually related to these other two phenomena. Throughout §1 we will see that the notion of a behaviour's *having a point* is of central importance to understanding how these three phenomena interrelate, and so – for our purposes – of central importance to understanding practical knowledge. Practical knowledge will be understood as formally distinct from theoretical ('ordinary') knowledge in two fundamental ways: it is undermined not by an unsuccessful judgment, but an

¹ Anscombe's view is often linked with Stuart Hampshire's: both think that practical knowledge is importantly linked to intention; both deny that it is based on epistemic reasons. I have chosen to concentrate on Anscombe's view for a couple of reasons. First, it is a very important feature of Anscombe's view as I understand it, and of mine, that practical knowledge is *constituted by* intention, and although Hampshire says much which suggests that this is his view (see *e.g.* Hampshire, *Freedom of the Individual*, 53), he also suggests in places that it is constituted by *belief* (see *e.g.* Ibid., 68), which makes interpretation tricky. Second, our investigation, like Anscombe's, begins with present-directed practical knowledge, whereas Hampshire seems to think that practical knowledge is in its fundamental form of *future* actions (Hampshire, "Reply to Walsh on Thought and Action," 411). A very interesting take on what Hampshire is up to is given in Roessler, "The Epistemic Role of Intentions."

unsuccessful performance, and rather than being ‘derived from the objects known’, it is ‘the cause of what it understands’.

In §2 I explain the benefits of Anscombe’s view. I will explain why the two fundamental formal features of practical knowledge enable her account to meet the *Apparent Features* and *Special Relationship* constraints, and to do so in a unified way. We will also set out a further constraint on any account of practical knowledge – the *Structure-Sensitivity* constraint – which requires an account to explain how practical knowledge is sensitive to the *internal* temporal and rational structure of its object, intentional action. I also consider explicitly how we should understand Anscombe’s claim that practical knowledge is knowledge ‘in intention’. I will explain that understanding her as thinking of practical knowledge as constituted by intention *rather than* by belief gives us a unified explanation of why practical knowledge would have the two formal features described in §1. And if this is right then her account also happily meets our new *Structure-Sensitivity* constraint too. Her account will be identified as a *Non-Cognitivist Intentionalism*, NCI.

But the view is not without problems, three of which I consider in §3. The first – that the account’s conception of the relation between intentional action and practical knowledge is overly strong – is dismissed. The second and third are deemed more serious. The second is that Anscombe’s account is bound to flout the *Why Knowledge?* constraint and the third is that it can only meet the *Intention-Content* constraint by flouting a new, final, constraint. The *Knowledge-Content* constraint requires an account of practical knowledge to leave it unmysterious why practical knowledge should be a kind of *propositional* knowledge. These two problems are left to be taken up in Chapter Six.

1. *Intentional Action, Practical Knowledge, ‘Why?’*

Associated with Anscombe’s *Intention* are two characterisations of intentional action, which is understood first as action “to which a certain sense of the question ‘Why?’ is given application”² and second, as the formal object of practical knowledge.³ *Prima facie* it is not obvious why these two characterisations should converge on a single phenomenon. In this section I want to explain why, on Anscombe’s view of things, they do converge, and in so doing elucidate the three notions – intentional action, the special sense of ‘Why?’ and practical knowledge – in terms of one another.

In (a) I relate the idea of ‘Why?’ questions to the idea of kinds of explanation, and explain the relevance to intentional action. In (b) I relate the applicability of Anscombe’s special sense of ‘Why?’ to an action’s *having a point*. In (c) I explain that for Anscombe, practical knowledge displays a formal feature (“Theophrastus’ Feature”) which means it can *only* take in behaviours which have a point. In (d) I identify intentional action as action which gets its point *in a special way* and, and in (e) I explain how this relates to a second formal feature of practical knowledge (“Aquinas’ Feature”). I sum up in (f). This section seeks only to set out Anscombe’s view; I turn to considering its merits in §2 and its challenges in §3.

² Anscombe, *Intention*, 9.

³ “Without [practical knowledge] what happens does not come under the description – execution of intentions – whose characteristics we have been investigating.” (Ibid., 87–8.; see also Ibid., 11.)

a. Intentional Action, the question 'Why?_{RA}' and a Special Kind of Explanation

In saying that intentional actions are “actions to which a special sense of the question ‘Why?’ has application”,⁴ Anscombe characterises intentional actions as actions to which a special kind of *explanation* is appropriate. For asking ‘why?’ of some phenomenon is asking for an explanation of it, and asking ‘why?’ in a certain sense is asking for a certain kind of explanation.

Asking “Why did Daisy blink?”, for example, might be asking for a *teleological* explanation of her blinking (asking for the *aim* of her blinking), or for an *efficient-causal* explanation (asking which prior events *prompted* her blinking). There are other kinds of explanation too. One of Anscombe’s key thoughts is that proper and peculiar to intentional action is a *sui generis* kind of explanation, the explanation of behaviour in terms of practical reasons, or reasons for action (I’ll use these expressions interchangeably). In the sense Anscombe is interested in, the question ‘Why?’ seeks an explanation of this kind. From now on I will use “Why?_{RA}” to indicate the sense of ‘Why?’ which seeks an explanation in terms of the agent’s *reasons for action* and “Why?” (*sans* subscript) to indicate the question in some other sense, or neutrally between its various senses - context will disambiguate.

If we want to understand intentional action, thinks Anscombe, we need to understand the special kind of *explanation* to which it is internally related, which means at the same time understanding the sense of ‘Why?_{RA}’. Although the special kind of explanation is explanation in terms of *reasons for action*, and ‘Why?_{RA}’ asks for these, saying this is not enough. Because just as there are various kinds of explanation, each linking with a different sense of ‘Why?’, there are also various kinds of *reason*, each being what might be given in explaining a phenomenon, each demanded by a different sense of ‘Why?’. So Anscombe must elucidate these concepts together, which she does by considering when the question ‘Why?_{RA}’ is refused, and when it is granted, application. For reasons of space I am going to concentrate on the responses which *grant* ‘Why?_{RA}’ application.

b. Granting Application to the Question 'Why?_{RA}'

To grant any ‘Why?’ question application is to implicitly characterise the queried phenomenon as *of the right kind* to be explained in the way in question. So granting ‘Why?_{RA}’, asked of a behaviour, is implicitly characterising that behaviour as of the right kind as to be explained by reasons for action (in a special sense which we will fully understand only at the end of (d)). Actions of this kind are intentional ones. So granting application to ‘Why?_{RA}’ asked of a behaviour characterises that behaviour as an intentional action.

There are two main ways in which ‘Why?_{RA}’ might be granted application: positively and negatively. I take each in turn.

Answering ‘Why?_{RA}’ positively involves stating or suggesting a reason for action of the special kind. Amongst these special reasons for action, Anscombe distinguishes two kinds. Stating or suggesting a reason might be stating or suggesting an *aim* or *intention*, or it might be stating or suggesting a *motive*. Both categories are once further bisected. In giving an intention one might indicate something *future* or something *present*. In either

⁴ Anscombe, *Intention*, 9.

case, one states an action or state of affairs at which the queried behaviour aims. A future intention is given “if you say ‘[Why_{RA}] are you crossing the road[?]’ and I reply ‘I am going to look in that shop window’”.⁵ A present intention is given when the question “Why_{RA} are you lying on that bed?” is met with “I’m doing yoga”.⁶

Motives are reasons which do not characterise an aim. The category is split into ‘backwards-looking’ and ‘interpretative’ motives (although a person’s motives can be a mixture of the two kinds⁷). Giving a *backwards-looking* motive is indicating that some event has happened in the past, which is (perceived by the agent as) good or bad in some respect and thus thereby invites a certain kind of response. “Why are you burning down Jones’ house?” might be met with “He killed my father,” thus showing the house-burning as motivated by the backwards-looking motive of *revenge*. Other examples of backwards-looking motives are “gratitude, and remorse, and pity for something specific”.⁸ To give an interpretative motive is not to suggest any particular past event to which the queried action is represented as an appropriate response, but “is to say something like ‘See the action in this light’”.⁹ Anscombe gives the examples of love, curiosity, despair and vanity.

One can grant application to ‘Why_{RA}’ *positively* – by explaining one’s action in terms of one of these kinds of reasons – without explicitly *stating* one’s reason. One might instead put one’s action in a context which makes its reason clear, to a greater or lesser extent. If I respond to the question “Why_{RA} are you working on a Sunday?” with “I have so much to do!” I make clear that I am working with the intention of *getting enough done*. And if our arsonist responds to “Why_{RA} are you burning that house down?” with “It’s Jones’ house”, he will make it clear to someone who knows what Jones has done that it is an act of revenge. Explaining that he is oiling the chain of the bike *because it’s Luke’s bike*, Peter makes it clear – to me at least – that his motive is love.

I distinguish these different kinds of answer to ‘Why_{RA}’ because I want to explain what they have in common. They all *explain* or *articulate* the *point* of the queried behaviour.¹⁰ Citing an aim or intention explains a behaviour’s point by representing it as a way of achieving that aim: to look in the shop window, I need to cross the road; lying on the bed like that is a way of doing yoga. Backwards-looking motives explain a behaviour’s point by representing it as an appropriate response to a particular (good or bad) past event. You might not agree that revenge is an appropriate response to Jones’ having killed my father – you might advocate the formal justice system instead. But in understanding my arson *as* revenge (and recognising that vengeance matters to me), you will understand the point of burning his house down *from my* perspective. You will understand that I see my arson as serving as a kind of leveller; as restoring justice between us. Finally, giving an interpretative motive explains the point of the queried behaviour by giving the light under which the behaviour’s point will be apparent. By giving *love* (say) as a motive for his action, Peter thereby suggests that love matters to him, and that the queried behaviour matters insofar as it is an expression of love. The same goes for less admirable motives:

⁵ Ibid., 35.

⁶ In Anscombe’s example the question is actually “*What* are you doing?” (Ibid., my emphasis). The possibility of switching between ‘what?’ and ‘why?’ is of great importance, as we shall see in §1(e) and again in §2(b).

⁷ Ibid., 21.

⁸ Ibid., 20.

⁹ Ibid., 21.

¹⁰ Anscombe characterises practical reasons as “reasons why it would be useful or attractive if the description [of the action] came true, not by evidence that it is true.” (Ibid., 6); we shall see in (d) that this is in fact only a partial characterisation of the reasons particular to intentional action.

if Luke says that his action is done out of *vanity* then although he might say that vanity *oughtn't* matter to him, he can't consistently deny that it *does*. The action has a point because it satisfies his vanity.

Let's turn to the *negative* way of granting application to 'Why?_{RA}'. Here, one answers in such a way as to state or suggest that one's action *was not* done for a reason:

Now of course a possible answer to the question 'Why?' is one like 'I just thought I would' or 'It was an impulse' or 'For no particular reason' or 'It was an idle action – I was just doodling'.¹¹

How can answering in this way be *granting* application to the question 'Why?_{RA}', given that the question seeks the reason for which one's action was done? The answer is that granting application to the question represents one's action as *of a kind* which *could* be explained in terms of its practical reasons. But an action's being of this kind does not in all cases mean that it *has* such an explanation.

We can use the notion of a behaviour's *having a point* to make the relevant distinction between a behaviour's *being of the right kind* to be explicable in terms of practical reasons, and a behaviour's *actually having an explanation* in terms of (being done for) practical reasons. Reasons for action articulate the point of an action. But not all actions have a point which *is articulate*; about which there is anything to be said in explanation of it. Sometimes an action will have a point which is so to speak 'atomic'.¹² Say I do an impression of Bob Dylan because I just thought I would: "I just felt like it." In responding to 'Why?_{RA}' thus, I imply that there is nothing to be said about *why I felt like it* but the *fact that I did* feel like it means that my *doing it* has a point. I do not do it *in order to satisfy some desire*, on Anscombe's view (I recognise that this is a controversial description of the case, but here I am restricting my attention to a *statement* of Anscombe's view as I understand it). In saying I did it simply because I felt like it I represent my action as having a point which is atomic; which does not admit of its own articulation or explanation. Actions like this have their *own* point. The action is done, as we are apt to say, *for its own sake*.¹³ I will return to these kinds of case in (c).

Let's sum up the important points. First, the question 'Why?_{RA}' is given a *positive answer* when a reason of some kind is stated or suggested. Reasons can be intentions, which can be future or present-directed, or they can be motives, which can be backwards-looking or interpretative (or they can be mixed). Reasons of all kinds *articulate the point* of an action, so giving a reason represents the queried behaviour *as having* a point. But some actions have an *atomic* point, a point which has no articulation, and such actions are ones in relation to which the question 'Why?_{RA}' is applicable, but for which it has no positive response. Such actions are done for their own sake and attract responses to 'Why?_{RA}' of the kinds Anscombe lists in the passage quoted above.

c. Having a point and Theophrastus' Feature

We have just learned that in granting application to the question 'Why?_{RA}', one represents one's behaviour as *having a point*. Here I want to suggest that only behaviours which *have a*

¹¹ Ibid., 25.

¹² Discussions of actions like this can be found in Raz, "Agency, Reason, and the Good" and in Alvarez, "Acting Intentionally and Acting for a Reason."

¹³ Cf. Alvarez, "Acting Intentionally and Acting for a Reason."

point could be objects of practical knowledge. Thus we will have made a start on explaining how actions to which ‘Why?_{RA}’ has application, and actions which are the objects of practical knowledge, might be actions of the same kind.

It is plausible and widely accepted that knowledge (at least knowledge of contingent facts, of which practical knowledge is a kind) presupposes the possibility of error; that “there is point in speaking of knowledge only where a contrast exists between ‘he *knows*’ and ‘he (merely) *thinks* he knows’.”¹⁴ We can link this idea with the idea of knowing as involving a certain kind of *success*: knowing is succeeding where one might have failed, whereas someone who merely thinks that they know has failed where they might have (and think they have) succeeded. Knowing is a kind of success and the avoidance of a certain kind of error.

A familiar way of thinking about the kind of success involved in knowing is in terms of one’s judgment matching the facts. A *successful* judgment is one which matches the facts, where the standard for success – and so for knowing – is set by how the facts are. If one’s judgment does not match the facts it is, for that reason, mistaken and so not knowledge. Anscombe observes that “modern philosophy” thinks of *all* knowledge as fitting this model, that:

Knowledge *must be* something that is judged as such by being in accordance with the facts. The facts, reality, are prior, and dictate what is to be said [judged], if it is to be knowledge.¹⁵

But in her view this is a mistake. Commonsense recognises a *different* way in which a fact-judgment mismatch might be normatively assessed:

In some cases the facts are, so to speak, impugned for not being in accordance with the words [judgment], rather than *vice versa*.¹⁶

On the assumption that knowledge involves a kind of success and the avoidance of a kind of error, the fact that there are two ways of normatively assessing a fact-judgment mismatch entails the logical possibility that there are two kinds of knowledge, each involving a *different kind* of success and the avoidance of a *different kind of* error. And exactly this thought is part of Anscombe’s conception of the difference between practical and theoretical knowledge. Theoretical (‘contemplative’¹⁷; ‘speculative’¹⁸) knowledge requires a successful matching between a person’s judgment and the facts, with the facts setting the standard for success, and the judgment being successful or otherwise depending on whether it meets this standard. Practical knowledge also requires a successful matching between judgment and fact, but here it is the *judgment* which sets the standard for success, and *the facts* which constitute a success insofar as they meet this standard, and a mistake insofar as they don’t.¹⁹

¹⁴ Anscombe, *Intention*, 14. *Knowing* is thus distinguished by Anscombe from *being able to say* (Ibid., 13).

¹⁵ Anscombe, *Intention*, 57; my emphasis.

¹⁶ Ibid., 4–5. See also Hampshire and Hart, “Decision, Intention and Certainty,” 9.

¹⁷ Anscombe, *Intention*, 57.

¹⁸ Ibid., 87.

¹⁹ I am being deliberately non-committal here about what ‘judgment’ means for Anscombe, because I want to describe the formal differences between practical and theoretical knowledge (as Anscombe sees them) independently of considerations about the identity of each kind of knowledge’s constituting attitude. Later on, as one might guess from the title of this Chapter, I will suggest that we ought to think of Anscombe as thinking that theoretical knowledge is constituted by belief and practical knowledge as constituted by intention, and that ‘judgment’ is to be read neutrally between these two attitudes.

But how could *the facts* be in error, or constitute a success, depending on someone's *judgment*? Surely the facts *just are as they are*, independently of what anyone thinks about them. How could they be answerable *to us*?

But this intuition ignores those facts which are about what we *do* - these facts *are* dependent on us. And it is, of course, *these* facts which are relevant to our investigation of practical knowledge. Indeed Anscombe elaborates her observation that sometimes 'facts' (rather than words or judgments) are impugned for a fact-judgment mismatch with an example in which it is clear just what *kind* of 'fact' can bear this responsibility:

e.g. I write something other than I think I am writing: as Theophrastus says (*Magna Moralia*, 1189b 22), the mistake here is in the performance, not in the judgment.²⁰

Read too literally, Anscombe's claim that 'the facts' can be impugned for not being in accordance with a person's judgment is misleading because it suggests that we should identify a person's performance with the fact that they are performing as they are, and this doesn't seem right. But Anscombe said only that in such cases the facts are "*so to speak*"²¹ impugned, indicating that she is speaking loosely. The idea is not the barely intelligible one that facts independent of human involvement are sometimes in error for not turning out how someone thought they would, but the truism that *a person* can get things wrong in two different ways; they can make a mistake in judgment, *or* in performance. The facts are 'so to speak' in error if I think I am writing 'causal relationship' but am in fact writing 'casual relationship'. In this case I fail to *know what I am writing* because I am *doing* something other than what I think I am doing.²²

I will refer to this feature of practical knowledge as 'Theophrastus' Feature'. Displaying 'Theophrastus' Feature' – being such as to be undermined by a mistake in *performance* – distinguishes practical knowledge *formally* from theoretical knowledge. Practical knowledge is undermined by a mistake in *performance* whereas theoretical knowledge is undermined by a mistake in *judgment*. Practical knowledge thus involves a *success in performance* whereas theoretical knowledge involves a *success in judgment*.

We can now start to link up Anscombe's characterisation of intentional action as action to which 'Why?_{RA}' has application, with her characterisation of it as an object of practical knowledge. We have already seen that in granting application to 'Why?_{RA}' one represents one's behaviour as *having a point*. And if practical knowledge displays 'Theophrastus' feature, then it can take in *only* action which has a point. Let's see why this is.

Some of the things creatures do constitute successful or mistaken performances. But not *all* the things creatures do can constitute successes or failures. Whilst typing I am making tapping noises with the keys. If one of the keys fails to tap when I hit it because a bit of fluff has got underneath, or because I haven't hit it quite hard enough, I don't thereby act *in error*. But writing 'casual' rather than 'causal' *was* acting in error. What's the difference?

The difference is that it doesn't *matter* to me whether the key makes a tapping noise or not, whereas it does matter whether I write 'causal' or 'casual'; making a tapping noise *has*

²⁰ Anscombe, *Intention*, 5.

²¹ *Ibid.*, 4; my italics.

²² It is much more natural here to say that I fail to know what I am writing because I am not writing *what I mean* or *intend* to be writing. And indeed this will be the paraphrase I will suggest in §2(d). (See again n. 19.)

no point whereas writing ‘causal’ does (*ex hypothesi*, of course – another example could be built in which making the tapping noise mattered, whereas what I wrote did not). Proof can be found in how I might act when I notice I have not done each of these things. Noticing I have not typed ‘causal’ (but ‘casual’), I go back and *correct* what I have written. But noticing that a particular key did not tap, I *do not* go back and correct this. There is nothing here *to* correct because whether I made a tapping noise or not didn’t matter; it didn’t have a point.

The thought holds equally for actions with an atomic, as with an articulate point. Let’s say the point of writing ‘causal’ is explained by saying that I am writing about a certain causal relationship, in order to defend a premise in an argument I am making. The point of writing ‘causal’ is here articulated – there are various things we can say about why it matters that I write ‘causal’ and not ‘casual’. There was nothing to say about why doing my Bob Dylan impression mattered – I *simply* felt like doing it; its point was *atomic*. But *doing the impression* nevertheless *did matter* to me in a way that making a tapping noise with the keys as I type does not. We can see this by noting that if my impression sounds more like Tom Waits than Bob Dylan it will be a mistake which I could try to correct (although I needn’t in fact go on to do so).

So far we are able to link the characterisation of intentional action in terms of the applicability of ‘Why?_{RA}’ with the characterisation as the object of practical knowledge *via* the notion of the relevant actions *having a point* as follows: intentional actions have a point, and behaviours are represented as having a point when ‘Why?_{RA}’ is granted application. And because practical knowledge displays Theophrastus’ Feature, it can take in *only* actions which have a point, since only actions which have a point can constitute a successful performance.

It is tempting at this stage to *identify* behaviours which *have a point* with *intentional* actions. But we cannot do this, because not all actions which have a point *are* intentional. In order to understand intentional actions properly, we need to understand the special way in which they *get* their points. To this I now turn.

d. Intentional Action and Egocentricity

Not all actions which have a point are intentional actions. But we can distinguish cases in which an action has a point independently of its agent’s recognising it as having this point, from cases in which an action’s having the point it has *is dependent* on the agent’s recognising it as having this point. When an action gets its point in the second way, I will call its point *egocentric*, and I will say that we ought to identify intentional actions with actions whose point is egocentric. Egocentricity is also a feature of the *reasons* particular to intentional action, and to the type of *explanation* one gives when one explains one’s action in terms of such reasons, as I shall explain.

Daisy the cow blinks. Her blinking has a point - to keep an advancing speck of dust out of her eye. Her blinking is not intentional. I wink. My winking has a point – to greet the approaching Jones. My winking *is* intentional. The difference is that Daisy’s blinking has the point it has independently of Daisy’s thinking of its having this point, whilst my winking has its point only because I think of it as having this point. My winking is informed by my recognition of the point it would serve, whereas Daisy’s blinking is not.

Daisy, being a cow, lacks the conceptual sophistication to recognise the point of her blinking. But it is not simply that she fails to recognise the point of her blinking which makes it non-egocentric in my sense. Compare a case in which *I* blink. I might be aware of the point of doing so, and even endorse my blinking on the basis of this recognition, thinking to myself what a good job it is I have this reflex, and how I value having clear and pain-free eyes. But even though I recognise and endorse the point of my blinking, it is still *like* Daisy's blinking and *unlike* my winking in that its *having* the point is independent of my recognition, and of my endorsement.²³ Positively, the point of my and Daisy's blinking is dependent on facts about biology and evolution, and not on facts about either of our conceptions of this point. The point of blinking is *not egocentric*, meaning that there is no sense in which the point of our blinking is *dependent on* our recognition of its having this point. The point of my winking is egocentric: my winking would not have the point it has unless I recognised it as a way of greeting Jones, something which I also recognise as worth my while to do. (Blinking *can* have an egocentric point, although this is not the usual case. For someone who has lost their blink reflex and has to blink intentionally or not at all, blinking will have an egocentric point; Hugh Grant might blink intentionally, whilst acting in a film.²⁴)

We could not identify intentional actions as actions which have a point. But we can identify them as actions whose point is egocentric. That is not to say that intentional actions *only* have an egocentric point. It is even possible for the point of an action to be *both* egocentric *and* non-egocentric. Say I take a drink of water, the point of which is to quench my thirst. There is *an* explanation of the point of drinking water which is, like the explanation of the point of blinking (in the normal case) non-egocentric; which will advert to facts about evolution and about my biology. The point of my drinking understood in *this* way is independent of my recognition of these facts. But unless I am sleep-drinking, we can *also* explain the point of my drinking in a way that *is* egocentric, which *does* advert to my recognition that drinking would quench my thirst. So the point is not that intentional actions have *only* an egocentric point; it is this: it is insofar as an action *has* an egocentric point that it is done intentionally.

Implicit in the previous paragraph is the thought that not only the *point* of a behaviour, but also its *articulation* (when it has one) can be egocentric or otherwise. I have suggested that giving a *reason* articulates a behaviour's point. And an *explanation* in terms of reasons can also be egocentric or otherwise. In one sense, Daisy and I *both* act for reasons, in blinking and winking respectively. There is *a* sense of 'reason for acting' in which Daisy blinks *for the reason that* it will keep dust out of her eyes. But her blinking's being done for this reason does not entail that Daisy recognises this as a reason for blinking, or that she could explain her blinking by citing this reason.

My winking, on the other hand, not only *is explicable* in terms of reasons (like Daisy's blinking), but further, its being so explicable is dependent on *my being able to explain it* in terms of these reasons. Daisy's reasons for blinking are not egocentric, whilst my reasons for winking are, and the *explanation* of Daisy's blinking is not egocentric whereas the *explanation* of my winking is. My winking at Jones is explicable in terms of my greeting

²³ Similar thoughts can be found in Velleman, "Replies to Discussion on the Possibility of Practical Reason," 12–13; and in Boyle and Lavin, "Goodness and Desire."

²⁴ In Alien Hand syndrome, the problem is that the action's point – to unbutton one's shirt, for instance – is *not* egocentric; it will be disavowed by the agent as *her* aim. For a nice discussion of various perspectives on this issue see Roessler and Eilan, "Agency and Self-Awareness: Mechanisms and Epistemology."

him *only because* I think of it as so explicable. The way in which the point of Daisy's blinking articulates is not dependent on her thinking of the point as so articulating, whereas the way in which the point of my winking articulates *is* dependent on my thinking of its point as so articulating. An action for which this is not true is not explicable in terms of those *reasons for action* which are proper and peculiar to intentional action. The special kinds of reasons for action which are given in explanations proper and peculiar to intentional actions are *egocentric*.

These ideas enable me to make a clarification about Anscombe's account. Some have thought of Anscombe's view as one on which intentional action is essentially subject to *teleological* explanation. But I don't think this is right. We have already seen that she distinguishes reasons which imply an aim or intention from motives, which can be interpretative or backwards-looking. Taken at face-value, here Anscombe is denying that *all* explanations in terms of practical reasons are teleological explanations. The *sui generis* kind of explanation relevant to intentional action is not essentially *teleological* (although it can be teleological, and is when an intention is cited);²⁵ the *sui generis* kind of explanation is explanation which articulates the point of a behaviour, where both the point and its articulation are egocentric in the way I have described.

Finally, recognising the importance of the egocentricity of intentional action's point and – where these apply – its reasons and explanations in terms of these, helps us understand a feature of Anscombe's presentation which we have so far glossed over. Anscombe's special question 'Why?_{RA}' is always addressed in the first instance *to the agent*.²⁶ This is not because a bystander *cannot know* about the agent's reasons or because she could not represent another's intentional action as having a point, and so as intentional. A bystander can do all this. But only the agent's *own* recognition of her action's point, and of the way (if any) in which it is articulated, plays a role in *determining the fact that* her action *has* this point, so articulated. Intentional action is action which has an *egocentric* point, and whose explanation, if any, is also egocentric. So the *agent's own* (sincere) responses to 'Why?_{RA}' are privileged: her own conception of the explanation of her action in terms of its reasons (if any) *determines* the fact that her action *is* so explicable. Intentional actions are actions for which this is true.²⁷

²⁵ For discussion see Müller, "Backwards-Looking Rationality and the Unity of Practical Reason" and Hursthouse, "Intention," 96. Cf. e.g. Newstead, "Interpreting Anscombe's Intention §§32FF."

²⁶ Some (e.g. Velleman, "Precis of The Possibility of Practical Reason," 228) have thought that her focus on what an agent might *say* about their action is evidence of a commitment by Anscombe to behaviourism. I hope to have made it clear here why she is not best understood in this way. Anscombe's own comments about there being controls on the truthfulness of an answer to the question 'Why?_{RA}' (Anscombe, *Intention*, 43–45) attest to this too.

²⁷ Does all this mean that animals cannot act intentionally, since they cannot give explanations of what they are doing? Some have accepted that intentional action (as opposed to action that is merely *purposive*) is possible only for language-users (see e.g. Hampshire, *Freedom of the Individual*, 413; Kenny, *Will, Freedom and Power*, 20; Hamilton, "Intention and the Authority of Avowals," 34). But Anscombe thinks that animals *can* act intentionally (Anscombe, *Intention*, 5). Is this inconsistent? I can't answer this fully here, but consider the following: A squirrel might try *various* ways of getting a nut out of a bird-feeder if the first attempt doesn't work, and although she can't answer a 'why?' question about her action, we do find it natural to say that she's *aware of* failing, *trying* to find a *better* way etc. *Maybe* there's a kind of primitive egocentricity here; an egocentricity of points and reasons which attaches to certain actions of non-linguistic creatures and distinguishes them from non-intentional but purposive actions like Daisy's blinking.

e. Egocentricity and Aquinas' Feature

In (b) I explained how because (on Anscombe's view) it displays Theophrastus' Feature, practical knowledge can only take in action which has a point. Here I want to explain why a second formal feature of practical knowledge (according to Anscombe) means that it will take in only action whose point is egocentric.

Following Aquinas, Anscombe distinguishes practical and theoretical knowledge as follows:

Practical knowledge is 'the cause of what it understands', unlike speculative knowledge, which is 'derived from the objects known'.²⁸

She carries on:

This means more than that practical knowledge is observed to be a necessary condition of the production of various results; or that an idea of doing such-and-such in such-and-such way is such a condition. It means that without it what happens does not come under the description – execution of intentions – whose characteristics we have been investigating.²⁹

Velleman also characterised practical knowledge as 'the cause of what it understands' (see again Ch. 3, §1(c)), but it is important to distinguish Anscombe's and Velleman's version of the characterisation. For Velleman, the idea was that the attitudes which constitute practical knowledge – intention-beliefs – are implicated in the aetiology of intentional action; Velleman's claim is about intentional action's *efficient causes*. For Anscombe, the claim is about intentional action's *formal cause*;³⁰ about what it is for an action to be intentional. The idea is that what it is for an action to be intentional is for it to be an object of practical knowledge.

But why should this be? Is the claim meant to be purely analytic – we *just do* call 'intentional' only actions which are objects of practical knowledge? Or is there more to the idea than this?³¹ There *is* more to the idea. For if intentional actions are actions with an egocentric point, this means that their being *the intentional action they are* is dependent on their agents' thinking of them *as such*. For simplicity's sake, and for now, let's consider why this is by thinking about actions whose point is *atomic*; I will consider the situation for actions whose points are articulate later on (in §2(b)).

I do my Bob Dylan impression just because I feel like it, or for its own sake, which means that the *only* point of it is *that it is a Bob Dylan impression – a Bob Dylan impression* is what I feel like doing. *Ex hypothesi* its point is egocentric since I do it intentionally, and this means that its *having* the point it has is dependent on my *thinking of it as having* this point. And this means that my intentionally *doing a Bob Dylan impression* is dependent on

²⁸ Anscombe, *Intention*, 87.

²⁹ *Ibid.*, 87–88.

³⁰ Richard Moran gives us an additional way of understanding the difference between Anscombe's and Velleman's claims that practical knowledge is 'the cause of what it understands': for Velleman the 'what it understands' is understood extensionally, whereas for Anscombe it is understood *intensionally* (Moran, "Anscombe on 'Practical Knowledge,'" 54). The extensionalist reading seems to rule out an explanation of why actions are intentional only under a description; I *think* this may be one of the points Anscombe makes in the hard-to-understand *Intention* §19. See also Hursthouse, "Intention," 88–89.

³¹ Thanks to Richard Holton for pressing me to be clear about this.

my recognising what I am doing *as* a Bob Dylan impression. My judgment about what I am doing is required for *what I am doing* to count as intentional, not for purely analytic reasons, but because my behaviour would not be an intentional action unless its point was egocentric, and the fact that its point is egocentric means that I must judge myself to be doing *a Bob Dylan impression* if I am to be doing something whose point is just that. If I judge that I am doing a Bob Dylan impression, *and I am doing so* then my judgment and my performance match up, the performance is successful and I have practical knowledge. My practical knowledge is ‘the cause of what it understands’.

I will extend these considerations later on (in §2(b)&(c)). In §2(b) we will see that actions whose points are *articulate* are themselves internally rationally and temporally structured. And in §2(c) I will try to explain how it is the agent’s practical knowledge which is responsible for her action’s having this internal structure. Thus we will see another dimension to Aquinas’ Feature.

f. Summing up

The aim of this section has been to elucidate three key notions in Anscombe’s picture: the notion of *intentional action*, the sense of ‘Why?_{RA}’ (and at the same time the special *kind* of explanation proper to intentional action) and, of most importance to us, the notion of *practical knowledge*. I came at this task by trying to explain how Anscombe’s characterisations of one of these notions – intentional action – in terms of the other two – the question ‘Why?_{RA}’ and practical knowledge – really do converge, despite it being unclear *prima facie* why they should.

The resulting picture is of intentional action as action which has a point, and of granting application to the special question ‘Why?_{RA}’ as showing that one’s action has a point. The applicability of ‘Why?_{RA}’ is dependent on its being granted application *by the agent* because the point of intentional action is *egocentric*. Because practical knowledge displays Theophrastus’ Feature it takes in only behaviours which *have a point* and because it displays Aquinas’ Feature, it takes in only actions whose point is *egocentric*. By bringing out the various formal features of intentional action and of practical knowledge, I have tried to show how they fit together; why intentional action should be the *formal* object of practical knowledge. And I have tried to do so at the same time as spelling out the special sense of the question ‘Why?_{RA}’, and trying to make clear why ‘Why?_{RA}’ should be applicable of all and only *intentional* actions; actions which are objects of practical knowledge.

I will now turn from setting out the details of Anscombe’s account as I see it, to explaining why I think the account is worth paying attention to.

2. Merits and Clarifications

We will consider here why a view on which practical knowledge is characterised by Theophrastus’ and Aquinas’ Features is a good one. In (a) I explain why an account of practical knowledge as having these two formal features meets the crucial *Apparent Features* and *Special Relationship* constraints, and does so in a *unified* way. In (b) I turn to a discussion in *Intention* which suggests that intentional action is *internally* rationally and temporally structured. One result of this internal structure is that under a given description, an intentional action can be represented as present *or* future. This leads me

to set out a new constraint on practical knowledge: the *Structure-Sensitivity* constraint requires an account of practical knowledge to be sensitive to the temporal and rational structure of its object, intentional action. In (c) I return to the idea of practical knowledge as knowledge ‘in intention’. I explain that seeing Anscombe’s account as one on which practical knowledge is constituted by intention, *and not* by belief, makes it clear how easily it meets the *Structure-Sensitivity* constraint and importantly, that it provides us with a further unification: a unification of ‘Theophrastus’ and Aquinas’ Features. I sum up in (d).

a. The Apparent Features and Special Relationship Constraints

I’ll start by explaining how Anscombe’s account meets the *Special Relationship* constraint and move on to considering the *Apparent Features* constraint.

The *Special Relationship* constraint requires an account to explain first, why practical knowledge is only ever knowledge of *intentional* action, and second, why this relationship looks particularly tight. The fact that practical knowledge is only ever knowledge of intentional action falls out of thinking of practical knowledge as having ‘Theophrastus’ Feature - as requiring a success not in judgment but in performance, since intentional action *just is* a kind of successful performance. But it is not any old kind of successful performance, but one the point of which, and so the success of which, is determined by the agent’s recognition of its having that point (its point is egocentric). And thinking of practical knowledge as having Aquinas’ Feature – as being ‘the cause of what it understands’ - means thinking of it as a kind of knowledge which determines its object *as* an intentional action. ‘Theophrastus’ Feature means that only *intentional* actions are the right shape to be its objects, and Aquinas’ Feature means that *if* an action is an object of practical knowledge (under a certain description), *then* it will be intentional (under that description). On Anscombe’s view, then, an action will be intentional *iff* it is an object of practical knowledge, and the relationship is a conceptual one. The relationship Anscombe sees existing between intentional action and practical knowledge is certainly a special one: intentional action is the *formal object* of practical knowledge. So her account meets the *Special Relationship* constraint.

And it meets it by more than simple fiat. To re-cap an idea from §1(e), the claim that practical knowledge is the cause of what it understands is not to be understood purely metalinguistically, as the claim that we – as a matter of brute fact – tend to *call* ‘intentional’ only actions which are objects of our knowledge, and that the knowledge in such cases is given the name ‘practical’. The claim that practical knowledge is the cause of what it understands is metaphysically thick – there is something we can say about *why* it is true, which adverts to something weightier than a merely analytic connection between ‘intentional action’ and ‘practical knowledge’. It is because of certain structural features of both intentional action and practical knowledge that the two fit – essentially - together. The point of *intentional* action is *egocentric*, which is to say that one’s action cannot be intentional unless one recognises its point, and this means that the agent must judge the action to be of a certain kind, where its being of *this* kind is necessary for it to be intentional. In the example above this meant that I could not be acting intentionally under the description *doing a Bob Dylan impression* unless I recognised what I was doing *as a Bob Dylan impression*. And – returning to ‘Theophrastus’ Feature – I *would* know this just in case I made no mistake in performance (just in case the impression sounds like Bob Dylan and not like Tom Waits), which meant that as long as I *was* doing a Bob Dylan impression, I would know that I was.

The *Apparent Features* constraint requires an account to accommodate or explain away practical knowledge's apparent *FPA*, *epistemic ungroundedness* and *psychological immediacy*. Aquinas' Feature explains why practical knowledge displays *FPA*: it is only because I have practical knowledge that I am typing, that I *am* typing intentionally. Nobody else's knowledge of what I am doing can play this role in constituting my typing as intentional. This underwrites a clear first-/third-person asymmetry in relation to facts about what a person is intentionally doing. The asymmetry does not mean that I am infallible about what I am doing intentionally. I would fail to know that I am typing 'causal' if I am typing 'casual' instead. And it needn't mean that I am necessarily any more *confident* about what I am intentionally doing than someone else might be. You may not be able to intelligibly doubt that I am intentionally typing if, say, you are watching me whilst I do it. But my authority survives even in such a case. It is, in Anthony Kenny's phrase "the authority of the judge, not of the witness";³² even an *expert* witness like you who are observing me type. My authority is not so much about my being in a *better* epistemic position than you are. It is rather that my epistemic position is – as yours is not – one on which my action's being intentional *metaphysically* depends.

Anscombe is famous for thinking that practical knowledge is not based on observation or evidence, but I did not consider this when setting out her view in §1 at all. I had a reason for that – it would have distracted from my task in §1, which was to set out the main *formal* features of practical knowledge, as (I think) Anscombe sees it. And I think it is misleading to think of epistemic ungroundedness as one of its fundamental formal features.³³ This feature of practical knowledge is, for Anscombe, a *clue* to what it must be like, i.e. knowledge which displays Theophrastus' and Aquinas' Features.

But with the formal properties of practical knowledge explained, we can *now* explain why it should be epistemically ungrounded. Consider why *theoretical* knowledge would be epistemically grounded in observation or evidence (or testimony or whatever else). For knowledge which involves a successful *judgment*, the standard for which is set by the facts, observation and the appreciation of evidence can act as means for judgmental success – for getting our judgments in line with the facts. But for knowledge which involves a successful *performance*, the standard for which is set by the agent's *judgment*, there is no need to rely on observation or evidence, because one's judgment doesn't need to *do* any lining-up; that is the responsibility of one's performance. Although (as we remarked in Chapter One (§2(b))) observation and evidence might guide one's performance itself, they do not – in so doing – act as reasons for believing *that one is doing* what one is doing. So they do not epistemically ground one's practical knowledge *that one is doing* what one is doing. If what I am aiming to do is to cycle into town, feeling that there is a strong crosswind, and having my eyes open the whole way will both be very helpful – perhaps *necessary* for success – but these bits of perceptual knowledge do not act as epistemic grounds.

Aquinas' Feature is relevant here too. For if practical knowledge is the cause of what it understands, then the object of practical knowledge is simply not available to be known on *any* epistemic basis unless the agent *already* has *practical* knowledge of it. I could *only* know that I was typing 'causal' on the basis of observation if I was *not* typing 'causal' intentionally. If I am, any observations I make must be secondary to my practical – non-

³² Kenny, *Will, Freedom and Power*, 6. Moran makes a similar point about practical knowledge's *FPA* in "Anscombe on 'Practical Knowledge,'" 68.

³³ Starting from here would lead us to look "hopelessly for the different *mode of contemplative knowledge* in acting, as if there were a very queer sort of seeing eye in the middle of acting" (Anscombe, *Intention*, 57.)

observational – knowledge. The same goes for any evidence I might have to think I’m typing ‘causal’.³⁴ Thinking of practical knowledge as having both ‘Theophrastus’ and Aquinas’ Features happily explains why it would – indeed *must* – be (and so appear) *epistemically ungrounded*.

Aquinas’ Feature also explains why practical knowledge should be (and so appear) *psychologically immediate*. Aquinas’ Feature means that if I *am* intentionally typing, it is *because* I know that I am. So there is no room for *considering* whether I am, consistently with my doing so intentionally and there is no room to *notice, find out, discover (etc.)* that I am intentionally typing. I can find out that I am *doing* something, but not that I am doing something *intentionally*. Aquinas’ Feature easily explains practical knowledge’s *psychological immediacy*.

So Anscombe accommodates all three of practical knowledge’s apparent features, and she needs to explain none away. Given our commonsense-human methodology, this is a further benefit of her account. And because both the *Apparent Features* and the *Special Relationship* constraint are met in virtue of thinking of practical knowledge as having the two fundamental formal features it has – ‘Theophrastus’ and Aquinas’ Features – Anscombe’s account meets these two constraints in a *unified* way. This is a further benefit of her account; it means that no further explanation needs to be given as to *why* the apparent features of practical knowledge cluster together in knowledge of intentional action; why knowledge with a certain content should have the *combination* of unusual features it has.

b. The Structure-Sensitivity Constraint

Here I will explain, with reference to a discussion within *Intention* which we have not yet come across, why it appears that when an intentional action is done for reasons, it is internally structured by those reasons. This is important to our account of practical knowledge, I will suggest, because it ought to make us expect *practical knowledge* to take in the agent’s future intentional actions as well as her ongoing ones, and to be sensitive to her practical reasons. These demands will feature in a new constraint on an account of practical knowledge, the *Structure-Sensitivity* constraint. We will see how Anscombe meets this constraint in (c).

Anscombe’s ideas about the internal structure of intentional action can be gleaned from her discussion of the following ‘concrete situation’:

A man is pumping water into the cistern which supplies the drinking water of a house. Someone has found a way of systematically contaminating the source with a deadly cumulative poison whose effects are unnoticeable until they can no longer be cured. The house is regularly inhabited by a small group of party chiefs, with their immediate families, who are in control of a great state; they are engaged in exterminating the Jews and perhaps plan a world war. – The man who contaminated the source has calculated that if these people are destroyed some good men will get into power who will govern well, or even institute the Kingdom of Heaven on earth and secure a good life for all the people; and he has revealed the calculation, together with the fact about the poison, to the man who is pumping.³⁵

³⁴ For similar thoughts see Moran, “Anscombe on ‘Practical Knowledge,’” esp. §4 and Hursthouse, “Intention,” 103. I will briefly consider the relationship between practical and theoretical knowledge later on (Ch. 6, §4(e)).

³⁵ Anscombe, *Intention*, 37.

Anscombe asks whether we should think of the various descriptions of Pumping Man's activity as each picking out a separate intentional action, or as distinct descriptions under which some *one* action is intentional. She answers in favour of the latter view: he is performing one action which is internally structured (or, in the language of §1, whose point is articulated) according to the man's intentions. I will try to explain here why the latter view seems more appropriate.

Imagine we ask Pumping Man why_{RA} he is behaving as he is, and he responds by giving the various intentions with which he is acting (we ignore what backwards-looking motives or interpretative motives he might give for simplicity). We ask why_{RA} he is (A) *moving his arm up and down* and he says that he is (B) *pumping water up to the house*. Successive 'Why?_{RA}' questions give the results that he is B-ing in order to (C) *replenish the water supply*; C-ing in order to (D) *poison the household* and D-ing in order to (E) *get good men into power*. In giving us his intentions he has given us a collection of action-descriptions. The descriptions form a chain, in which each description is linked to later ones by 'in order to'.

An important observation is that not all of the descriptions in this chain relate to one another in the same way:

... there is a break in the series of answers that one may get to [the question 'Why_{RA} are you X-ing?']. Let the answer contain a further description Y, then sometimes it is correct to say not merely: the man is X-ing, but also: 'the man is Y-ing' – if, that is, nothing falsifying the statement 'He is Y-ing' can be observed. E.g. 'Why are you pumping?' – 'To replenish the water supply'. If this was the answer, then we can say 'He *is* replenishing the water supply'; unless indeed, he is not.³⁶

The descriptions before the break characterise what he *is doing* whereas after it the chain is made up of descriptions of what he *is going to do* (for simplicity we assume that his intentions will get executed), and at which his current action is aimed. For example, his claim that he is (B) *pumping water up to the house* in order to (C) *replenish the water supply* might entail that he *is replenishing the water supply*, whilst his claim that he is (C) *replenishing the water supply* in order to (E) *get good men into power* might *not* entail that he *is getting good men into power*. The break would here occur somewhere between (C) and (E).

There are two important things to note about the break. First, it is not *sharp*:

... is there much to choose between 'She is making tea' and 'She is putting the kettle on in order to make tea'? Obviously not.³⁷

In relation to Pumping Man's action, he might say equally coherently whilst moving his arm up and down that he *is replenishing the water supply* or that he *is going to* replenish the water supply, and that he is moving his arm up and down *in order to* do so. Importantly this means that *replenishing the water supply* can be represented equally as present *or* as future - within a certain boundary, what someone *is currently doing* and what they are *going*

³⁶ Ibid., 38–9. Anscombe explains that the qualification 'unless he is not' does not render her claim tautologous. The reason is that *what* we infer is not simply that *replenishing the water supply* is something he *is doing*, but further, that he is doing it *intentionally*. What we infer thus has more content than what we inferred it from.

³⁷ Ibid., 40.

to do in the future can be one and the same. The ‘break’ is vague, and descriptions within its penumbra apply to action conceived as present *and* as future.³⁸

Second, the break is not *static* (Anscombe does not say this explicitly but it is suggested by her commentary): its position moves over time along the chain of descriptions. At a certain point in time – when he’s just started doing so – we might not be happy to infer from the fact that he is moving his arm up and down *in order to* replenish the water supply to the conclusion that he *is replenishing* it. But at a later time, we *might* be happy to make this inference. At the earlier time the break comes between (A) and (C), and at the later time it comes after (C). In a slightly different example, I am (A*) *putting on the kettle* in order to (B*) *make a cup of tea* in order to (C*) *give Bev a cuppa* in order to (D*) *cheer her up*. Here you might not now be happy to infer from the fact that I am (A*) putting on the kettle in order to (D*) cheer up Bev that I am (now) cheering up Bev, but you might a bit later be happy to infer from the fact that I am (C*) giving Bev a cuppa in order to cheer her up that I am (D*) cheering her up. Here the break moves from somewhere *before* (D*) at the earlier time, to somewhere *after* (D*) at the later time. (The examples are slightly different because in the second example I *no longer* count as (A*)-ing at the later time, whereas in the first I *still* count as (A)-ing at the later time. I won’t consider further implications of these different kinds of case.) That an agent was not doing something at an earlier stage which they are doing at a later stage is not an interesting point – although it is of course true. What is interesting is that ‘the break’ Anscombe speaks of has a different position *in the same string of action-descriptions* (or: intentions), depending on how far things have progressed.

These considerations help us see why Anscombe wants to say that Pumping Man is performing just *one* action which is intentional under various descriptions, rather than many different intentional actions. Because the break is *not sharp*, there are cases in which the answer to the question *what he is doing* and the question *what he is going to do* converge on the same description. And because the break is *not static*, there are some temporal perspectives from which these questions *don’t* converge on the same description, and other temporal perspectives from which they *do*. With this in mind, the point of saying that each description picks out a *distinct intentional action* breaks down. If each action were distinct we would have to think of them as merging into and out of one another over time, and commonsense doesn’t think of things in this way. Instead it sees intentional actions as *unfolding* over time, and from a given temporal perspective, having parts which are past, present and future. This is much more consistent with the idea that there is just one action here, which is not only intentional under various descriptions, but which is conceived as present or future depending on the description under which it is considered.

Intentional action is best seen as *internally* temporally structured by the various descriptions which apply to it. Within the penumbra of the break, Pumping Man is not performing several accidentally related intentional actions at the same time, but his moving his arm up and down *is* (in the circumstances) his pumping, and so on. And the movability of the break suggests that his present and future actions are not merely accidentally related either. The fact that he is now pumping with the intention of (*e.g.*) getting good men into power means that if he succeeds in getting good men into power by pumping, he will do so intentionally.

It is by responding to ‘Why_{RA}’ in the way he does that Pumping Man makes clear the internal structure of his action. This shows that it is his *reasons* (in this case his *intentions*)

³⁸ For more discussion see Falvey, “Knowledge in Intention,” and Thompson, “Naive Action Theory.”

which are responsible for his action's internal structure: his action is structured *by* his practical reasons. By expressing the *reasons*, he articulates the point of his (*e.g.*) pumping, thus telling us both what else he *is doing* in pumping (here his intention is present-directed), and telling us what he *will be doing* (here his intention is future-directed). So in saying *why_{RA}* he is acting, he is also telling us *what* he is, and will be doing.

I can now keep my promise to explain a further dimension of Aquinas' Feature. In §1(e) I explained how in relation to an action whose point is atomic, knowing *what* one is doing under a certain description ('doing a Bob Dylan impression' in that case) renders that action intentional. But here we can add that many actions are intentional under more than one description, the descriptions and the way in which they link together being determined by the agent's practical reasons. An action's point is articulated by the agent's practical reasons, and the internal structure of this articulation is the very internal structure of intentional action itself (again, recall we are assuming no mistake in performance is made). This internal structure gives the identity of an action, and so the action's identity is determined by what the agent thinks of themselves as doing, which in turn is determined by the agent's reasons for action.

Again, then, practical knowledge is 'the cause of what it understands', but *here* – unlike in the Bob Dylan impression example – the 'what it understands' is not given by a single action-description, but by a chain of action-descriptions, tied together *as* descriptions of *the same* action by the agent's recognition of the means-end relations between them or, by the explanation the agent would give of what they are doing in terms of their practical reasons. In complex intentional actions like Pumping Man's, the present and the future are bound together, and they are bound together *by the agent's self-conception*. It is internal to intentional actions that they *can* be so structured, because their points can be articulated. The question 'Why?_{RA}' is the tool we have for uncovering this structure.

If *intentional action* can have present and future aspects which are inextricably linked, we ought to expect practical knowledge to take in its object – intentional action – under both present and future aspects. And given that this structure is determined by the agent's practical reasons, we should expect practical knowledge to be sensitive to these too. This is why Anscombe says that "the notion of 'practical knowledge' can only be understood if we first understand 'practical reasoning'."³⁹ We can now set out a further constraint on a good account of practical knowledge:⁴⁰

The Structure-Sensitivity Constraint

A good account of practical knowledge ought to see it as sensitive to the internal rational and temporal structure of intentional action.

³⁹ Anscombe, *Intention*, 57. This comment introduces Anscombe's discussion of Aristotle's account of practical reasoning in *Intention* §§33-44, but at the end of that discussion, Anscombe says: "... the interest of [Aristotle's] account is that it describes an order which is there whenever actions are done with intentions; *the same order as I arrived at* in discussing what 'the intentional action' was, when the man was pumping water." (Ibid., 80.) This is a reiteration of the point that when someone acts for reasons, her action is itself structured by those reasons. For discussion see *esp.* Anscombe, "Thought and Action in Aristotle: What Is Practical Truth?" and "Practical Inference," Mothersill, "Anscombe's Account of the Practical Syllogism" and Vogler, "Anscombe on Practical Inference."

⁴⁰ Some remarks of Rosalind Hursthouse's suggest that being able to meet the *Structure-Sensitivity Constraint* may require rejecting *CTA*. (Hursthouse, "Intention," 95.) This is intriguing but I cannot pursue it here.

If intentional action has a certain shape, we should expect this to be picked up by the faculty of practical knowledge, given that the faculty of practical knowledge takes in action *qua* intentional. We will see exactly how Anscombe's account does this in the next section, after I have filled in a final, and very important, piece of Anscombe's account.

c. Knowledge "*in Intention*"

It is tempting to read 'judgment', wherever it occurs, as 'belief', or as some mental act which precipitates a belief. But here I want to suggest that we are better off reading 'judgment' as it appears in *Intention* neutrally between 'belief' and 'intention'. The successful judgments which constitute *theoretical* knowledge are indeed to be understood as beliefs, but the judgments which set standards on successful performance – the judgments relevant to *practical* knowledge – are not beliefs but intentions. This is how I understand Anscombe's claim that practical knowledge is knowledge "*in intention*":⁴¹ an agent's practical knowledge is constituted by her intention, on the condition of successful performance. Practical and theoretical knowledge differ *at bottom* in being constituted by *distinct psychological attitudes*. Anscombe is presenting a *Non-Cognitivist* version of *Intentionalism* - NCI.

Accepting NCI gives us a unified account of why practical knowledge displays Theophrastus' and Aquinas' Features. We have already seen how these features explain practical knowledge's apparent features and its special relationship to intentional action. Below we will see that NCI very easily meets our new *Structure-Sensitivity* constraint.

On Anscombe's understanding of things, intentions are distinct from beliefs in terms of the normative relation each attitude bears to what it represents.⁴² Beliefs are attitudes whose correctness depends on how things are in the world, whereas intentions are attitudes which are themselves standards of correctness for how things (performances) are to turn out. It is easy to see how practical knowledge might involve a successful performance and not a successful judgment if the judgment involved is intention, because intention is precisely a kind of judgment to which performances are answerable, and in relation to which performances can be successful or otherwise depending on whether they match it.

Thinking of the judgment involved in practical knowledge as intention also explains why practical knowledge would display Aquinas' Feature. The object of practical knowledge is, as we know, an intentional action. This is a quick way of saying that the object of practical knowledge is action *which is intentional under some description*. The descriptions under which someone has practical knowledge of what they are doing are exactly the descriptions under which their actions are intentional. And it is the agent's intention which determines the descriptions under which her action is intentional. This means that it is the agent's *intentions* which formally characterise the object of practical knowledge. If the agent's intentions are also the judgments involved in practical knowledge, then it is easy to see how her knowledge will be the (formal) cause of what it understands: the agent's intention determines the description under which an action must fall if it is to be intentional. And if an action *does* fall under this description, because there is no mistake in performance, the agent has knowledge (Theophrastus again).

⁴¹ Anscombe, *Intention*, 57.

⁴² The point is usually made by saying that belief and intention have different 'directions of fit', and I have used this language in previous Chapters. But Anscombe doesn't use this terminology. There are reasons to worry about the 'direction of fit' terminology; see Ch. 6, n. 1.

It is for related reasons that *NCI* happily meets our new *Structure-Sensitivity* constraint. Our discussion of Pumping Man showed that his action was internally temporally and rationally structured by his intentions, which he gave in his answers to our successive ‘Why?_{RA}’ questions. We can ask a similar question about his intentions to the question we asked about his intentional actions: are there as many intentions as there are descriptions given in his responses to ‘Why?_{RA}’? Or is there just one intention, with a complex content; an intention *to do A in order to B, in order to C, in order to D* (or: *to E by D-ing, by C-ing, by B-ing, by A-ing*)? Here is what Anscombe says:

... when we speak of [*n*] intentions, we are speaking of the character of being intentional that belongs to the act in each of the [*n*] descriptions; but when we speak of one intention, we are speaking of [...] the intention *with* which the act in each of its other descriptions was done, and this intention so to speak swallows up all the preceding intentions *with* which earlier members of the series were done. The mark of this ‘swallowing up’ is that it is not wrong to give D as the answer to the question [‘Why?_{RA}’] about A; A’s being done with B as intention does not mean that D is only indirectly the intention of A...⁴³

Just as we can talk of one action which is internally structured, we can talk equally of one internally complex intention: the last intention in the chain ‘swallows up’ – and so *contains* – the preceding ones in the chain. So we can think of his action as being the intentional internally complex action it is in virtue of being the execution of his intention which is itself internally complex; the internally complex intention is manifested in an isomorphically internally complex intentional action.

The *Structure-Sensitivity* constraint demanded an account of practical knowledge to be sensitive to the internal rational and temporal structure of intentional action. And if practical knowledge is constituted by *intention* (if *intentions* are the judgments involved in practical knowledge), and if intentions are themselves internally structured isomorphically to the action which is their execution, no further questions remain about how practical knowledge manages to be sensitive to intentional action’s internal structure. Indeed if practical knowledge is constituted by intention, understood in this way, it *could not but* be sensitive to intentional action’s internal structure because it *determines* this structure (again: it is the cause of what it understands – only action which comes under the descriptions internal to the intention will *be* an intentional action; will *be* a successful performance and so will *be* an object of practical knowledge).

I will not argue that a conception of practical knowledge as belief-constituted (either *CI* or a non-*Intentionalist* account like O’Brien’s or Paul’s) *could not* meet the explanatory challenge set by the *Structure-Sensitivity* constraint. But a view of practical knowledge as belief-constituted *would* have to explain why someone’s beliefs about what they *are* doing would tend to go along with beliefs about what they were *going* to do, and to explain the interconnectedness of these beliefs with agent’s practical reasons (if any). It’s not that such an explanation must be hard to come by, it’s just that on *NCI* no such explanation even needs to be *given* – on the contrary, practical knowledge *could not but* take in these various kinds of object, if it is constituted by intention, understood non-Cognitively and as itself internally structured.

Before summing up this section it is worth making a clarification. Some have interpreted Anscombe as thinking of practical knowledge as knowledge that one has of *what* one is

⁴³ Anscombe, *Intention*, 46.

doing in virtue of knowing *why* one is doing it. But if *NCI* is Anscombe's view, this is not quite right. *NCI* is a view on which it is in virtue of *having intentions* (which may or may not be further articulated by one's practical reasons) that one has practical knowledge of one's action (assuming no mistake in performance).

The view of practical knowledge as knowledge *of* one's reasons would make it impossible to have practical knowledge of those intentional actions we do for no reason; whose point I called 'atomic'. But on the view I am attributing to Anscombe, we *can* and *do* have practical knowledge of these intentional actions. I have practical knowledge that I am doing my Bob Dylan impression just in case I am doing it in execution of my intention. It is because I intend to do it that this action has a point. And it is because it has a point that it can constitute a successful performance. It is successful *just* insofar as it is the execution of my intention, and if it is, I have practical knowledge.

We are encouraged to reject – as an interpretation of Anscombe – the view on which practical knowledge is (*inter alia*) knowledge *of* reasons (including intentions) by Anscombe's commitment to the view that knowledge presupposes the possibility of brute intelligible error. I fail in theoretical knowledge if my belief is mistaken; I fail in practical knowledge if my *performance* is mistaken, and intelligible examples of both kinds of case abound. But being wrong about *why* I am acting tends to suggest self-deception, confusion or some other assault to my intelligibility. So we ought not interpret Anscombe as thinking that we know what we are doing (will do) by *knowing why* we are doing (will do) it. Rather, we know *what* we are doing and will do in virtue of intending to do it. When these intentions are executed (whether internally structured or atomic) – when there is no mistake in performance – they constitute practical knowledge. This is *NCI*'s conception of practical knowledge.

d. Summing up

Anscombe's conception of practical knowledge met the *Apparent Features* and *Special Relationship* constraints in a unified way, in virtue of displaying two fundamental formal features which distinguished it from theoretical knowledge: 'Theophrastus' and Aquinas' Features. Thinking of practical knowledge as constituted by intention, understood non-cognitively, explains why it should display both features, and leaves no explanatory work to be done in order to meet our new *Structure-Sensitivity* constraint. And it does all *this* in a unified way. So *NCI* has significant explanatory power. Additionally, on *NCI*, all three of practical knowledge's *apparent* features are accommodated as *actual* features. And given that our guiding principle throughout is to do justice, where possible, to our commonsense-human understanding of things, this is a further significant benefit of *NCI*.

Let's zoom out a bit and think about *NCI* in comparison to the other accounts we have been considering. These all conceived practical knowledge as constituted by belief, and in so doing they put themselves under a certain explanatory demand: for a belief to constitute *knowledge* it must not only be true, but must be held in favourable epistemic circumstances, as demanded by the Standard Metaphysics of Knowledge. What are these circumstances? O'Brien, Paul, Velleman and Setiya all accepted the need to respond to this question, and responded to it in different ways. But none of their answers were entirely adequate.

Anscombe side-steps this explanatory demand. According to *NCI*, practical knowledge is not epistemically grounded in *any* way. Recall Velleman's worry about such a suggestion, that it makes "epistemological exceptions" in practical knowledge's favour by denying that it is subject to the usual epistemological requirements on knowing.⁴⁴ Velleman interprets Anscombe as simply *not being able to see* what evidence practical knowledge might be based on.⁴⁵ But this is a mistake. Anscombe thinks of practical knowledge as a kind of knowledge formally such as to *rule out* an evidential – or any other epistemic – basis. Knowledge which displays 'Theophrastus' and Aquinas' Features not only *need not* but *could not* be grounded in any kind of epistemic justification.

Anscombe gives us a way of thinking about practical knowledge on which its epistemic ungroundedness (as well as its *FPA* and *psychological immediacy*) are not puzzles to be solved, but on which they are necessary concomitants of the fact that practical knowledge is fundamentally formally distinct from theoretical knowledge. Underwriting these differences – I have suggested – is a difference in the attitude constituting practical and theoretical knowledge. Whereas theoretical knowledge is constituted by belief, practical knowledge is constituted by intention.

But *NCI* is not without its challenges. I will consider three in the next section.

3. Three *Prima Facie* Problems

The first challenge takes up some worries we have already come across when discussing Lucy O'Brien's view of practical knowledge, and I mention them here mostly to put them to one side. The second two worries are more serious. I will not seek to solve them here – that will be the work of Chapter Six, where I will defend a view of practical knowledge very closely related to *NCI*, which I think can meet all of our constraints.

a. *Too Special a Relationship?*

I said that Anscombe did well on the *Special Relationship* constraint because of her view that practical knowledge displays Aquinas' Feature. Thinking of practical knowledge as 'the cause of what it understands' seems to entail a commitment to the idea that an action will be an object of practical knowledge under *all* descriptions under which it is intentional. But many have thought – as we have already seen (Ch.1, §3(b); Ch. 2, §1(c)), that this claim is too strong; it seems to admit of counterexamples.

As we saw in Chapter One, the quintessential version of this kind of example is Davidson's example of the man making ten carbon copies, but I am going to concentrate on O'Brien's example of Bev playing snooker instead. First, because we have already discussed O'Brien's example in some detail, but second because I think that O'Brien's case is actually harder for a view like *NCI* to cope with than Davidson's, for there is a kind of response to Davidson's case which is not available in relation to O'Brien's. The relevant difference is that *making ten carbon copies* is something you can still be doing even if your first attempt doesn't work out. You'd simply take the successful sheets off the top and carry on. The response is that if this counterfactual is true of the carbon copier then he *does know* that he is making ten carbon copies – even if he's surprised that *his first*

⁴⁴ Velleman, *Practical Reflection*, 25.

⁴⁵ "... the evidence [on which practical knowledge is based] has escaped Anscombe's attention because the agent's conception of his action isn't and could not have been derived from it." (Ibid., 105.)

attempt worked (perhaps he is not intentionally *making ten carbon copies in one go* but this needn't mean he is not intentionally *making ten carbon copies*). And if the counterfactual is false, he's not making the ten copies intentionally. Michael Thompson makes roughly this argument.⁴⁶ But potting the (*this*) pink is an action one only ever gets one go at, which means that this kind of counterfactual is not relevant in O'Brien's example, so this kind of response cannot be given.

O'Brien thought Bev could intentionally pot the pink but fail to be authoritative about the fact she was doing so and so lack practical knowledge that she was. I complained there that O'Brien wrongly generalises from this example in at least two ways (§1(d)). But Anscombe's critics do not need an example which will generalise; they just need one case in which an agent acts intentionally under some description, but lacks practical knowledge under that description.

My thought is that when we consider the case in detail, it becomes less obvious that Bev really does both intentionally pot the pink, *and* lack practical knowledge of what she is doing. And indeed for any case, the extent to which we are suspicious about whether the agent has practical knowledge that they are doing something under some description, roughly equals the extent to which we are suspicious about whether they are acting intentionally under that description, at least when we think it through.

Compare again Bev's pink-potting with Ronnie's. Imagine again that on point of the physics of the table, things are just the same in Bev's and in Ronnie's shots. And recall that Bev doesn't know, as she strikes the cue-ball, whether she is potting the pink, whereas (*ex hypothesi*) Ronnie does. Bev doesn't know this *because* she doesn't know the pink will go in, whereas Ronnie knows the pink will go in *because* he knows he's potting it. This was what I suggested in Chapter Two (§1(d)).

These differences are underwritten by the fact that for Bev, luck plays a rather large role, and expertise and skill a very minor one, whereas for Ronnie it is the other way round. Given that Ronnie is skilled and Bev is not, this is not a contingent feature of the case. Being skilled in some domain *just is* narrowing down the role which luck is likely to play in one's successes in that domain, and being unskilled *just is* having one's successes in that domain owe a large debt to luck. For these reasons I think we might be differentially happy to say of Bev and of Ronnie, that they pot the pink *intentionally*, for the concept of *intentional* action is one which to a certain extent squeezes out that of *luckily* achieving one's aims.

I can aim to win the lottery – I do so when I buy a ticket. So why can't I win the lottery *intentionally*? Because there is no way of achieving this aim which does not owe a very large debt to luck. What about winning at Chess? I can aim at this too, but can I do this intentionally? Well it depends on the case. If I am playing an averagely intelligent eight-year old child, I can win intentionally because winning *here* is something I can do without relying very much on luck. But if I am playing Deep Blue, winning is almost impossible for me to *do*, but it is completely impossible for me to do without relying much on luck. And it is completely impossible for me to do *intentionally*. Can I cycle into town, make a cup of tea, write a thank-you letter or finish a chapter of my Thesis intentionally? Typically, yes. It needn't be *easy* to do all these things (which is why I include the final example), but being able to do them *without relying on luck* doesn't require being able to do

⁴⁶ Thompson, "Anscombe's Intention and Practical Knowledge," 209–210.

them at the drop of a hat. Again, doing something intentionally seems to require meeting one's aims without relying very much on luck.

Bev's potting the pink is not like my beating Deep Blue at chess. But nor is it like my beating the eight-year-old. There is a *scale* running from those aims we achieve purely by luck to those aims we achieve independently of luck playing any role at all (Given that we are necessarily hostage to worldly impingements, we might expect the terminal points of this scale to be unoccupied by any *actual* action.) And the closer to the 'pure luck' end of this scale the achievement of an aim sits, the less happy we would be to say that it represents the *execution* of an intention; that the aim was achieved *intentionally*.

Further, I think that the degree to which we are happy to say that someone's *knowledge* of what they are doing explains their knowledge that they *will achieve* their aim in acting, maps on to the position of the achievement of that aim on our scale. More concretely, because we are happy to say that Ronnie knows that the pink is going in *because* he knows that he is potting it, we think of his potting the pink as rather close to the 'no-luck' end of the scale. Whereas because Bev needed to *wait and see* whether the pink was going in, in order to know that she *was potting* it, we are happier to say that luck played a relatively central role in her potting it, and so – I claim – less happy to say that she *intentionally* potted it.

I don't want to say that there are *no* reasons to say that Bev potted the pink intentionally. If I said that it would leave it mysterious why examples like O'Brien's and Davidson's have seemed so compelling. Instead the thought is that there are *some* reasons to deny that Bev's pink-potting is intentional which do not exist in Ronnie's case, *viz.* that luck plays a relatively large role in Bev's pink-potting but not in Ronnie's. And for this reason, Bev's pink-potting is not *straightforwardly* an example in which someone does something which is intentional under some description but is not an object of practical knowledge under that description. I will leave the case to one side for now, but will have more to say about it in Chapter Six (§2(b)), for my discussion there will make available a slightly more nuanced understanding of the special relationship between practical knowledge and intentional action.

b. Two big (prima facie) problems: Our three remaining Constraints

The second and third problems cannot be put aside so easily. I have not considered in the foregoing how *NCI* fares on two of our constraints – the *Why Knowledge?* constraint and the *Intention-Content* constraint. And the worries are that it will not meet the former, and that meeting the latter leaves it mysterious why practical knowledge should be – as it is – a kind of *propositional* knowledge. The second problem will lead me to add a final constraint to our set.

It is important to *NCI* that practical and theoretical knowledge are formally distinct, since it is their formal differences which explain practical knowledge's philosophically interesting features. The features turn out, on Anscombe's view, to be puzzling only given what she calls our "incurably contemplative conception of knowledge".⁴⁷ Practical knowledge is not puzzling as long as we free ourselves from the idea that *all* knowledge is constituted by belief, requires a successful judgment, and is 'derived from the objects known'. In the language I have been using throughout this Thesis, what Anscombe

⁴⁷ Anscombe, *Intention*, 57.

thinks we need to do is to free ourselves from the assumption that *all* knowledge is structured by the Standard Metaphysics of Knowledge.

But the “contemplative conception of knowledge” - the view that all knowledge fits the Standard Metaphysics - needn’t be held on the basis of prejudice. It is likely to simply seem hard to understand how a kind of *knowledge* could possibly be as Anscombe describes it. First, how could knowledge *fail* to be constituted by a true belief held in epistemically favourable circumstances – how could something which fails to fit the Standard Metaphysics of Knowledge be a kind of *knowledge* in anything but name? And second, how could knowledge be constituted by intention just on the condition that that there is no failure of performance? Isn’t Anscombe, in holding this, equating *knowing that one is doing something* with *intentionally doing it*? But intentionally doing something is not an epistemic phenomenon – it is an action-theoretic one. So don’t we *need* to hold on to the idea that whatever practical knowledge looks like, it *must* display the Standard Metaphysics? Isn’t that the only way to understand *any* phenomenon as a *knowledge*-phenomenon?

If we can’t explain how executing an intention could possibly be a kind of *knowledge* then we will have to reject *NCI* as failing to meet the *Why Knowledge?* constraint. And note that its failure will be much more serious than the failures we have seen so far in relation to this constraint. For our other protagonists (Paul, Velleman and maybe Setiya) failed to meet it because they failed – in their own ways – to give a complete description of the favourable epistemic conditions under which a true belief about what one is doing gets elevated to knowledge-status. But they were nonetheless all working within a paradigm - that of the Standard Metaphysics of Knowledge - in which it is clear how, *pending a filling-in of the favourable epistemic condition*, the structure they are describing would be the structure of a kind of *knowledge*. But Anscombe’s ‘*Why Knowledge?*’ problem is of a different order. The problem for *NCI* is not in the details, but in the whole idea that *knowledge* might be constituted in a belief-free way, simply by the execution of intention.

The second problem also relates to the idea that practical knowledge is intention-constituted, and concerns how to understand intentions. I have not explicitly attributed to Anscombe the view that intentions are *not* propositional attitudes, and nothing in her account commits her to Propositionalism. But if *NCI* is to meet the *Intention-Content* constraint, it must conceive of intentions as at least in their fundamental form, non-propositional.⁴⁸ The worry is that meeting the *Intention-Content* constraints seems *prima facie* inconsistent with meeting a final constraint, which we can now state:

The Knowledge-Content Constraint

An adequate account of practical knowledge ought not to leave it mysterious why practical knowledge is a kind of propositional knowledge.

How could a *non-propositional attitude* possibly constitute a kind of *propositional knowledge*?⁴⁹ This is a worry for Anscombe’s view if we build into it a rejection of Propositionalism. But it is also a worry for *any* version of *Intentionalism*. For an *Intentionalism* which accepts

⁴⁸ Anscombe certainly thinks that intentions can be *expressed* using indicative statements, but as I pointed out at the end of Chapter Four (§3(b)), the content of an attitude does not in general match the content of its linguistic expression.

⁴⁹ McDowell thought that it couldn’t, recall Ch. 4, n. 3.

Propositionalism will flout the *Intention-Content* constraint, whereas an *Intentionalism* which rejects Propositionalism looks set to flout the *Knowledge-Content* constraint.

An account which flouts any of these three constraints; the *Why Knowledge?* constraint, the *Intention-Content* constraint or the *Knowledge-Content* constraint must be rejected. In the next Chapter I will introduce and defend a stronger version of *NCI* than the one here described, one which explicitly rejects Propositionalism about intention. And I will develop it in a way which – I think – meets these three constraints.

Chapter Six:

Non-Propositionalist Intentionalism

Introduction

This Chapter clarifies and defends *Non-Propositionalist Intentionalism* – *NPI*. *NPI* agrees with Anscombe’s *NCI* in thinking of practical knowledge as constituted by intention, as long as there is no mistake in performance or, as I shall put it here, on the condition that the intention is executed (understood in the way described in Ch. 4, §1(b)). According to both views, I have present-directed practical knowledge that I *am typing* because my intention to do so is getting executed, and if I have future-directed practical knowledge that (e.g.) I *will go to the cinema tomorrow*, it is because I now having an intention to go, which will get executed when the time comes.

The difference between *NPI* and *NCI* is just in the way in which the views understand intentions. Anscombe was explicit only that intentions are unlike beliefs in relation to the normative relationship they bear to ‘the facts’: intentions set standards on ‘the facts’ (strictly, on *performances*) whereas beliefs have standards set for them by how things are with the facts.¹ The idea of intention internal to *NPI* is of a kind of attitude which differs from belief along this normative dimension, but also along a *semantic* dimension: intentions are, at least in their most fundamental form, non-propositional attitudes; they are attitudes towards *acting* in a certain way, and their contents make no reference to their subjects.²

¹ It is worth explaining briefly why I avoid the language of ‘directions of fit’ here. There are two reasons why this terminology is unhelpful in this context. First, the idea of two attitudes having different ‘directions of fit’ suggests that both attitudes *might* have the same content, differing *only* in their direction of fit. But belief and intention could sit in such a relation only if intentions were propositional attitudes, and I have rejected this view. Second, the ‘direction of fit’ terminology seems to suggest that the normative structures of intention and belief *mirror* one another, but they do not, not perfectly at any rate. With belief, there is a sharp distinction between whether blame for falsity is on the side of the world or the mind; it is on the side of the mind. With intention, there is no such sharp distinction – if an intention is not executed, then the fault is with the performance; something at once mental *and* worldly. The ‘direction of fit’ language papers over important structural *differences* between the ways in which success and failure relate to intention and to belief. Similar worries are brought up in Alvarez, *Kinds of Reasons*, 66–72 and in Boyle and Lavin, “Goodness and Desire,” 170–174. Both Alvarez and Boyle & Lavin link the notion of direction of fit with *Propositionalism about desire*; Alvarez discusses additional worries.

² I made the assumption there that some intentions were propositional in order to start my argument from a position the Propositionalist would accept. But it is consistent with my *conclusion* in Ch. 4 that *no* intentions are propositional attitudes; that all intentions are intentions *to do* something (*to send* Bev to prison, for instance, rather than *for* Bev *to go* to prison). These issues have implications for the scope of the discussion of this Chapter: if intentions are *never* propositional then the account I give here is an account of *all* practical – *i.e.* intention-constituted – knowledge. If some intentions are propositional,

Denying *Propositionalism* about Intention entails denying *Cognitivism* about Intention, which means that *NPI* includes or entails *NCI*. And this means that any benefit of *NCI* is *eo ipso* a benefit of *NPI*. The benefits of *NCI* were argued in Chapter Five to be that it meets the *Apparent Features* and *Special Relationship* constraints in a unified way, and that it met the new *Structure-Sensitivity* constraint with minimal effort. The first two constraints were met by thinking of practical knowledge as displaying Theophrastus' and Aquinas' Features (Ch.5; §2(a)). These features in turn were explained as stemming from thinking of practical knowledge as constituted by intention, understood *non-Cognitively*. And thinking of practical knowledge as intention-constituted was also responsible for *NCI* meeting the *Structure-Sensitivity* constraint particularly easily (see again Ch. 5, §2(c)). Because *NPI* entails *NCI* we can treat the question of whether *NPI* meets these three constraints as having already been answered.

There were two major problems with *NCI* (Ch.5; §3(b)). First, it was unclear whether it could meet the *Intention-Content* constraint consistently with a new – and final – constraint, the *Knowledge-Content* constraint; it was unclear how *if* we think of intentions as non-propositional, then we can think of the knowledge they constitute as a kind of propositional knowledge. Because *NPI* contains a rejection of Propositionalism about Intention, *NPI* will meet the *Intention-Content* constraint. *NPI* meets this constraint trivially, but not in an *ad hoc* way, since a rejection of Propositionalism was motivated by my argument in Chapter Four. With this in mind, the problem for *NPI* is whether it can meet the *Knowledge-Content* constraint. I will refer to this as the *Knowledge-Content* problem.

The second problem for *NCI* was to explain how it meets the *Why Knowledge?* constraint. Anscombe seemed to think it dogmatic to *deny* that someone executing their intention *thereby* knows that she is acting as intended,³ but her opponent is bound to find it equally dogmatic simply to insist that they *do*. Acting as one intends is acting intentionally. Why think that *this* would be a *knowledge*-phenomenon as well as an action-theoretic one? According to *NCI*, practical knowledge is not structured according to the Standard Metaphysics of Knowledge. Why should its metaphysical structure – that of executed intention – be thought of as a form which *knowledge* can take? I will refer to this as the *Why Knowledge?* problem, and this problem is inherited by *NPI*.

The primary aim of this Chapter is to respond to the *Knowledge-Content* and the *Why Knowledge?* problems, on behalf of *NPI*. My strategy for doing so will involve developing two ideas in addition to the core claim of *NPI*. The first idea is about how to think of *knowledge per se*. *NPI* requires thinking of knowledge in a way that is neutral between its practical and theoretical manifestations, because on *NPI* – as on *NCI* – practical and theoretical knowledge are *formally* distinct. I will characterise knowing that something is the case in terms of being *mentally engaged* with a fact, in such a way that one has certain specifically epistemic *capacities* relative to this fact. Knowing that something is the case fundamentally involves the capacity to express, and to make (practical and theoretical) inferences from what one knows. My response to the *Why Knowledge?* problem will involve explaining how intending – just as much as believing – can put a person in

then the account is of an importantly *privileged* subset of practical knowledge, given the fundamentality of *non-propositional* intentions explained at the end of Chapter Four. In the latter case, I would of course need to say more about what practical knowledge looks like when constituted by *propositional* intentions, and how this knowledge relates to knowledge constituted by non-propositional intentions.

But this task goes beyond the scope of what I attempt here.

³ Anscombe, *Intention*, 57.

possession of these capacities: it *will* do so just in case the execution condition on practical knowledge is met.

The second idea is that *practical knowledge*, *intentional action* and the *execution of intention* are to be understood as *non-reductively* related to one another. I will describe these phenomena as three different facets of a single underlying phenomenon. These three facets are conceptually related, and the underlying phenomenon of which they are facets has them *as* facets essentially. The concepts which pick out these three facets must be conceptually elucidated in terms of one another. This idea will help me respond to the *Knowledge-Content* problem.

I have said that *NPI* entails *NCI* but I would like to make a couple more comments about the relationship between the two views before continuing. First, it might be that Anscombe is best thought of in *Intention* as putting forward *NPI* and not the weaker *NCI* which I have attributed to her. I won't take a stand on this here. I attributed *NCI* to Anscombe because she is *not* explicit that intentions are *non*-propositional attitudes, not because she *is* explicit that they *are* propositional attitudes. I suspect that Anscombe would accept the stronger *NPI*, if my interpretation of her is half-right, but I don't seek to argue for this here, and to some extent *NPI* was developed independently of interpreting Anscombe in that way that I did in the previous Chapter.⁴ Second, although *NPI* takes over *NCI*, my discussion of *NPI* will abstract away from certain of the issues which we concentrated our attention on in Chapter Five. I will not, for example, consider the details of the internal rational and temporal structure of intentional action and of practical knowledge.

The plan for this Chapter is as follows. In §1 I develop *NPI* by explaining the ways in which, on the view, practical and theoretical knowledge are analogous and disanalogous. This section contains the beginnings of a response to the *Knowledge-Content* problem (in (d) & (e)), and it does the ground-work for my response to the *Why Knowledge?* problem, which I begin to give in §2. §2 clarifies what is at issue in the *Why Knowledge?* Problem; I reject two versions of the problem by being clear about the relationship *NPI* sees as holding between intentional action, intention-execution and practical knowledge. This is where the 'three-facets' idea described above comes in. In §3 I respond to the clarified *Why Knowledge?* problem. I list four capacities which commonsense epistemology links with knowing, and explain why someone who is executing (or will execute) an intention displays each of these capacities. We will then have *pro tanto* reason for accepting that *NPI* describes a kind of *knowledge*.

In §4 I consider some reasons to think that knowing entails believing, and consider how harmful this conclusion would be to *NPI*. I argue that knowing requires being *mentally engaged* with a fact in such a way as to underwrite the epistemic capacities considered in §3, and I argue that although believing *can* embody the relevant kind of mental engagement, intention can too. I will then have removed a final obstacle to accepting that *NPI* describes a kind of knowledge. I will also come back to a new version of the *Knowledge-Content* problem, defending *NPI* against this as well. I will conclude that *NPI* is a view of practical knowledge which meets all six of our constraints, and is worth taking seriously.

⁴ That is not to say that the ideas in this Chapter are independent of my understanding of *Intention*; it's rather that the particular way in which I have interpreted Anscombe in the previous Chapter also owes a debt to the development of the view I describe and defend in this one.

1. Non-Propositionalist Intentionalism

NPI can be clarified by describing the analogies and disanalogies which, on the view, exist between practical and theoretical knowledge. So here I will state and explain these analogies and disanalogies, without aiming to argue for *NPI*. During the course of this section we will make a start on responding to the *Knowledge-Content* problem (in (d) & (e)), but I will have nothing to say about the *Why Knowledge?* problem. For this reason the problem is likely to seem acute throughout this section. But I ask the reader not to worry too much about this here; I will address these concerns directly in later sections.

a. 'Knowledge-Elevation' in Practical and Theoretical Knowledge

According to *NPI*, intention and belief are different kinds of psychological state, but are both equally potentially knowledge-constituting. Only *potentially* knowledge-constituting, because both belief and intention constitute knowledge only if further conditions are met; if they are not met a given intention or belief falls short of knowledge. I will refer to whatever meets these further conditions as 'knowledge-elevators'. Knowledge-elevators are factors which elevate a potentially knowledge-constituting attitude to actual knowledge-status. Different factors elevate a *belief* to *theoretical* knowledge, and an *intention* to *practical* knowledge: theoretical and practical knowledge have different knowledge-elevators.

Theoretical knowledge's knowledge-elevators are familiar from traditional epistemology as *truth* and *favourable epistemic circumstances*. A belief will be elevated to knowledge-status when it is characterised by these factors. People argue about the exact nature of the *epistemically favourable circumstances* which elevate a (true) belief to knowledge-status - I will come back to this idea in (c).

Practical knowledge's knowledge-elevator is – for us – familiar from the previous chapter. Anscombe held that intentions constituted knowledge as long as there was no *mistake in performance*, (Theophrastus' Feature) which means that intentions are elevated to knowledge-status by their *execution*: *execution* is intention's knowledge-elevator.

Isn't there a problem for *NPI* here? If an intention is elevated to knowledge-status by its *execution*, then how can we ever have intention-constituted knowledge of *future* actions? How can I know that I *will go* to the cinema tomorrow if my knowledge requires *execution* and there is no *execution* yet on the scene?

But *NPI* does not say that intentions constitute knowledge only *during* their execution. It says that intentions constitute knowledge *given* their execution. So I can know that I will go to the cinema tomorrow, in virtue of now intending to go, as long as I execute my intention *tomorrow* – when the time comes. If we don't think it is problematic for me to now know that the film I want to see will be on tomorrow, virtue of (*inter alia*) a now future condition (the film's being on tomorrow) then we shouldn't find a problem with my knowing that I *will go* in virtue of a now future condition (my executing my intention tomorrow). *NPI* involves the assumption that someone can *now* have knowledge of the future in virtue of relating to some *future* condition, but this assumption seems *no more* problematic in relation to practical than in relation to theoretical knowledge.

But what if I now intend to go to the cinema tomorrow afternoon, but change my mind this evening, only to change it back again tomorrow morning, ending up going to the cinema as originally planned tomorrow afternoon? It doesn't seem that in such a case I *know* now that I will go to the cinema tomorrow afternoon. Is this a problem for *NPI*?

No. *NPI* does not deliver the verdict that in this case I know (now) I will go to the cinema tomorrow afternoon. For although we *can* say that I have one intention which I drop this evening and take up again tomorrow morning, intentions ought to be identified not only by reference to their subjects and their contents, but also in a way which respects the fact that they exist for the period of time between being formed and being either executed or dropped (because of a change of mind, laziness or forgetfulness). With this in mind, it turns out that the intention I execute tomorrow afternoon is distinct from the one I have now. The one I have now does not get executed, so it is not elevated to knowledge-status. So *NPI* agrees with common-sense that I *don't* now know that I will go to the cinema tomorrow afternoon.

I will consider some more complexities which attach to future-directed practical knowledge later on (§3(c)), but now let's consider some analogies and disanalogies between practical and theoretical knowledge's knowledge-elevators.

b. Knowledge-Elevation and two kinds of mistake

I have just suggested that theoretical knowledge has *two* knowledge-elevators (truth and favourable epistemic circumstances), whereas practical knowledge has *one* (execution). So here we find a *disanalogy* between practical and theoretical knowledge on *NPI*. But there is nevertheless an analogy between practical and theoretical knowledge's knowledge-elevators. For both kinds of knowledge can be undermined in *two* different ways. In theoretical knowledge one of the two knowledge-elevators might fail to obtain, whereas in practical knowledge the knowledge-elevator can fail to obtain in two different ways.

Consider first how *theoretical* knowledge might be undermined. Bev believes that Pam's new handbag is a knock-off. Her belief might fail to constitute knowledge by being *false* – if Pam's new handbag is not a knock-off but is legit. But even if it's true, it might fail to constitute knowledge if the epistemic circumstances are not right; if Bev's belief is, e.g. based on the false assumption that Pam bought it down the market and not from a reputable stockist.

Practical knowledge will be undermined if a person's intention is not executed, and intentions can fail to be executed in two ways. My intention to water the tomatoes might not be executed because I don't do what I intend – I forget to water them, or change my mind, or water the aubergines by mistake. But I also might fail to execute my intention even though I *do*, in a sense, do what I intend. Recall from Chapter Four (§1(b)) that in the case where I trip over the step, sloshing water over the grow-bags, my intention is *satisfied* but not *executed*. I fail to have practical knowledge in both cases. (I may or may not have *theoretical* knowledge in the second case – if I do I will have a true belief that I am watering the tomatoes which will most likely be perceptually justified. Here *perception* constitutes the epistemically favourable circumstances required for theoretical knowledge.)

The two ways in which the execution condition might fail to be met map on to *two strands* of the way in which intentions set the standard for successful performance. One strand is

determined by the *content* of the intention: if the content is not satisfied then I don't execute my intention. The other strand is determined by the fact of its being an *intention* that I have, and the fact that intentions are attitudes whose execution requires more than a *mere* match between content and outcome. Execution is a *special kind* of success, and it is success of this kind which we aim at in intending. We don't merely aim to get our intentions' contents satisfied. We aim to *execute* them; only in meeting *this* aim do we act intentionally.

The two kinds of potential mistake can be described neutrally between the practical and the theoretical cases. One kind of mistake occurs when a person lacks knowledge because she has a potentially knowledge-constituting attitude whose content fails to relate her to the facts (false belief; unexecuted intention). The other kind occurs *despite* the fact that the content of a person's potentially knowledge-constituting attitude *does* relate her to the facts, and *because* it relates her in the wrong way for knowledge (true but unjustified belief; satisfied but unexecuted intention).

So despite an initial disanalogy between practical and theoretical knowledge – practical knowledge has one knowledge-elevator (*execution*), whereas theoretical knowledge has two (*truth* and *favourable epistemic circumstances*) – the cases are nonetheless analogous in terms of the complexity of the kind of success knowledge involves in each case. This complexity is shown up in the fact that there are two ways in which both practical and theoretical knowledge can be undermined.

c. Arguing about Knowledge-Elevation

I mentioned above that people argue about one of theoretical knowledge's knowledge-elevators, *viz.* how to understand the favourable epistemic circumstances in which a true belief constitutes knowledge. I will explain what these arguments look like, and describe some analogous arguments surrounding *execution* on the practical side. Traditionally, these latter arguments have been conceived as arguments about the conditions on intentional action, and they are. But *NPI* conceives them as *also* being arguments about the conditions on practical knowledge.

First, there is an argument about *reduction*. Some think that the epistemically favourable circumstances in which a true belief constitutes knowledge can be given a universal reductive description, and some deny this. A universal reductive description is a description of what makes for *favourable epistemic circumstances* which a) applies to *all* and *only* cases in which a true belief constitutes knowledge and which b) does not include (implicitly or explicitly) the concept of knowledge itself. This is one way of understanding the dispute between what has come to be known as the 'knowledge-first' approach to understanding knowledge and the more traditional 'JTB'-style accounts.

Second, there are arguments which come up most naturally within the reductive paradigm. These are arguments about exactly what constitute the epistemically favourable circumstances which elevate a true belief to knowledge-status. Must such beliefs be internalistically justified, and if so what does this require? Or do true beliefs constitute knowledge when they track the truth, or when they are caused by the facts they concern, or when they are produced by reliable mechanisms? Ought we be trying to give a reductive account of *justification*? Or does the *justification* condition need replacing with other kinds of condition which stand a better chance of being universally and reductively

described? The arguments between these positions are familiar in the post-Gettier literature, and I won't detail them here.

Both arguments – the one about reduction and the one internal to the reductionist position – have analogies in the practical sphere. Familiar from the action-theoretic literature is a debate between those who think that a universal reductive description of the requirements on *intention-execution* can be given, and those who don't. The most familiar version of the reductive account of intention-execution is a reductive version of *CTA*, which says that an intention is executed *iff* it suitably causes some bodily movement which matches the content of the intention, where *suitability* is assumed to be describable without falling back on the concepts of intention-execution or intentional action.

Within this reductionist position, arguments abound about the exact terms in which the reduction is to be effected; about what features must be displayed by a causal chain between intention and bodily movement matching the intention's content, if the intention is to confer 'intentionalness' on to the bodily movement; for it to count as *executed* and not merely *satisfied*. Again, these debates are familiar, this time from the literature on the problem of causal deviance in the context of action-theory, and again they needn't be detailed here.

Both debates in the practical sphere are about how to understand *execution* (as distinct from *satisfaction*). Because intentional action is the upshot of executing our intentions, the debates are at the same time about how to understand *intentional action*. But according to *NPI* both debates are *also* about how to understand *practical knowledge*, because it is *execution* which elevates an intention to knowledge-status. Thought of in this way the debates are about whether *practical knowledge* can be given a universal non-reductive description and if it can, in what terms. And so thought of in this way, both practical and theoretical knowledge admit of the same kind of arguments. Later on (see §2(b); §4(d)) I will come down on the *non-reductive* side of the debate as it attaches to practical knowledge; this will be important to my response to the 'Why Knowledge?' problem.

d. Factivity and propositional structure in Practical and Theoretical Knowledge

The reader might be worrying about the idea that practical knowledge has only one knowledge-elevator – *execution*. The worry I have in mind is that because knowledge is *factive*, *truth* must be a condition on any attitude's constituting knowledge. But I have said that according to *NPI*, *truth* is *not* amongst practical knowledge's knowledge-elevators.

This worry relates to the *Knowledge-Content* problem, which was the problem of explaining how a *non-propositional* attitude (intention) could possibly constitute a kind of *propositional knowledge* (practical knowledge). The problems are related because it is *propositional* knowledge which is *factive*, yet *NPI* holds that practical knowledge is *both* constituted by a *non-propositional* attitude, *and* does not include *truth* amongst the conditions this attitude must meet in order to constitute knowledge. Indeed, because according to *NPI*, practical knowledge is constituted by a non-propositional attitude (intention), *truth could not* be a condition on which this attitude constitutes knowledge. Truth can't be a condition on an attitude which is incapable of being true.

It is true that according to *NPI*, *truth* is not a condition on practical knowledge. But although this may seem incoherent, it is not, because it doesn't undermine the *factivity* of

practical knowledge. For the factivity of practical knowledge comes *not* from a property of its constituting attitude, but a property of its *knowledge-elevator*. Let me explain.

Theoretical knowledge is factive in virtue of the fact that its constituting attitudes – beliefs – are true. Truth is one of theoretical knowledge’s knowledge-elevators for the reason that only by being true will a belief relate a person to a fact, and since knowledge is factive we can demand that the beliefs involved in theoretical knowledge are true. My belief that the pangolin is endangered can be knowledge only if it is a fact that the pangolin is endangered. If it is, my belief might relate me to this fact in the right way for knowledge (it *will* do so just in case the epistemic circumstances in which I hold the belief are favourable).

Practical knowledge is factive not in virtue of being constituted by an attitude (belief) with a true content, but in virtue of involving the execution of the intention, where executing an intention precipitates a fact. I intend to be typing, and I am executing this intention; the upshot is the obtaining of *the fact that I am typing*. In practical knowledge, just as in theoretical knowledge, a person is related to a fact. But the way in which they are related to this fact is different in practical and theoretical knowledge, according to NPI. In the latter, we are related to the facts we know in virtue of (*inter alia*) having beliefs about them; in the former we are related to the facts we know in virtue of *precipitating* them. This precipitation of a fact is what happens when someone *executes* their intention.

The content of theoretical knowledge is carried over directly from the content of the constituting attitude (belief), whereas the content of practical knowledge is determined by the constituting attitude, *and* the fact that the execution condition is met. The relationship between the content of my intention and the content of my knowledge is not one of identity, but it is nonetheless a tight and non-arbitrary relationship. We can understand this relationship better by recalling some thoughts from Chapter Four. There I argued that intentions (at least in their fundamental forms) do not contain reference to their subjects in their contents. The content of my intention is not *that I am typing* or *for me to type*, but is simply *to type* or *typing*. When I execute my intention, I in some sense join *myself* with *typing*, and through doing this I create the fact *that I am typing*; this fact is the content of my knowledge. I will have more to say about the idea of ‘joining myself’ with *typing* in (e).

We can relate these ideas to Anscombe’s thought that practical knowledge displays Theophrastus’ and Aquinas’ Features. Theophrastus’ Feature means that it is by being *executed* (no mistake in performance) that an intention constitutes knowledge. And I have just explained how execution secures an *object* for practical knowledge; the fact that one is acting. And Aquinas’ Feature means that practical knowledge is *the cause of what it understands*. NPI holds that *execution* is internal to practical knowledge; it is practical knowledge’s knowledge-elevator. And since execution is responsible for the precipitation of the fact which is practical knowledge’s object, this means that the phenomenon which gets an agent practical knowledge *of the fact that they are acting* (*viz.* the execution of intention) is the very same phenomenon which precipitates the object of this very knowledge (the ‘what it understands’).

Denying that practical knowledge requires having a true attitude towards the fact known might seem an overly procrustean way of dealing with the *Knowledge-Content* problem. Why should we make these contortions to save the idea that practical knowledge is

constituted by a non-propositional attitude? Well, I have already argued (Ch. 4) that a Propositionalist conception of intention is untenable, which means that *if* any form of *Intentionalism* is correct, then it must be a *non-Propositionalist Intentionalism*. And by rejecting O'Brien's and Paul's accounts (Ch. 2), I made the case that the two main alternatives to *Intentionalism* don't look particularly likely to meet our constraints on a good account of practical knowledge. I hope that by the time I have responded in detail to the *Why Knowledge?* problem, the comments I have made in this section will come to seem less procrustean than they might at present, and that *NPI* will seem worth taking seriously.

e. *Practical and Theoretical Predication*

In (c) I said that in executing my intention to type, I in some sense 'join myself' with typing, thus precipitating *the fact that I am typing*, securing at the same time, knowledge and the object of knowledge. I will call this 'joining' 'practical predication' and will compare it to a kind of predication which is involved in theoretical – belief-constituted – knowledge.

We sometimes talk about predication as something which formally characterises the relationship between constituents of a proposition. In the proposition *that the pangolin is endangered*, *being endangered* can be said to be predicated of *the pangolin*. That the predicate *is predicated of* the subject marks the difference between a proposition and a mere collection of senses. But there is another sense of 'predication', where it is thought of as something a thinker can (in a thin sense) *do*. One way in which *we* predicate is in judgment. In judging *that the pangolin is endangered*, I myself can be said to predicate *being endangered* of *the pangolin*. In so doing I unify distinct representations into a propositionally shaped whole, *viz.* the content of the belief which is precipitated by my judgment. I will refer to predication in this sense – the predication a thinker does in unifying sub-propositional representations into a propositional representation – as 'theoretical predication'. Where knowledge is constituted by belief, it involves theoretical predication, since belief involves theoretical predication.

But just as – according to *NPI* – theoretical knowledge is not the only kind of knowledge, equally theoretical predication is not the only predicating a person can do. In what I will call *practical* predication, a person unifies sub-propositionally-shaped constituents not into a propositionally-shaped representation or mental content, but into a propositionally-shaped worldly item; into a *fact*. I practically predicate *typing* (an act-type) not with a *representation of* myself, but *with myself*. In so doing I unify *typing* with *me*, thus *instantiating* typing and precipitating *the fact that I am typing*. Practical knowledge involves practical predication, and *executing* an intention can be thought of as – practically - predicating the intended action of oneself.

The upshots of practical and theoretical predication – facts and mental (e.g. belief-) contents respectively - share the possibilities of logical shape, but they have different ontological statuses, and for this reason different (although at one level analogous) relationships to the knowledge they are involved in. Whereas theoretical predication is essential to theoretical knowledge's *constituting attitude*, practical predication is essential to practical knowledge's *object*. And where theoretical predication is independent of theoretical knowledge's knowledge-elevators, practical predication is dependent on practical knowledge's knowledge-elevator. Indeed, 'practical predication' is really only a description of what goes on *in executing* an intention. I will have more to say about the idea of practical predication in §4(c), where I will use the idea to defend *NPI* against the

Why Knowledge? problem, and against a new incarnation of the *Knowledge-Content* problem too.

f. Summing up

I have not tried to defend *NPI* in this section, just to describe it in some detail. But I have along the way explained my strategy for responding to the *Knowledge-Content* worry, which involves prying apart issues about the content of the attitude which constitutes a bit of knowledge from issues about the content of the knowledge itself. In theoretical knowledge, these contents match, but they don't – according to *NPI* – in practical knowledge. I explained that the content of practical knowledge was a function of both the constituting attitude *and* the knowledge-elevator, and linked this idea back to Anscombe's idea that practical knowledge displayed Theophrastus' and Aquinas' Features; an idea *NPI* takes over from *NCI*.

I approached the task of explicating *NPI* by considering how the view compares and contrasts practical (intention-constituted) and theoretical (belief-constituted) knowledge. We saw in our discussion of Anscombe that thinking of practical and theoretical knowledge as formally distinct offered a unified way of meeting two of our key constraints (*Apparent Features* and *Special Relationship*) and sat happily with a third (*Structure-Sensitivity*). But it was exactly this idea of a *formal* difference between the two kinds of knowledge that gave rise to the *Why Knowledge?* problem: if practical knowledge isn't structured by the Standard Metaphysics of Knowledge, what makes it a kind of *knowledge*?

I haven't tried to respond to this here, but I have done some groundwork. For one thing we observed throughout this section was that although they look formally different at a certain level of description, zoom out a bit and we can find descriptions which are neutral between the two phenomena. The possibility thus opens up that these neutral descriptions are descriptions of the structure of knowledge *per se*, and that the formally different, lower-level descriptions are descriptions of the different ways in which knowledge can manifest itself. Indeed, this is just what *NPI* holds.

2. The *Why Knowledge?* Problem: Clarifications

It is not controversial to suggest that there is a level of description at which analogies of the kind I described in §1 exist between the practical and the theoretical sphere. Consider for example, this description of the way in which 'knowledge' mirrors 'action' from Timothy Williamson:

Knowledge and action are the central relations between mind and world. In action, world is adapted to mind. In knowledge, mind is adapted to world. When world is maladapted to mind, there is a residue of desire. When mind is maladapted to world, there is a residue of belief. Desire aspires to action; belief aspires to knowledge. The point of desire is action; the point of belief is knowledge.⁵

⁵ Williamson, *Knowledge and Its Limits*, 1. (Williamson is a Knowledge-Firster, but this kind of comparison between knowledge and action can be accepted by JTB-ers too.) Similar sentiments can be found in Danto, *Analytical Philosophy of Action* (Ch. 1); Armstrong, "Non-Inferential Knowledge," 181; Peacocke, *Holistic Explanation: Action, Space, Interpretation*.

What is likely to be controversial is my characterisation of the analogies as analogies between two kinds of *knowledge*. My opponent is likely to insist that knowledge *must* be constituted by belief; that intention cannot even *aspire* to knowledge, let alone constitute it.

If left undefended, the supposition that only belief could aspire to knowledge is dogmatic. But it would also be dogmatic for *me* to claim, without defence, that belief is *not* the only attitude which might aspire to knowledge; that we were supposed to simply *buy* this from Anscombe was one of the main problems I raised for her view, or at least for her presentation of it as I understood things. At least, if Anscombe does provide an argument in favour of the epistemic credentials of *NPI*'s conception of practical knowledge, it is not as clear as one might like.

So I intend to consider directly why the phenomenon I have been *calling* 'practical knowledge' does indeed deserve to be thought of as a kind of *knowledge* – as an epistemic phenomenon; as (in Velleman's words) 'worthy of the name' "knowledge". My strategy has three stages. First, (this section) I try to get clearer on where the problem is supposed to lie. Doing so leads me to clarify the details of the relationship between intentional action, intention-execution and practical knowledge, according to *NPI*. Second (§3), I draw up a list of capacities which commonsense-epistemology links with knowing, and I argue that an intender who is executing or will execute her intention displays all of these capacities, concluding that this gives us *pro tanto* reason to take seriously the idea that intention constitutes practical knowledge, on the condition that it is executed.

In §4 I remove a final obstacle to accepting this conclusion by considering the claim that believing must be essential to knowing, because knowing requires being mentally engaged with the fact known, in a way only believing that fact could account for. I accept the picture of knowing as mental engagement with a fact, but deny that only belief could provide the right kind of engagement: I argue that intention can too. A new version of the *Knowledge-Content* problem comes up here and I respond to it by returning to the idea of practical predication. By the end of §4 I will have completed my response to both the *Why-Knowledge?* and the *Knowledge-Content* problems.

a. What is the Why Knowledge? Problem?

What exactly is the worry about *NPI*'s epistemic credentials? Perhaps it is represented in the following argument:

Why Knowledge? (ident.)

- P_i1) *NPI identifies practical knowledge with the execution of intention, which is itself identical with intentional action.*
- P_i2) *Intentional action and intention-execution aren't knowledge phenomena (concepts) but action-theoretic ones.*
- C_i) *What NPI calls 'practical knowledge' is not a knowledge-phenomenon (concept) but an action-theoretic one.*

I can easily rebut this version of the *Why Knowledge?* problem by being clear that I *do not* want to identify *intentional action*, *the execution of intention* and *practical knowledge*. Instead, I

imagine the three phenomena as sitting in constitution-relations to one another, and as importantly conceptually related.

In a bit more detail, *NPI* sees practical knowledge as constituted by intention, on the condition that it is executed, and *the execution of intention* and *intentional action* are also constituted by intention, on the condition that it is executed (the former trivially so). Given these constitution-relations, which I will describe more fully in (b), the three phenomena are essentially related, and the expressions which pick them out are essentially co-referential. But the three phenomena are not identical, and the expressions which refer to them are not synonymous. So (P_i1) is false and *Why Knowledge? (ident.)* does not go through.

But there is another way of thinking about the *Why Knowledge?* problem; something like the following:

Why Knowledge? (explan.)

- P_e1) *NPI explains practical knowledge in terms of the execution of intention and intentional action.*
- P_e2) *Intentional action and intention-execution aren't knowledge-phenomena (concepts) but action-theoretic ones.*
- C_e) What *NPI* calls 'practical knowledge' isn't a knowledge-phenomenon (concept) but an action-theoretic one.

This argument makes a weaker assumption about the relationship *NPI* asserts between *practical knowledge*, *intentional action* and *intention-execution*. And it is true that *NPI* sees *practical knowledge* as explicable in terms of the other two phenomena (concepts). But this version of the argument also misses the mark. How it does so depends on how we read "explains ... in terms of ..." in (P_e1).

The argument might go through if my plan were to effect a conceptual reduction from practical knowledge to one or the other or both of the other two concepts, since then I would be trying to reduce the concept of practical knowledge to a set of *non-epistemic* concepts, and it would indeed be unclear how what I called 'practical knowledge' really was a kind of *knowledge*. But as I am developing it, *NPI* does not offer this kind of conceptual reduction. So on this reading (P_e1) is false and the argument does not go through.

So what *is* the explanatory relationship between *intentional action*, the *intention-execution* and *practical knowledge*, according to *NPI*? An *elucidatory* one. Elucidation (as I am using the term here) is a kind of non-reductive conceptual explanation. It is a kind of *conceptual* explanation because it involves explaining a concept by tracing its conceptual relationships to others; by thus locating it in a web of concepts. And it is a kind of *non-reductive* explanation in the sense that the norms which govern elucidatory explanation do not impugn elucidations in which the *explanantia* themselves require explanations in terms of the *explananda*.

In general, elucidating a concept of one commonsense conceptual type may well require tracing its relations across commonsense-type boundaries, to concepts of other

commonsense-types. *NPI* elucidates a concept of *commonsense epistemology* in terms of concepts of *commonsense psychology* and *commonsense action-theory*. This in itself ought not to be a controversial procedure. Indeed, it seems to be a feature of the way our commonsense-human understanding is structured that although there are various commonsense conceptual types, these do not have sharp boundaries. And this is likely to be accepted by *NPI*'s opponent, who will likely accept that *belief*, for example (a commonsense-psychological concept) ought to be understood in terms of not only other commonsense-psychological concepts (*desire*, *intention*) but also commonsense action-theoretic ones (of *intentional action*), commonsense-epistemological ones (*knowledge*, *justification*), commonsense-semantic ones (*truth*) and so on.

Elucidation is a kind of conceptual, non-reductive explanation, and elucidating commonsense-human concepts *in general* involves tracing conceptual relations over commonsense-type boundaries, without threatening the *explanans*' claim to be of its initial commonsense-type. Because I am suggesting that we *elucidate* the concept of *practical knowledge* in terms of *intention-execution* and *intentional action* and not that we effect a reduction between these concepts, on this reading of *Why Knowledge? (explan.)*, (P_e1) and (P_e2) are true, but (C_e) doesn't follow.

b. The 'Base Phenomenon' and its three Facets

I have just said something about how *NPI* conceives of the explanatory relationships between the *concepts* of *intention-execution*, *intentional action* and *practical knowledge*. I turn here to considering *NPI*'s account of the relationships between the phenomena which these concepts pick out.

It might be misleading to suggest, as I did above, that all three of these phenomena are constituted by *intention* on the condition that it is executed (the first trivially so), because this suggests that *intention-execution* is the fundamental category, with the other two somehow supervening on it. But I don't want to suggest a picture like this. We should understand *intention-execution*, *intentional action* and *practical knowledge* as phenomena on the same explanatory and ontological level; as each equally a facet of some one essentially three-faceted phenomenon.⁶

For example, currently, a phenomenon is instantiated which can be referred to equally as *my intentionally typing*, or as *my executing my intention to type* or as *my knowingly typing* - where the knowledge in question is practical (*i.e.* knowledge which has the special features attaching to practical knowledge we described in Chapter One).

It will be useful to be able to refer to this three-faceted phenomenon neutrally between its three facets; I will call it 'the base phenomenon'. Sometimes 'base phenomenon' will be used to refer to a *type* and sometimes to a *token*. I will say that a person *instantiates* the base phenomenon (*type-sense*) when they are the person whose *practical knowledge*, *execution of intention* and *intentional action* make up the base phenomenon's (*token-sense*) three facets. From now on I will let context disambiguate between the *type*- and *token*-senses.

The base phenomenon is essentially three-faceted, which means that it is not properly instantiated unless all three of its facets are instantiated. And the facets are conceptually

⁶ Adrian Haddock also suggests that practical knowledge is in some sense an *aspect* of intentional action itself (Haddock, "The Knowledge That a Man Has of His Intentional Actions," 162). One way of thinking about the *Why Knowledge?* problem is as the problem of explaining how this can be so.

related, which means that someone is capable of recognising one facet only if they are capable of recognising its two others. It would not be possible, for example – for me or anyone else – to think of me as *intentionally feeding Egbert* unless it were also possible to think of me as having knowledge of doing so which displayed practical knowledge's three apparent features, or to think of me as *executing* rather than merely satisfying my intention – although of course it is not a requirement that one must attach the expressions 'practical knowledge' and 'execution' to these concepts, or indeed any particular expressions at all.

In saying that the base phenomenon is essentially three-faceted, I am saying *both* that *intentional action*, *intention-execution* and *practical knowledge* cannot be pulled apart – whether metaphysically or explanatorily – *and* that they are nonetheless *distinct* – again, metaphysically and explanatorily. How is this supposed to work? I will try to explain why we should think of these phenomena as distinct, and then why we should think of them as nonetheless essentially bound up with one another.

We can see that the three phenomena are *distinct* facets of the base phenomenon by considering their contrast-classes. To return to an earlier example, I might be thought of as typing 'causal relationship' *intentionally* rather than *accidentally*, or I might be thought of as *executing* an intention rather than failing to do so in either of the two ways we considered above (*no* satisfaction or *mere* satisfaction). Finally, I might be thought of as typing 'causal relationship' *knowingly* rather than *unawares*, or (e.g.) of knowing 'all along' rather than knowing because it is something I have noticed.

Under which of these aspects an instance of the base phenomenon is thought of will depend on our interests in thinking about it. For instance, the notion of *intention-execution* is well-suited to highlight an agent's steadfastness, particularly in a context where such steadfastness is surprising or noteworthy: "She said she intended to run the marathon and I didn't believe that she would – but credit where credit's due: she *went through with* it." It is obviously *true* that she didn't run the marathon *by accident*, or *unawares*, but pointing to these facets of the base phenomenon would not help our speaker make his point.

We might think of the base phenomenon under its *intentional action* aspect when the kind of issues which played such a central role in our discussion of Anscombe are to the fore; when we want to highlight its purposive or goal-directed character and its means-end structure. Thinking of the action as an object of practical knowledge, or as done in execution of an intention doesn't seem to pick up on these features of the base phenomenon. In thinking of the base phenomenon under its *intention-execution* aspect, one seems to characterise it in a *backwards-looking* way, whereas in thinking of it under its *intentional action* aspect, one seems to characterise it in a *forwards-looking* way.

I'd like to suggest that we think of the base phenomenon under its *epistemic* aspect; as *practical knowledge*; when we are interested in highlighting the agent's *cognitive involvement* with what she is doing, when we want to highlight the special kind of *ownership* which we have over what we do intentionally. In thinking of someone as *knowingly* engaged in (say) kicking the cat, we see her kicking the cat as a product of her rational mind and as something she identifies herself with; something which she thinks of as expressing her motivations and perhaps her personality in a broader sense too. Thinking of people as cognitively involved in what they are doing in this way also seems important to understanding the *degree* and *kind* of personal, moral and legal responsibility they bear in relation to their actions.

Thinking of the base phenomenon's three facets as distinct also helps us to see more clearly what's going on in the 'potting the pink' example. Because the three facets of the base phenomenon are *different*, there will be cases in which we are happier to apply some of the concepts than others. How happy we are to apply one of the concepts will depend on how happy we are to apply related concepts. For example, because *acting intentionally* is closely related to *intending*, *aiming* and *trying*, and to *being responsible* for an outcome, and because we are happy to think of Bev as trying, aiming and intending to pot the pink, and as to *some* degree responsible for it's going in, we have *some* reason for saying that she potted it intentionally (even if not as much reason as we have in Ronnie's case, because – as I said at the end of Chapter Five – she owes a large debt to luck). Because *practical knowledge* is closely related to e.g. *knowing all along*, and *being confident* about the fact that something is taking, or will take, place, and because these concepts *do not* apply to Bev, (she had to *wait and see* how the balls dispersed until she knew it would go in the pocket; she wasn't *sure* it was going to go in whilst taking her shot), we have no reason at all to say that she potted it with practical knowledge.

The fact that the three essential facets of the base phenomenon can be differentially instantiated does not undermine thinking of practical knowledge as an aspect of intentional action (and of the execution of intention). All it means is that the base phenomenon can be instantiated imperfectly. In Bev's case, the base phenomenon is like a car which cannot get one from *A* to *B* (see again Ch. 1, §2(b)). Such a car still counts as a car because it has many features which are essential to the concept of being a car. But it is a *dud* car in that it lacks this one. Bev instantiates the base phenomenon in virtue of the fact that certain concepts relating to her *intentionally potting the pink* are in place - she intended to, she tried to, and she pulled it off; her doing so was not *entirely* down to luck. But the base phenomenon she instantiates is a dud, in virtue of the fact that certain other concepts do not apply to it – she did it more out of luck than out of *skill*, she wouldn't be able to *repeat* it, she *didn't know she was doing it*. And even though we *might* think of her as potting the pink intentionally, hers is not a central or straightforward case of intentional action.

So the three facets really are different. Why, then, think of them as facets of *the same phenomenon*, as conceptually related? I will consider just one reason here, *viz.* that thinking of them as conceptually related seems to explain certain kinds of exchange which would otherwise be puzzling. Consider the following two snippets of conversation:

- “Did you mean to upset Dan?”
 - - - “Goodness, no! I didn't realise I had until he'd left.”
- “Do you know where Superman is?”
 - - - “Clark Kent is at the *Daily Planet* offices.”

The cases feel quite different: the first feels quite natural whereas the second seems odd. But why don't *both* feel odd? Both look like cases in which a question is asked about one subject-matter but answered in relation to another. It feels as if some conversational maxim is flouted in the second exchange but not in the first. Why not?

According to *NPI*, both are cases in which a question is asked about an object under one of its aspects, and answered in relation to the same object, under another of its aspects. According to *NPI*, the difference is that only in the first case is a conceptual relationship

recognised between the two aspects: possessing the concepts of practical knowledge, intentional action and intention-execution involves recognising their interconnections, and this is why I can answer a question about whether I *intended to* upset Dan by saying something about my *knowledge* of what I was doing, or lack of it. In suggesting that one aspect of the base phenomenon (*practical knowledge*) was uninstantiated I *thereby* suggest that its other aspects (*intentional action* and *execution of intention*) were uninstantiated too.

Even stipulating that in the second case, Clark Kent's identity with Superman is common knowledge, the respondent still does not appear to be properly answering the question asked. This is because although both aspects of the object are known by both parties, the aspects are not conceptually related (so *a fortiori* not recognised as such). This means that responding to the question by giving *Clark Kent's* whereabouts requires an implicit "... and (as we know) Superman *is* Clark Kent", which lends an unnecessarily complicated air to the response. But no further implicit qualification is needed for the response "I didn't know I was upsetting Dan" to count *as* a response to the question, because the relations between the concepts of executing an intention and having practical knowledge are built in to the interlocutors' grasp of the concepts themselves. They recognise that practical knowledge and intention-execution are essentially aspects of a single underlying phenomenon; what I have been calling the 'base phenomenon'.

c. The Revised 'Why Knowledge?' Problem

I have been trying to say more about the exact relation I take there to be between the concepts and phenomena of *intentional action*, the *execution of intention* and *practical knowledge* as a way of showing why my development of *NPI* is not susceptible to the two versions of the *Why Knowledge?* problem we considered in (a).

In so doing I have come down on the 'non-reductive' side of the debate about intention-execution which I discussed in §1(c). This means, by the way, that *NPI* is likely to appeal more to someone who is already attracted to the non-reductive side of the debate about *theoretical* knowledge; a Knowledge-First-er rather than a JTB-er. I don't think there are any logical barriers on accepting a reductive theoretical epistemology and a non-reductive practical epistemology, but it would be a strange choice. At any rate here is not the place to pursue such questions; I mention the importance of my non-reductive development of *NPI* only for the sake of transparency.

But even building non-reductionism into *NPI*, a version of the *Why Knowledge?* problem remains. For although I have stated a view of the relationship between the three phenomena/concepts, on which there is no *conflict* between thinking of one of the three as epistemic and the other two as not so, I still need to convince the reader that the base phenomenon *has* an epistemic aspect; that there is a phenomenon which has not only *intentional action* and the *execution of intention* as facets, but *practical knowledge* too. I am in a position to do this now.

3. Why Practical Knowledge is Knowledge

My argument that the base phenomenon - the phenomenon which is identifiable as both acting intentionally *and* as executing one's intention – also has an epistemic aspect will involve showing that the capacities and feature associated by commonsense with *knowing* are all possessed by someone instantiating the base phenomenon. I will thereby have traced the conceptual relationships between instantiating the base phenomenon, and instantiating the phenomena associated with knowing, and this will give us a *pro tanto* reason for accepting that intentions really *can* constitute knowledge, when the execution condition is met.

So what does knowing involve? Our methodology demands that this question be answered by *commonsense*, not *philosophical*, epistemology. To answer this question is to elucidate the commonsense concept of knowledge. The following list represents a plausible collection of the central epistemic capacities possessed by a person *qua p*-knower.

- a) Knowing that *p* puts one in a position to “provide a correct answer to a potential question”,⁷ and to proffer this information in providing accurate testimony.
- b) Knowing that *p* puts one in a position to use the fact that *p* as a reason for “[m]aking a decision or an inference, believing, doubting, performing a calculation in one's head, doubting or believing or hoping something, conceiving a desire or forming an intention”.⁸
- c) Knowing that *p* allows one to stop thinking, worrying or considering whether *p*.⁹
- d) If someone knows that *p*, then she can explain how she knows it, and can thereby give an interlocutor reasons for themselves believing that *p*.

Obviously I need to be careful at this stage not to beg any questions, either in favour or against *NPI*. So I have derived this list from a consideration of what *NPI* (along with Anscombe's *NCI*) call *theoretical* knowledge; what *NPI*'s opponent thinks of as knowledge *per se*. I will argue that someone instantiating the base phenomenon straightforwardly displays the capacities described in (a)-(c), and displays one element of that in (d). But I will also explain that the other element of (d) is not in all cases to be expected of someone who has theoretical knowledge (again, for my opponent: who has knowledge *per se*), so *NPI* will remain unharmed.

My opponent is bound to complain that I have left an item off my list which is both central to a commonsense conception of knowing *and* damaging to *NPI*: doesn't commonsense epistemology link knowing with *believing*? I will leave this question to one side until §4. (This is why I said above that at the end of this section we will have a *pro*

⁷ Hyman, “How Knowledge Works,” 436. Hyman is here quoting White, *The Nature of Knowledge*, 119. White seems to think of knowledge as *defined* in the way described. I am not using the idea in this way here. Instead, I am characterizing *being able to provide a correct answer to a potential question* as a plausible commonsense concomitant of knowing that *p*.

⁸ Hyman, “How Knowledge Works,” 439. For Hyman this condition characterises knowledge, but as with the capacity in (a) I am treating this capacity only as a concomitant of knowing.

⁹ Thanks to Alison Fernandes for this suggestion.

tanto reason for thinking that NPI can overcome the *Why Knowledge?* problem.) Let's consider the items on our list in turn.

a. Correct answers and accurate testimony

Does instantiating the base phenomenon put one in a position to 'provide a correct answer to a potential question', and to proffer this information in providing accurate testimony?

Yes. Intentions can be expressed by making statements about what one will do, or is doing, and there is no barrier to thinking that someone can, in so doing, provide an interlocutor with accurate testimony. Her testimony will be *accurate* on the condition that the intention she expresses is being, or will be, executed.

An example: "What are you doing tomorrow at 3pm, Bev?" asks Henry. "I'm going to the hairdresser's", says Bev. *If* this is a case in which Bev's intention (to go to the hairdresser's) will be executed, then Bev has provided a correct answer to Henry's question; she does so by expressing her intention. Fast-forward to that evening: "What are you up to at the moment Bev?" asks Henry on the phone. "Getting ready for a night out with the girls", says Bev. *If* this is a case in which Bev's intention *is getting* executed, then again Bev provides a correct answer to Henry's question; and again by expressing her intention.

It is her having and executing her intention which, in both the present- and the future-directed case, explains how Bev is able to provide a correct answer to Henry's questions, and these correct answers constitute accurate testimony. And although in our example Bev is prompted by Henry's questions, there is no *need* for her to be. She *could* tell Henry about her intentional actions, present and future, quite spontaneously – again, by expressing her intentions, and again thereby providing testimony which will be accurate just in case her intentions are being, or will be, executed.

So instantiating the base phenomenon puts Bev in a position to provide a correct answer to a question, and to provide (spontaneous) accurate testimony.

b. Using what one knows as a reason

The idea in (b) comes from John Hyman, who wants to *identify* knowledge (*per se*) as:

... the ability to do things, or refrain from doing things, or believe, or want, or doubt things, for reasons that are facts. [...] [T]he list – doing, refraining from doing, believing, wanting, doubting – is not exhaustive. Any verb can be added which can occur in a sentence of the form '*A*'s reason for [Ψ]-ing was that *p*'.¹⁰

Whether or not Hyman is right to *identify* knowledge in this way, it seems very plausible that someone who knows that something is the case will have the capacity Hyman describes. Indeed, it is hard to see what the *point* of knowing that *p* could be if it did not furnish one with a capacity like this. So does instantiating the base phenomenon give one this capacity?

¹⁰ Hyman, "How Knowledge Works," 441. I have changed Hyman's ' φ ' to ' Ψ ' because ' φ ' has been used throughout my discussion differently to how Hyman uses it in his.

Yes. We have already seen that Bev can accurately testify about what she is doing just on the basis of instantiating the base phenomenon, and being able to express her intention. The examples in (a) are in fact *already* examples in which she uses a fact as a reason for Ψ -ing in some way. For one of her reasons for *telling Henry that she is getting ready for a night out with the girls* is the fact that she *is doing so*. This is not her *only* reason for telling Henry this, of course, but it needn't be to meet Hyman's condition.

By instantiating the base phenomenon, Bev is able to use the fact that she is getting ready for a night out as a reason for Ψ -ing in various other ways too: she can use it as a reason to *put her lucky dress on*; to *call Pam* and *decide* how to get to the club, and to *intend* to drink a litre of water before bed. And the same goes for the future-directed case. Bev could use the fact that she *will go* to the hairdresser's later as a reason for *getting* money from the cash-point, for *worrying* about her bank-balance, for *believing* that by 4pm she'll look a million dollars, and so on.

So it seems that instantiating the base phenomenon puts Bev in a position to use the fact that she is acting or will act as intended, as a reason for Ψ -ing in the various ways in which Hyman has in mind.

c. Closing down enquiry

The thought behind (c) is that if someone knows that p , then the question of whether or not p is the case can be closed for them; *enquiry into whether p* is no longer necessary for someone who *knows* that p . Does instantiating the base phenomenon allow one to stop thinking worrying or considering whether p ?

In many cases, it clearly does. It is hard to make sense of someone *currently* doing something intentionally, but continuing to enquire about whether they are doing it. They may wonder whether they are doing it *well*, or whether they will *complete* it, but for reasons we have already come across (see Ch. 2, §1(d) & §2(c)), enquiring into these things is different from enquiring whether one *is doing* what one is doing.

What about people who instantiate the base phenomenon in virtue of now having an intention which *will be* executed? In many cases here too such a person will not make further inquiries into whether they will act. Bev seemed to regard it as *settled* that she would go to the hairdresser's later when she expressed her intention to Henry. She didn't hedge her claim with saving clauses, and we didn't get the feeling that she ought to have done.

Indeed it seems as if it *must* be the case that forming an intention puts a stop to further enquiry about whether or not one will act, at least typically. It's hard to see what the point would be of forming intentions to do things if we then always, or even usually, still took it as up in the air as to whether we would act as intended. And it's hard to see how expressing our intentions could play the role it does play in our lives if expressing an intention for the future were only ever expressing a tentative guess about what one was likely to do. You and I would not be able to make plans together, for example, if we both understood the truth-value of my statement "I will be at the cinema at six" as requiring further investigation up until I executed the intention. So I think we can safely assume that at least typically, enquiry will be closed for someone instantiating the base phenomenon.

But if instantiating the base phenomenon is a way of knowing, which is what I am arguing on *NPP*'s behalf, then it needs to be the case that in instantiating the base phenomenon, one *thereby* possesses the capacities commonsense links with knowing. If (c) is indeed one of these capacities, then instantiating the base phenomenon must in all cases underwrite this capacity. So all instantiations of the base phenomenon must give an agent the capacity to cease enquiring into whether she is acting or will act. And the worry is that *certain* cases of what *NPI* classes as future-directed practical knowledge *don't* look like cases in which the agent ceases enquiring into whether she will act as intended.

An example: I intend to call Henry this evening, and (we will assume) my intention will get executed when the time comes. Thus I now instantiate the base phenomenon and according to *NPI* I have practical knowledge. So I ought to cease enquiry about whether I will call Henry. However, calling Henry is something I often 'forget' to do, especially when Julia is around and the sun is out and the pubs are open. I rack my brains trying to remember Julia's schedule for this week – did she say she was free tonight? Or was she working late? Given the relationship between Julia's being free and the likelihood of my 'remembering' to call Henry, this can count as a kind of 'further enquiry' into whether or not I will call him. So, the thought goes, the base phenomenon obtains because I *have* an intention to call Henry, which *will (ex hypothesi)* get executed (Julia is, it turns out, working late), but (c) is not met, since I have not ceased enquiring into whether I will call him.

The first thing to say in response to this worry is that the same *kind* of case arises in agreed cases of knowledge as standardly understood – what *NPI* classes as 'theoretical' knowledge. Consider a version of the familiar 'unconfident schoolboy' example.¹¹ Our boy is very insecure, but pretty competent, and is in the midst of revising for his history exam. He is doing an online test which consists of ten questions. He fills in his answers, getting (say) 8/10 correct (he doesn't know this yet). At this stage he can click a button to get his results, but his lack of confidence about his own ability leads him to hesitate, return to his books and check his answers - he doesn't want his online score to be poor because it'll be bad for his confidence and lessen his chances of doing well in the exam. Looking up the answers, he realises – with surprise – that he has got 8/10. This lad – it is plausible to say – *knew* the answer to (say) question 6, "In what year was the Battle of Hastings?", but carried on his enquiry (by checking his books) nevertheless.

The description of the case as one in which he *knew* the answer to question 6, yet continued his enquiry, is allowed for by our commonsense-human concepts, which means that (c) appears too strong as a commonsense condition on, or concomitant of knowing even as *standardly* understood. Having theoretical knowledge that *p* is consistent with a kind of epistemic low self-esteem which might well lead the *p-knower* to carry on their enquiry as to whether *p*. This won't be the typical case, but it's perfectly intelligible.

What should we conclude? I think it would be an overreaction to reject (c) altogether as a capacity linked by commonsense with knowing, because (c) does seem to capture something important. For although our boy *did keep* enquiring into when the Battle of Hastings was fought, there is a sense in which he *could have* ceased doing so. The example is precisely of someone who is *over-cautious*, and it seems very plausible that we think of him as *over-cautious* in checking his answer to question 6 precisely because we think of him as *knowing* it. Observing him we might want to tell him to calm down, stop worrying and trust his first instincts: "Come on Chris, you *know* this stuff!" And we would be right

¹¹ An argument to the effect that confidence is not necessary for knowing is given in Woozley, "Knowing and Not Knowing."

to do so – his score attests to that. The happier we are to think that he knows the answer to question 6, the more otiose and neurotic his double-checking appears and the more reasonable further enquiry appears, the less happy we would be to think of him as having *knowledge*.¹²

So I think that there is an important relationship between *knowing* and *ceasing enquiry*, but it is not an entailment relationship but a permissive-normative one. Someone who knows that *p* is in a position to cease enquiring into whether *p* – any such enquiry is in a certain sense otiose. But in certain cases a *p*-knower may well *de facto* continue enquiring. Knowledge is consistent with occasional lapses of epistemic self-esteem.¹³

Chris still enquires into when the Battle of Hastings was fought, and I still enquire into whether I will call Henry this evening. And this is consistent with enquiry being, in both our cases, *otiose*. In further enquiring into the facts that we know, we engage in some kind of irrationality; in epistemic overkill. Sure it might seem perfectly reasonable from *my* perspective to consider whether I will call Henry, but it seemed perfectly reasonable from Chris's perspective to check his answers in his text-books. Further enquiry can *be* otiose without *seeming* so.

But there is a disanalogy between Chris's case and mine, for there seems to be a *special* kind of difficulty for the suggestion that having an intention which *will be* executed, renders further enquiry otiose. The worry is that it just doesn't seem as plausible that because I *now* intend to call Henry, and because this intention *will be* executed (*ex hypothesi*), further enquiry *is* otiose. Intuitively, I *do* need to keep enquiring because I don't know that my intention will get executed, and because I have fairly good reason to think it may not.

Is it otiose for me to keep enquiring into whether I will call Henry this evening, on the assumption that I intend to and that I will in fact do so? That depends on the case. One question is whether I am best described as straightforwardly intending to call him. I have described myself as 'forgetful' when it comes to calling him – especially when I'd rather be in the pub. The quote-marks indicate that this not a forgetfulness I try very hard to combat. If the case is like this, then pressure is put on my claim to - *really* – intend to call Henry. Alexander knows what I'm like; "Do you *really* intend to call Henry?" he might ask. And the most honest answer might be "-ish", or "Sort of".¹⁴ This would mean that I don't straightforwardly instantiate the base phenomenon (I might instantiate it *-ish* or *sort of*), which would allow that practical knowledge is not straightforwardly instantiated either.

Let's say I *do* straightforwardly intend to call Henry – I answer Alexander's question by saying "Yes, I really do intend to call him – it's just that I know I often forget." In this

¹² David Lewis describes this kind of case as one in which Chris *knows* but does not *believe* that *p* (Lewis, "Elusive Knowledge," 556; see also Williamson, *Knowledge and Its Limits*, 42). If he is right then this would be a hindrance to my view, in that it conflicts with my understanding of *theoretical* knowledge as *belief-constituted*. But I am not sure of Lewis' intuition about this case. Perhaps he has in mind that we can't be confused about what we believe, but this claim doesn't seem convincing. It might not be *rational* for Chris to fail to recognize his own belief, but it doesn't seem *impossible*.

¹³ Woozley again: "... to say of a man that he knows that *p*, although he is not sure of it is to criticize him for [...] not being sure of something of which he ought to be sure." (Woozley, "Knowing and Not Knowing," 159.)

¹⁴ I do not have, in Hampshire's phrase, a "serious practical intention" (Hampshire, *Freedom of the Individual*, 77).

kind of case, however, conducting further enquiry into whether or not I will in fact call him doesn't seem the appropriate thing to do. The appropriate thing to do if I *really do* intend to call him would be to find a way to mitigate my known forgetfulness; for instance to set an alarm on my phone saying "Ring Henry!" Having done this, though, further enquiry *does* start to seem otiose and (c) is met.

So I think that instantiating the base phenomenon – at least *straightforwardly* doing so – does render further enquiry otiose. This is the case both for someone who *is executing* their intention, and for someone who *will execute* their intention, although the latter case brings up a complexity not present in the former. The complexity has highlighted that the base phenomenon may fail to be straightforwardly instantiated. But in such a case, I have suggested, all three of its constituting phenomena will be instantiated in a suspect way. So this doesn't cause a problem for my claim that the base phenomenon has an epistemic aspect, and that its having this aspect follows from the fact that instantiating it gives one the capacities which commonsense attaches to knowing. The happier we are to say that the base phenomenon is straightforwardly instantiated, the happier we are to say that further enquiry is otiose. So someone instantiating the base phenomenon displays capacity (c).

d. Knowledge and reasons

The final member of our list links *knowledge* with *reasons*. The idea has two parts. First, someone who knows that *p* can give the grounds on which their knowledge is based, and second, they can give someone else a reason to believe that which they know. Call these two parts of (d), (d') and (d'') respectively. So, does someone instantiating the base phenomenon possess these capacities?

Someone instantiating the base phenomenon clearly seems to possess (d''), for a reason we have already seen. Bev could *testify* to Henry about what she was doing, and about what she was going to do. And testifying to what one knows is a way of giving someone a reason to believe what one knows.

But someone instantiating the base phenomenon will not possess (d'). According to *NPI* Bev knows that she is getting ready to go out *simply* in virtue of instantiating the base phenomenon. Her knowledge is not constituted by belief, and it *has no* epistemic grounds. Bev can't *give* the epistemic basis of her knowledge, because it *has none*. Although Bev can testify that she will go to the hairdresser's later, and this testimony can be a reason *for Henry* to believe that Bev will, it is not *Bev's* reason.

This would be a problem for *NPI* if (d') appeared – from the commonsense-human perspective – to be a necessary concomitant of knowing. (Indeed, it would be a problem for any conception of practical knowledge which seeks to *accommodate* rather than to *explain away* its *epistemic ungroundedness*.) But (d') is *not* a capacity possessed even by all bearers of agreed pieces of *theoretical* knowledge. A first kind of case is when someone might seem to have *some* grounds, but perhaps because there are so many minor grounds it is not obvious what to say when asked for them. My knowledge of what my name is, when I was born, where I live, that I have a brother, that cats typically have tails, looks like this. I wouldn't know what to say if asked how I know each of these things, since anything I could point to would appear to be so incomplete an answer as to be misleading.

A second kind of case is one's knowledge of one's own psychological goings-on. How do I know that I *believe* Egbert needs feeding, that I *feel* slightly hungry, that I *hope* it's not raining when I have to walk into town later? There might in some circumstances be things to say in response to these questions ("I noticed myself looking at her chips!"), but there *needn't* be anything to say. And most will agree that for all that, these are cases in which I have knowledge.

These cases might seem unhelpful to some for reasons stressed by Wittgenstein, who worried about the idea that in ordinary kinds of case, we can be straightforwardly said to *know* propositions like those in the examples above. I won't get embroiled in those considerations here. For anyone finding the examples in the previous paragraph dubious, I can point to examples in which we know that someone must have *some* grounds, but in which these grounds are not apparent to the knower:¹⁵

"How do you know she's upset? She doesn't *look* sad."

- - - "I can just tell – don't ask me how".

"How do you know so much about the Kodiak bear?"

- - - "I really can't remember."

In light of these considerations, we cannot require that *if* the base phenomenon has an epistemic aspect, *then* anyone instantiating it would possess the capacity to give the grounds on the basis of which they know that they are acting or will act. If this was required for some phenomenon to count as a kind of knowing, then even agreed cases of theoretical knowledge would not pass the test.

The capacities which come under (*d'*) and (*d''*) come apart in these theoretical cases too. For even someone whose theoretical knowledge has grounds, which they cannot give, can give *someone else* grounds for believing what they know: "I know it's true, you'll just have to take my word for it." Circumstances will dictate whether this purely testimonial basis is a *good* reason to believe the thing in question, but the requirement we are considering is only that a knower is able to give *a* reason.

So we should drop (*d'*) from our list, and maintain (*d''*). And instantiating the base phenomenon seems to put someone in possession of (*d''*).

e. Summing up

I think, then, that instantiating the base phenomenon underwrites the possession of *all* the capacities on our list which seem essential to *knowing* as commonsense epistemology understands it. Only (*d'*) did not seem to come along with instantiating the base phenomenon, but (*d'*) did not seem to essentially attach to knowing as standardly understood (to what *NPI* calls 'theoretical' knowledge). Possession of the four capacities (excepting (*d'*)) appears to accrue to someone who is executing (or who will execute) their intention no less than to someone who has a true belief held in favourable epistemic circumstances.

We thus have *pro tanto* reason for accepting that the base phenomenon has an epistemic aspect; that *intentionally φ -ing* and *executing one's intention to φ* are two facets of a

¹⁵ See also Lewis, "Elusive Knowledge," 551.

phenomenon of which *knowing that one is φ -ing* is a third. The idea is that *simply* in virtue of instantiating the base phenomenon, an agent is in the position of a knower, since simply instantiating the base phenomenon puts someone in possession of a set of capacities which seem to be linked by commonsense, essentially, to knowing. If this is right then knowing *per se* can be manifested in two ways – by belief *or* by intention; by instantiating the Standard Metaphysics of Knowledge *or* by instantiating what I have been calling the base phenomenon.

The reason is only *pro tanto* because *NPI*'s opponent can object by finding some *other* feature of a person who is classified by commonsense as a *p*-knower, which is *not* a feature of someone instantiating the base phenomenon; I haven't claimed that my list of capacities is exhaustive, just that the item on it are *central*. In the next section we will consider whether *p*-knowers must be such in virtue of being *p*-believers. If they are, then *NPI* will have been undermined.

4. Knowledge and Mental Engagement

In this section I will consider whether commonsense links knowing to believing in a way that is damaging to *NPI*. *NPI* will not be damaged by the conclusion that knowing *entails* believing unless we have further reason to think that this entailment tracks a constitution-relation between knowing and believing. So I will explicitly consider whether we have any reason to think that the commonsense conception of knowing that *p* is one which sees knowing as *constitutively involving* belief rather than intention.

I will start in (a) by thinking about how we might get to the idea that all knowledge has to constitutively involve belief. The idea will be that knowing that *p* requires being *mentally engaged* with the fact that *p*, where this *mental engagement* is what explains how a knower possesses the epistemic capacities considered in the previous section. I then go on to consider in (b) and (c) whether a person can be mentally engaged with the fact that she is φ -ing in virtue of intending to φ , and executing this intention. A *prima facie* problem for the idea that she can comes in the form of a new incarnation of the *Knowledge-Content* problem. As it arises here the worry is that we can't make sense of an agent being able to *express* her knowledge that she is φ -ing by expressing an intention, because her intention does not represent the fact *that she is φ -ing*, it represents at most *φ -ing*. I explain why I think that this worry stems from a suspect conception *knowledge-expression*. I suggest a different way of understanding *knowledge-expression* which is both not suspect and which leaves the idea of being able to express *propositional* practical knowledge in virtue of having a *non-propositional* attitude (intention) intact. In (d) I respond to a final objection. I sum up in (e).

By the end of this section I will have completed my case in favour of *NPI* and answered both the *Why Knowledge?* problem and the revised *Knowledge-Content* problem. I will conclude in §5 that *NPI* ought to be taken seriously, at least as a basis for further investigation.

a. Why think knowing entails believing? (Mental engagement with a fact)

I am not aiming here to rule out that someone instantiating the base phenomenon also believes that she is acting as she is. Perhaps we even *always* have beliefs about our present and future intentional actions. What I want to argue against is the idea that when

someone's knowledge is practical knowledge, the fact that they know that they are acting as they are is to be explained *inter alia* in terms of the fact that they *believe* that they are acting as they are.

I have already set out my case in §3 for thinking that someone will possess all the epistemic capacities which commonsense links to knowing *simply* in virtue of instantiating the base phenomenon. But it is consistent with everything I said in §3 that someone instantiating the base phenomenon *only* has these capacities because she *also* believes that she is acting as intended.

It is hard to avoid begging questions here. We need to find a way to arbitrate the dispute between *NPI* and its opponent without assuming *either* that all of the epistemic capacities in §3 attach to the base phenomenon *independently* of any beliefs an agent might have about what she is doing *or* that they attach to the base phenomenon only because the agent has such beliefs in addition to instantiating the base phenomenon.

Here's the tactic: I will consider an agreed case of *belief-constituted* knowledge, and think about what contribution to the person's possession of the epistemic capacities in §3 is made by the fact that she believes that the fact known obtains. We can then ask whether the same kind of contribution can be made independently of belief, in virtue of instantiating the base phenomenon. That sounds a complicated task but it is made much simpler once we recognise that in fact the four capacities I listed in §3 actually all seem to be grounded in two more fundamental capacities; the capacity to *express* what one knows and the capacity to draw *inferences* – both practical and theoretical – from what one knows. So we really only need to ask, for an agreed case of theoretical knowledge, what it is about *believing that p* that explains the person's capacity to *express* what she knows, and to *draw inferences* from it.

I know that the pangolin is endangered, *inter alia* in virtue of believing that the pangolin is endangered. This is our starting-point. What is it about my belief which underwrites my capacity to *express* this knowledge, and to *draw inferences* from it?

I think the answer must be that in *believing* that the pangolin is endangered, I am *mentally engaged* with the fact that the pangolin is endangered. What do I mean by *mental engagement*? Two things about my belief seem important: first, my belief *represents the fact* that the pangolin is endangered, and second, my belief *commits me to the obtaining of this fact*.

Having a *representation of* the fact that the pangolin is endangered seems important to my capacity to express my knowledge. Why is this? Well, to put the thought rather metaphorically, I can't express the *fact* that the pangolin is endangered; I can only express my own 'internal' features. We might think that it is possible for me to express my *knowledge* that the pangolin is endangered only by expressing some internal condition which relates in the right way to the fact that I know. My belief is just such a condition: it relates to the fact I know *by representing it*, and it is an internal condition of me, and so something I can express. Having a belief which represents the fact that I know thus seems central to my capacity to express my knowledge.

Being *committed to* the obtaining of the fact that the pangolin is endangered seems important to my capacity to make inferences from my knowledge. Why is this? Well, compare a state which *represents* the fact that the pangolin is endangered but which does not commit one to the fact's obtaining. Say that Ellisif *imagines* (but does not believe) that

the pangolin is endangered, and because the pangolin *is* endangered, Ellisif thereby represents a fact. It seems that I can draw practical and theoretical inferences from the proposition that the pangolin is endangered which Ellisif cannot. Consider the case for theoretical inference-drawing. Ellisif might recognise the same logical relations between the proposition that the pangolin is endangered, the proposition that the pangolin is the only scaled mammal and the proposition that the only scaled mammal is endangered. But recognising these relations alone does not render Ellisif able to *conclude* that the only scaly mammal is endangered. I *can* draw this conclusion. The difference seems to be exactly that in *believing* that the pangolin is endangered I am committed to the obtaining of the fact that it is, whereas in *imagining* that the pangolin is endangered, Ellisif is *not* so committed.

So my belief that the pangolin is endangered seems to explain my capacity to *express* and *infer from* what I know in virtue of its being a kind of *mental engagement* with the fact that the pangolin is endangered. And my belief seems to deliver the requisite mental engagement in virtue of *representing* the fact that I know and *committing me* to its obtaining.

b. Intentions as commitments

Now we can try to arbitrate the dispute between *NPI* and its opponent. The opponent's idea is that *only* a belief could provide the relevant kind of mental engagement with a fact to underwrite the capacity to express and to infer from what one knows. *NPI* must deny this. I suggest here that it is relatively uncontroversial that intentions *commit* a person to the obtaining of the fact that they are acting, or that they will act. In (c) I move on to a discussion of *representation* and *knowledge-expression*, where things will be a bit more complicated.

I don't see any reason to deny that intentions are themselves a kind of commitment to the obtaining of a fact. First, we already know that one can express intentions by making a claim about what one is doing or about what one will do. "I am going to the cinema tomorrow" looks like it expresses a commitment to the obtaining of the fact that I will go to the cinema. My reader might worry that I am here assuming something my opponent will not accept. Perhaps – they might suggest – I can only make assertions like this one in virtue of *believing* that I will go to the cinema. I think that in disagreeing on this fact about intentions my opponent may be taking things a step too far away from our commonsense-human methodology. (For a related argument recall my objection to Castañeda's thoughts on the relation between an attitude's linguistic expression and its content in Ch. 4, §3(b).)

At any rate, I don't think *NPI*'s opponent can make the consideration stick. For quite apart from how we express intentions, consideration of the kind of attitudes they are seems to entail that they commit their subjects to the fact that they are acting, or will act. Intentions aim at execution.¹⁶ And the execution of an intention (on *anyone's* view) precipitates a fact. We would not bother forming intentions unless we implicitly accepted this about them. So I think we can happily accept that intentions commit their subjects to the obtaining of facts. My intention to go to the cinema tomorrow commits me to the fact that I will go to the cinema tomorrow, and my intention to be typing commits me to the obtaining of the fact that I am now typing.

¹⁶ According to *NPI* this means that they also aim at intentional action, *and* at practical knowledge; but I can't pursue this here.

I want to be clear why I am not contradicting myself in saying this. I have argued against a view (Cognitivism about Intention) according to which intentions are commitments to the obtaining of facts *in the same way* that beliefs are. Beliefs commit a person to the obtaining of a fact by committing her to the *truth* of the *propositional content* of the belief. And according to *NPI* beliefs don't *have* propositional contents, so they cannot be commitments to the obtaining of facts in quite this way. But thinking of intentions and beliefs as both committing their subjects to the obtaining of facts does not require thinking of them as doing so in just the same way. I have outlined the idiosyncratic way in which *intentions* commit their subjects to facts in the previous couple of paragraphs.

c. Knowledge-Expression (*The Knowledge Content problem again*)

Being committed to the obtaining of a fact seemed important for being able to draw inferences, and I have suggested that both intention and belief commit their subjects to the obtaining of facts. What about the capacity to *express* what one knows? In (a) we followed a train of thought which seemed to threaten *NPI*'s claim that instantiating the base phenomenon is enough on its own to count as knowing: knowing requires being able to express what one knows, but expressing what one knows requires expressing an internal condition which *represents the fact known*. And intentions do not represent facts.

This is a new version of the *Knowledge Content* problem. The earlier worry was the sense of a conflict between the idea that practical knowledge is *propositional* and the idea that it is *constituted by intention*, a non-propositional attitude. I responded to the earlier worry by trying to drive a wedge between the issue of the *knowledge's* content and the issue of the content of its *constituting attitude*. But now we have seen a reason to worry that this there is no space into which this wedge can be driven: if I am to *know* that the pangolin is endangered I must be able to express this, but now expressing knowledge seems to require expressing a *representation of* that fact – something which is isomorphic to the fact I know but unlike this fact, 'inside' me. I needed a *belief* that the pangolin is endangered in order to express my knowledge of the fact. But my intention to type *does not* represent the fact that I am typing. So don't I, after all, need a *belief* that I am typing to express my knowledge of *this* fact?

The first thing to say is that the idea of *expressing one's knowledge* as requiring some inner facsimile of the fact which one knows, which is – in being expressed – somehow ejected into the world, is a metaphor, and it is unclear what it is a metaphor *for*. We should be suspicious of this metaphor at least for the reason that expressing an attitude, or a bit of knowledge, does not leave the subject *without* the feature expressed. Whatever 'expressing' an attitude or a bit of knowledge involves, it can't be like – say – expressing the pip from a lemon by squeezing it. For the lemon pip goes from being *inside* the lemon to *outside* it, and so *not* inside it; whereas when we express our attitudes and bits of knowledge, they remain with us.

If the idea that expressing knowledge requires expressing an isomorphically-shaped attitude is encouraged by a model of expression like the lemon-pip model, then we needn't take it very seriously, for the lemon-pip model is problematic for the reason just given. The question for *NPI* is whether a better model for understanding knowledge-expression will *also* require us to think of the knowledge as expressed by expressing an attitude whose content is isomorphic to the fact known.

A better model of knowledge-expression is made available by our language of predication introduced in §1(e). I spoke there of theoretical predication – which is internal to theoretical knowledge by being internal to its constituting attitude (belief) – and of practical predication which is internal to practical knowledge, *without* being internal to its constituting attitude (intention). But another kind of predicating a person can do is *linguistic*. I linguistically predicate *being endangered* of *the pangolin* when I assert: “The pangolin is endangered”.

With this in mind, an alternative model of knowledge-expression becomes available. The idea is that one expresses one’s knowledge when one’s linguistic predication captures the predication internal to one’s knowledge. In the theoretical case this would mean linguistically predicating *F* of *a* and thereby capturing what one knows in virtue of *theoretically* predicating *F* of *a* in one’s judgment. My knowing that the pangolin is endangered involves – we said above – my theoretically predicating *being endangered* of *the pangolin* and thus securing a belief-content which is a unification of representations of the pangolin and of being endangered into a propositionally-shaped representation. This theoretical predication constitutes knowledge *iff* theoretical knowledge’s knowledge-elevators are in place. My knowledge is *expressed* when I *linguistically* predicate isomorphic linguistic constituents into an assertion: “The pangolin is endangered”. Call this the *predication-mirroring* model of knowledge-expression.

We can tell an analogous story for practical knowledge. I can express my practical knowledge by *linguistically* predicating *ϕ-ing* of myself, where I thereby capture what in *this* case is known in virtue of my *practically* predicating *ϕ-ing* of *myself*. My knowing that *I am typing* involves my practically predicating *typing* of *myself*. This is just what I do in executing my intention to type and so precipitating *the fact that I am typing*, which is the object of my knowledge. On the predication-mirroring model, I can express my practical knowledge by *linguistically* predicating *typing* of *myself* and thus capturing the practical predication internal to my knowledge that I am typing.

Both theoretical and practical knowledge can be expressed by linguistically predicating in a way that captures the predication internal to one’s knowledge. In theoretical knowledge, this predication will also be internal to the knowledge’s constituting attitude – *belief* – because theoretical predication is the unification of thought-contents into a belief-content. But in practical knowledge, this predication will *not* also be internal to the knowledge’s constituting attitude – *intention* – because practical predication, as explained in §1(e), occurs in the *execution* of intention, not in its formation.

My opponent here might demand that I explain *how* I could be capable of linguistically predicating *typing* of *myself* without having theoretically (“internally”) predicated *typing* of *myself*. But part of what I want to suggest is that the feeling of a demand here stems from a suspect metaphor like the lemon-pip model of knowledge-expression. Once we have rejected this metaphor, there is no reason left to think that one can express a piece of knowledge only by expressing an attitude isomorphic to it, and we are freed up to adopt the predication-mirroring model of knowledge-expression. At the very least, the burden is on *NPI*’s opponent to explain why expressing a piece of knowledge is possible only by expressing an isomorphically shaped psychological attitude.

d. A Red Herring?

I will consider a final complaint before concluding this chapter. The complaint is that all this detail about practical predication is a red herring. For can't someone practically predicate something – snoring, for example – of oneself *without* having the expressive and inferential capacities which we have been treating as essentially linked with knowing? Can't one even practically predicate *the content of one's intention* of oneself without possessing these capacities – as when I water the tomatoes by tripping and sloshing, and don't realise what I am doing?

The answers to both questions are yes. And in light of this my opponent is likely to suggest that we *still* need to think of my knowledge that I am typing as belief-constituted: only by thinking of me as *believing* that I am practically predicating *typing* of myself (or just: that I am typing) can we understand how I am mentally engaged with the fact that I am typing in such a way as to underwrite the relevant expressive and inferential capacities (and so the higher-level epistemic capacities familiar from §3).

I have already given a sense of how I want to respond to this kind of objection, for it is a new slant on the one I considered at the end of (c); the objection that I have not explained *how* I manage to be capable of linguistically predicating *typing* of *myself* without having theoretically ('internally') predicated *typing* of *myself*. Here I will only add some comments to the earlier discussion, and be clear about the relevance of the non-reductive aspect of *NPI* to my approach to this kind of objection.

We need to be clear that practical predication is – because execution is – to be understood *itself* in relation to practical knowledge. This is because practical predication is another name for intention-execution, and we have already said that intention-execution is one facet of the base phenomenon, of which practical knowledge is another facet. So far we have been using the idea of intention-execution to explain practical knowledge, but that was solely for dialectical reasons – it was *practical knowledge* that we have been trying to understand. Here, though, we want to understand intention-execution or – as we are now also calling it – practical predication. And to do this we need to bear in mind its relationship to *practical knowledge*.

Attention to the non-reductive character of *NPI* (at least as I am developing it here), shows that I *do not* practically predicate *watering the tomatoes* of myself in the bad case. At least, I do not do so in the sense that is required for – and internal to – practical knowledge. The sense in which I practically predicate *watering the tomatoes* of myself in the bad case is analogous to the sense in which Ellisif theoretically predicates *being endangered* of *the pangolin* when she *imagines* that the pangolin is endangered. In so doing she does not predicate *being endangered* of *the pangolin* in a way apt to constitute *knowledge* that the pangolin is endangered; we can see this by observing that *even if* theoretical knowledge's knowledge-elevators are in place, Ellisif – who merely imagines that the pangolin is endangered – does not *know* that it is.

There is more to say, of course, about what is involved in the theoretical predicating I do when I *judge* that the pangolin is endangered, how this kind of theoretical predication differs from that which Ellisif does in merely *imagining* that the pangolin is endangered. Whatever more there is to say about this is analogous to what more needs to be said about the practical predication of *watering the tomatoes* which I do of myself in *executing* my intention, as opposed to the practical predication of *watering the tomatoes* which I do of

myself in merely satisfying it (when I trip and slosh). I don't have room to pursue either case here.

The take-away point is that whatever more needs to be said will not be sayable without making mention of practical knowledge itself. This is the upshot of my non-reductive conception of the relations between the base phenomenon's three facets, and the concepts which pick them each out. Given the centrality of the non-reductive characterisation of the relationship between *practical knowledge*, *intentional action* and *intention-execution* to *NPI*, my opponent cannot expect me to describe the kind of practical predication which is sufficient for practical knowledge, *independently* of talking about practical knowledge itself. This commitment to non-reductionism is likely to put someone with a certain philosophical mind-set off *NPI*. But it is a central and whole-hearted part of the picture I have been defending: a feature of the view which I am happy to lean on, and not a consequence to shy away from.

e. Summing up

I have tried to explain in this section why I think that someone can possess the capacities associated with knowing ((*a*)-(d')) I virtue of instantiating the base phenomenon, *rather than* in virtue of instantiating the Standard Metaphysics of Knowledge, and I have done so by trying to first uncover, and then undermine, the motivations for thinking that knowledge (*per se*) *must* be constituted by belief. What I have *not* done is deny that someone who instantiates the base phenomenon – and so has practical knowledge – will have a belief about what they are doing. Perhaps we typically have beliefs about what we are intentionally doing. Perhaps we *always* do. Both possibilities are consistent with what I have argued here – that practical knowledge is not *constituted by* belief. It is worth making some very quick comments about how if we *do* have beliefs about what we do intentionally as well as having intention-constituted practical knowledge, the two phenomena relate.¹⁷

First, according to *NPI*, any beliefs we have about our intentional actions will be dependent on practical knowledge for their truth. My belief that I am intentionally typing will be true only if I am intentionally typing, but I will be intentionally typing only if I have practical knowledge that I am. And second, any beliefs we have about our intentional actions will – plausibly – be dependent on our practical knowledge for their *justification*. Here I borrow a thought from Johannes Roessler, who, in a discussion of Hampshire's account of practical knowledge, remarks that “it can hardly be suspect to believe that *p* when one knows that *p* and has good reason to think one knows that *p*”.¹⁸

What are my good reasons to think I know that I am typing? Perhaps they are that I know that I have the relevant epistemic capacities in relation to the fact that I am typing. Sure, my knowing that I possess them is dependent on my knowing that I am typing, but this needn't rule out that I *do* know that I possess them. Compare: my knowing that the sun is out is dependent on knowing that I am not dreaming, but this does not rule out that I *do* know that the sun is out; at least if we think that global scepticism must be false, and this is an assumption of our commonsense-human methodology.

¹⁷ See also Anscombe's comments on “two knowledges [in any operation, of which] one is practical, the other speculative.” (Anscombe, *Intention*, 88–89.)

¹⁸ Roessler, “The Epistemic Role of Intentions,” 53.

Another option would be to reject the requirement in Roessler's second conjunct, and say simply that 'it can hardly be suspect to believe that p when one knows that p '. This seems hard to deny. In either case, having *practical* knowledge that I am typing, which is constituted by intention and not by belief, can act as a ground for *theoretical* knowledge that p . Because I instantiate the base phenomenon, any belief to the effect that I am typing will be elevated to (theoretical) knowledge-status: it will be *true*, and it will be held in favourable epistemic circumstances, these being just the circumstances in which I instantiate the base phenomenon; in which I have *practical* knowledge.

5. Conclusion

In this chapter I have elucidated and defended *NPI*. *NPI* took over from *NCI* the capacity to meet the *Apparent Features*, *Special Relationship* and *Structure-Sensitivity* constraints, doing all this in virtue of thinking of practical knowledge as constituted by intention, understood *non-cognitively*. *NPI* adds in that intentions are to be understood *non-propositionally*, and so trivially meets the *Intention-Content* constraint. I have tried to explain why *NPI* also meets the *Knowledge-Content* constraint and the *Why Knowledge?* constraint, despite first appearances. We can allow that it meets the former if we are happy to allow that someone can have knowledge of a fact without representing that fact to themselves, and I explained how I see this working. And we can allow that it meets the latter if we think of knowing that p as involving being mentally engaged with the fact that p in such a way as to underwrite certain epistemic capacities; those detailed in §3.

NPI's capacity to meet both constraints relies on a conception of knowing *per se* as being mentally engaged with a fact in such a way that one possesses the capacities (a)-(c) & (d'') from §3, which I said in turn relies on the possession of the capacities to *express* and to *make inferences from* one's knowledge. Perhaps knowing that p *explains* a person's having this collection of capacities in relation to a fact, or perhaps it *just is* their having this collection of capacities in relation to a fact. I have remained neutral on this question throughout this section. Whichever way we develop the idea, my thought in this Chapter has been that although having a true belief in favourable epistemic circumstances is *one* way in which someone might possess the capacities which go along with knowing, it is not the *only* way, and so that instantiating the Standard Metaphysics of Knowledge is not the only way of knowing. When *what is known* is the fact that one is acting or will act, and the actions in question are intentional, a person will possess these capacities in virtue of having an intention, on the condition that it is executed; in virtue of instantiating what I have been calling the base phenomenon. And this was the key claim of *NPI*.

Conclusion

By way of concluding, I want to re-cap the considerations which led me to develop and defend *NPI*, to flag up just a few of the loose ends I have left, highlighting issues which would particularly benefit from further investigation.

a. The lay of the land

Chapter Six aimed to remove the most obvious obstacles to accepting the key claim of *NPI*, that practical knowledge is constituted by intention, understood non-propositionally (and so non-cognitively), given that it is executed. It was of central importance to *NPI*'s capacity to meet all six of our constraints that it did *not* fit the Standard Metaphysics of Knowledge (true belief held in epistemically favourable circumstances). So along with *NPI* as an account of *practical* knowledge, came a view about knowledge *per se*, as a structure which can be manifested in two different ways: theoretically by *believing* or practically by *intending*. Both intention and belief are equally potentially knowledge-constituting states, because both can put one in the possession of a set of capacities which commonsense links with knowing. There is simply no reason to deny that someone with all these capacities is a knower, or so I have suggested.

But for the reader unconvinced by *NPI*, it is worth recalling how we came to it, for if she wants to reject it, she will have to reject one or more of my premises. The first (Chapter One) was that a good account of practical knowledge must meet the *Apparent Features* and *Special Relationship* constraints. The second (Chapter Two) was that two main alternatives to *Intentionalism* (the consciousness-based and inferentialist accounts) fail to meet these constraints, and do so in a way which motivates considering *Intentionalism*. The third premise (Chapters Three and Four) was that a certain kind of *Intentionalism* – *CI* – wouldn't fit the bill. The fourth (Chapter Five) was that a different kind of *Intentionalism* – *NCI*, which I attributed to Anscombe – met the *Apparent Features* and *Special Relationship* constraints very well, and the new *Structure-Sensitivity* constraint; and it did this *because* of rejecting Cognitivism about Intention. The fifth (Chapter Five) was that *if* this version of *Intentionalism* could not also meet the *Why Knowledge?* constraint and meet the *Intention-Content* constraint consistently with the new *Knowledge-Content* constraint, then it would have to be rejected. The Sixth (Chapter Six) was that a development of *NCI* – *NPI* – could both meet the *Intention-Content* constraint (trivially), and the *Why Knowledge?* and *Knowledge-Content* constraints. I concluded that *NPI* is worth taking seriously.

Two extra assumptions helped the argument along. One was that we ought to adopt what I have been referring to as a commonsense-human methodology in our investigations into the notion of practical knowledge, because it is a commonsense-human concept. And the other was that no account of practical knowledge which sees it as grounded somehow in perception, including proprioception, was likely to accommodate its seeming features. I did not consider any such accounts, a strategy which has since been vindicated by the considerations in Chapter Five (§2(b)) which led to the *Structure-Sensitivity* constraint: if practical knowledge must take in *future* as well as *present* intentional action, then it cannot be grounded in perception, since perception cannot be of what is future. If there are any good perception-based accounts of a special kind of awareness we have of what we are up to, it is not an account of exactly the phenomenon we have

been calling ‘practical knowledge’ (I attached a similar thought, recall, to ‘consciousness-based’ accounts in Chapter Two).

I hope the above will serve as an aid to the reader who is not convinced by *NPI*; as a way of pinpointing where the argument of the Thesis seems to have gone astray. I will close by considering some loose ends, suggesting some avenues for further investigation.

b. Loose ends

One of the main ideas of *NPI* is that in practical knowledge one knows that one is φ -ing without *representing oneself as* φ -ing (Ch. 6, §1(d); §4). This claim has some resonance with ideas which can be found in Anscombe and in Lewis. In ‘The First Person’, Anscombe denies that ‘I’ is a referential expression, suggesting instead that we use ‘I’ to express “unmediated agent-or-patient conceptions of actions, happenings or states”.¹ And in ‘Attitudes *De Dicto* and *De Se*’, Lewis argues that *de se* attitudes are not attitudes to propositions *about* oneself, but are instead self-ascriptions of properties. But things are complicated by the fact that Anscombe’s and Lewis’ remarks concern self-ascription rather generally, whereas my remarks about *NPI* are specific to practical knowledge, in virtue of its being constituted by *intention*; a non-propositional attitude. So it would be good to think more about how the non-representational element of *NPI* might relate to these ideas of Anscombe’s and Lewis’, and additionally about how their ideas relate to one another.

NPI is wholeheartedly non-reductive and for that reason sits much better with a non-reductive conception of theoretical knowledge. But there are various ways in which one might be a non-reductionist about theoretical knowledge. Consider just two ways of developing the ‘Knowledge-First’ idea. On one, all knowledge is belief-constituted, albeit not *reducible* to belief: knowing is believing *par excellence*. On another, knowing is being in a primitive relation to a fact.² If we assume that both of these are supposed to be accounts of knowledge *per se*, then only the second is consistent with *NPI*. Whether or not, in the end, *NPI* should be accepted will depend *inter alia* on whether it is consistent with an acceptable account of knowledge *per se*. I started to make the case, in Chapter Six (§§3&4) that it is, but a full defence of *NPI* would require a more thorough investigation of how to think about knowledge *per se*.

If my response to the *Why Knowledge?* problem on *NPI*’s behalf works, then an interesting possibility is opened up for thinking about psychological self-knowledge. The strategy in Chapter Six (§§3&4) was to show that instantiating the base phenomenon put one in possession of a collection of distinctively epistemic capacities, independently of one’s instantiating the Standard Metaphysics of Knowledge. But notice that a parallel argument might be available in relation to psychological self-knowledge. My belief that the pangolin is endangered seems to go hand-in-hand with my being able to provide accurate testimony about whether I believe it, to use the fact that I believe it as a reason, to stop enquiring about whether I believe it, and to give someone else a reason for believing that I believe it. In light of this, why not say that I *know* that I believe that the pangolin is endangered? This would be a kind of expressivism about self-knowledge, but not an expressivism to worry about. For it is not one which *contrasts* my ability to express my belief with my capacity to know that I have it; on the contrary, the suggestion is that my

¹ Anscombe, “The First Person,” 159.

² A discussion of these options can be found in Williamson, *Knowledge and Its Limits*, Ch. 1, §5.

capacity to express my first-order belief *that the pangolin is endangered* is central to my possession of the epistemic capacities which go along with knowing that I believe that the pangolin is endangered. This would be an expressivism with epistemic credentials, so is worth exploring.

What goes for belief seems also to go for sensations, emotions and other attitudes such as desire, hope, fear and intention. The latter is worth commenting on explicitly. Anscombe, I suggested (Ch. 5, §2(c)), denied that knowing *what* one is doing is in any way grounded in knowing *why* one is doing it: we don't have practical knowledge by *knowing* our reasons or intentions, but by *having* them. *NPI* agrees with this: practical knowledge is *constituted by* intention. But if the idea just floated is right, then typically someone who has practical knowledge will have knowledge of her intention too. For having practical knowledge entails *having* an intention, and having an intention (like having a belief) puts one in possession of the capacities relevant to knowing. I don't know that I am typing *because* I know that I intend to be typing; my intention itself constitutes – in different ways – my knowledge that I am typing *and* my knowledge that I intend to be doing so.

The concept of practical knowledge plays a complex and important role in our lives and in our understanding of the kinds of creatures we are. I have alluded to this role in places, saying rather simplistically that the objects of practical knowledge form a legally and morally relevant class. But there is more to practical knowledge's role than this. Plausibly, for example, it also helps to constitute the psychological make-up of individuals, binding together experiences into *the experience of a life*; helping to link one's future self with one's present and past selves; making the idea of logically distinguishable time-slices hard to maintain, grounding a person's temporal unity. These issues deserve further attention. It is plausible that it is in virtue of its apparent features that practical knowledge can play roles like these. If this is right, then rejecting any or all of these as *merely* apparent will have implications for the reality of the notion of a human being, and of a distinctively human kind of life. To repeat a thought from my *Introduction*, it does not seem clear that there is a coherent viewpoint from which to doubt the reality of concepts as fundamental as these to who we take ourselves to be.

Johannes Roessler ends his paper 'The Epistemic Role of Intentions' by advocating a project in which we do not try to "offer a philosophical defence of the commonsense view of practical self-knowledge, but question whether it can coherently be challenged."³ I have just been suggesting that it *cannot* coherently be challenged. But in this Thesis I have offered a 'philosophical defence' of it nonetheless.

³ Roessler, "The Epistemic Role of Intentions," 55.

Bibliography

- Alvarez, Maria. "Acting Intentionally and Acting for a Reason." *Enquiry* 52, no. 3 (2009): 293–305.
- . *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford University Press, 2010.
- Anscombe, G. E. M. *Intention*. Harvard University Press, 1957.
- . "Thought and Action in Aristotle: What Is Practical Truth?" In *New Essays on Plato and Aristotle*, edited by J. R. Bambrough, 143–58. London: Routledge and Kegan Paul, 1965.
- . "Under a Description." *Noûs* 13, no. 2 (1979): 219–33.
- . "The First Person." In *Self-Knowledge*, edited by Quassim Cassam. Oxford and New York: Oxford University Press, 1994.
- . "Practical Inference." In *Human Life, Action and Ethics: Essays by GEM Anscombe*, 109–47. Exeter: Imprint Academic, 2005.
- Armstrong, D. M. "Non-Inferential Knowledge." In *Belief, Truth and Knowledge*, 162–83. Cambridge: Cambridge University Press, 1973.
- Baier, Annette C. "Act and Intent." *Journal of Philosophy* 67, no. 19 (1970): 648–58.
- Bishop, John. "Is Agent-Causality a Conceptual Primitive?" *Synthese* 67, no. 2 (1986): 225–47.
- Boyle, Matthew. "Active Belief." *Canadian Journal of Philosophy, Supplementary Volume* 35 (2009): 119–47.
- . "'Making Up Your Mind' and the Activity of Reason." *Philosophers' Imprint* 11, no. 17 (2011).
- Boyle, Matthew, and Doug Lavin. "Goodness and Desire." In *Desire, Practical Reason and the Good: Classical and Contemporary Perspectives*, edited by Sergio Tenenbaum. New York: Oxford University Press, 2010.
- Bratman, Michael. *Intention, Plans, and Practical Reason*. London: Harvard University Press, 1987.
- Broome, John. "Normative Practical Reasoning." *Aristotelian Society Supplementary Volume* 75, no. 1 (2001): 175–93.
- Burge, Tyler. "Content Preservation." *The Philosophical Review* 102, no. 4 (1993): 457–88.
- Castañeda, H-N. "Intentions and Intending." *American Philosophical Quarterly* 9, no. 2 (1972): 139–49.
- Chisholm, Roderick M. "Freedom and Action." In *Freedom and Determinism*, edited by K Lehrer. New York: Random House, 1966.
- . "The Structure of Intention." *Journal of Philosophy* 67, no. 19 (1970): 633–47.
- Clark, Philip. "The Action as Conclusion." *Canadian Journal of Philosophy* 31, no. 4 (2001): 481–505.
- Coope, Ursula. "Aristotle on Action." *Aristotelian Society Supplementary Volume* 81, no. 1 (2007): 109–211.
- Crane, Tim. "The Intentional Structure of Consciousness." In *Consciousness: New Philosophical Perspectives*, edited by A Jokic and Q Smith. Oxford and New York: Oxford University Press, 2003.
- . *The Objects of Thought*. Oxford: Oxford University Press, 2013.
- Dancy, Jonathan. *Practical Reality*. Oxford: Oxford University Press, 2000.
- . "On Knowing What One Is Doing." *Philosophical Studies* 121, no. 3 (2004): 239–47.

- . “Action, Content and Inference.” In *Wittgenstein and Analytic Philosophy: Essays for PMS Hacker*, edited by PMS Hacker, H-J Glock, and J Hyman, 278–98. Oxford: Oxford University Press, 2009.
- Danto, Arthur. *Analytical Philosophy of Action*. Cambridge: Cambridge University Press, 1973.
- Davidson, Donald. “Actions, Reasons, and Causes.” In *Essays on Actions and Events*, 2nd ed. Oxford: Clarendon Press, 2001.
- . “Agency.” In *Essays on Actions and Events*. Oxford: Clarendon Press, 2001.
- . “Intending.” In *Essays on Actions and Events*. Oxford: Clarendon Press, 2001.
- . “Problems in the Explanation of Action.” In *Problems of Rationality*. Oxford: Oxford University Press, 2004.
- Davis, Wayne. “A Causal Theory of Intending.” *American Philosophical Quarterly* 21, no. 1 (1984): 43–54.
- Dokic, Jérôme. “The Sense of Ownership: An Analogy between Sensation and Action.” In *Agency and Self-Awareness*, edited by Johannes Roessler and Naomi Eilan, 321–44. Oxford and New York: Oxford University Press, 2003.
- Donnellan, Keith S. “Knowing What I Am Doing.” *Journal of Philosophy* 60, no. 14 (1963): 401–9.
- Dorr, Cian. “Non-Cognitivism and Wishful Thinking.” *Nous* 36 (2002): 97–103.
- Falvey, Kevin. “Knowledge in Intention.” *Philosophical Studies* 99, no. 1 (2000): 21–44.
- Fodor, Jerry. “Propositional Attitudes.” *The Monist* 61 (1978): 501–23.
- Goldman, Alvin I. *A Theory of Human Action*. Princeton University Press, 1970.
- Grice, H. P. “Intention and Uncertainty.” *Proceedings of the British Academy* 5 (1971): 263–79.
- Grünbaum, Thor. “Perception and Non-Inferential Knowledge of Action.” *Philosophical Explorations* 14, no. 2 (2011): 153–67.
- Haddock, Adrian. “At One with Our Actions, but at Two with Our Bodies.” *Philosophical Explorations* 8, no. 2 (2005): 158–72.
- . “The Knowledge That a Man Has of His Intentional Actions.” In *Essays on Anscombe’s Intention*, 147–69. Harvard University Press, 2011.
- Hamilton, Andy. “Intention and the Authority of Avowals.” *Philosophical Explorations* 11, no. 1 (2008): 23–37.
- Hampshire, Stuart. “Reply to Walsh on Thought and Action.” *The Journal of Philosophy* 60, no. 14 (1963): 410–24.
- . *Freedom of the Individual*. New York: Harper and Row, 1965.
- Hampshire, Stuart, and H. L. A. Hart. “Decision, Intention and Certainty.” *Mind* LXVII, no. 265 (1958): 1–12.
- Harman, Gilbert. “Practical Reasoning.” *The Review of Metaphysics* 29, no. 3 (1976).
- Hornsby, Jennifer, and Jason Stanley. “Semantic Knowledge and Practical Knowledge.” *Aristotelian Society Supplementary Volume* 79, no. 1 (2005): 107–30.
- Humberstone, Lloyd. “Direction of Fit.” *Mind* 101, no. 401 (1992): 59–83.
- Hursthouse, Rosalind. “Intention.” *Logic, Cause and Action: Essays in Honour of Elizabeth Anscombe (Royal Institute of Philosophy Supplement)* 46 (2000): 83–105.
- Hyman, John. *Action, Knowledge and Will*. Oxford: Oxford University Press, 2015.
- . “How Knowledge Works.” *The Philosophical Quarterly* 49, no. 197 (1999): 433–51.
- Kenny, Anthony. *Will, Freedom and Power*. UK: Basil Blackwell, 1975.
- Langton, Rae. “Intention as Faith.” In *Agency and Action*, edited by Helen Steward and John Hyman, 243–58. Cambridge: Cambridge University Press, 2004.
- Lewis, David. “Attitudes de Dicto and de Se.” *Philosophical Review* 88, no. 4 (1979): 513–43.
- . “Elusive Knowledge.” *Australasian Journal of Philosophy* 74, no. 4 (1996): 549–67.

- Marcel, Anthony. "The Sense of Agency: Awareness and Ownership of Action." In *Agency and Self-Awareness*, edited by Johannes Roessler and Naomi Eilan, 48–93. Oxford and New York: Oxford University Press, 2003.
- McDowell, John. "What Is the Content of an Intention in Action?" *Ratio* 23, no. 4 (2010): 415–32.
- Montague, Michelle. "Against Propositionalism." *Noûs* 41, no. 3 (2007): 503–18.
- Moran, Richard. "Making Up Your Mind: Self-Interpretation and Self-Constitution." *Ratio* 1, no. 2 (1988): 135–51.
- . *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton University Press, 2001.
- . "Anscombe on 'Practical Knowledge.'" In *Agency and Action (Royal Institute of Philosophy Suppl. 55)*, 55:43–68. Cambridge University Press, 2004.
- . "Self-Knowledge, 'Transparency', and the Forms of Activity." In *Introspection and Consciousness*, edited by Declan Smithies and Daniel Stoljar, 211–36. Oxford: Oxford University Press, 2012.
- Mothersill, Mary. "Anscombe's Account of the Practical Syllogism." *Philosophical Review* 71, no. 4 (1962): 448–61.
- Müller, Anselm. "Radical Subjectivity: Morality Versus Utilitarianism." *Ratio* 19 (1977): 115–32.
- . "Backwards-Looking Rationality and the Unity of Practical Reason." In *Essays on Anscombe's Intention*, edited by A. Ford, J. Hornsby, and F. Stoutland, 242–69. London: Harvard University Press, 2011.
- Newstead, Anne. "Interpreting Anscombe's Intention §§32FF." *Journal of Philosophical Research* 34 (2009): 157–76.
- O'Brien, Lucy. "On Knowing One's Own Actions." In *Agency and Self-Awareness*, 358–82. Clarendon Press, 2003.
- . "Self-Knowledge, Agency, and Force." *Philosophy and Phenomenological Research* 71, no. 3 (2005): 580–601.
- . *Self-Knowing Agents*. New York: Oxford University Press, 2007.
- Paul, Sarah K. "How We Know What We're Doing." *Philosophers' Imprint* 9, no. 11 (2009): 1–24.
- . "How We Know What We Intend." *Philosophical Studies* 161, no. 2 (2012): 327–46.
- Peacocke, Christopher. *Holistic Explanation: Action, Space, Interpretation*. Clarendon Press, 1979.
- . "Our Entitlement to Self-Knowledge: Entitlement, Self-Knowledge, and Conceptual Redeployment." *Proceedings of the Aristotelian Society* 96 (1996): 117–58.
- . *Being Known*. Oxford: OUP, 1999.
- . "Three Principles of Rationalism." *European Journal of Philosophy* 10 (2002): 375–97.
- . "Action: Awareness, Ownership, and Knowledge." In *Agency and Self-Awareness*, edited by Naomi Eilan and Johannes Roessler. Oxford: Clarendon Press, 2003.
- . "Mental Action and Self-Awareness (I)." In *Contemporary Debates in the Philosophy of Mind*, edited by Lucy O'Brien and Matthew Soteriou, 358–76. Oxford University Press, 2009.
- . "Mental Action and Self-Awareness (II): Epistemology." In *Mental Actions*. Oxford University Press, 2009.
- Pickard, Hanna. "Knowledge of Action Without Observation." *Proceedings of the Aristotelian Society* 104, no. 1 (2004): 205–30.

- Proust, Joëlle. "Perceiving Intentions." In *Agency and Self-Awareness*, edited by Johannes Roessler and Naomi Eilan, 296–320. Oxford and New York: Oxford University Press, 2003.
- Radford, Andrew. *Transformational Grammar: A First Course*. Cambridge: Cambridge University Press, 1988.
- Raz, Joseph. "Agency, Reason, and the Good." In *Engaging Reason: On the Theory of Value and Action*. Oxford: OUP, 2002.
- Roessler, Johannes. "Intentional Action and Self-Awareness." In *Agency and Self-Awareness: Issues in Philosophy and Psychology*, edited by Johannes Roessler and Naomi Eilan, 383–404. Consciousness and Self-Consciousness. Oxford and New York: Oxford University Press, 2003.
- . "The Epistemic Role of Intentions." *Proceedings of the Aristotelian Society* CXIII (2013): 41–56.
- Roessler, Johannes, and Naomi Eilan. "Agency and Self-Awareness: Mechanisms and Epistemology." In *Agency and Self-Awareness*, 1–47. Oxford and New York: Oxford University Press, 2003.
- Ryle, Gilbert. *The Concept of Mind*. London: Penguin Books, 2000.
- Schwenkler, John. "Perception and Practical Knowledge." *Philosophical Explorations* 14, no. 2 (2011): 137–52.
- Searle, John R. "The Intentionality of Intention and Action." *Enquiry* 22, no. 1–4 (1979): 253–80.
- Setiya, Kieran. "Explaining Action." *The Philosophical Review* 112, no. 3 (2003): 339–93.
- . "Practical Knowledge." *Ethics* 118, no. 3 (2008): 388–409.
- . "Practical Knowledge Revisited." *Ethics* 120, no. 1 (2009): 128–37.
- . "Knowledge of Intention." In *Essays on Anscombe's Intention*, edited by A. Ford, J. Hornsby, and F. Stoutland, 170–97. USA: Harvard University Press, 2011.
- . "Knowing How." *Proceedings of the Aristotelian Society* CXII, no. 3 (2012).
- Shah, Nishi, and Matthew Silverstein. "Reasoning in Stages." *Ethics* 124, no. 1 (2013): 101–13.
- Sosa, Ernest. "Intellectual Virtue in Perspective." In *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press, 1991.
- Stanley, Jason, and Timothy Williamson. "Knowing How." *Journal of Philosophy* 98, no. 8 (2001): 411–44.
- Taylor, Richard. *Action and Purpose*. Englewood Cliffs, NJ: Prentice-Hall, 1966.
- Thompson, Michael. "Naïve Action Theory." In *Life and Action*. Harvard University Press, 2008.
- . "Anscombe's Intention and Practical Knowledge." In *Essays on Anscombe's Intention*. Harvard University Press, 2011.
- Velleman, David. *Practical Reflection*. Princeton: Princeton University Press, 1989.
- . *The Possibility of Practical Reason*. Oxford University Press, 2000.
- . "Precis of 'The Possibility of Practical Reason.'" *Philosophical Studies* 121 (2004): 225–38.
- . "Replies to Discussion on the Possibility of Practical Reason." *Philosophical Studies* 121, no. 3 (2005).
- Vermazen, Bruce. "Objects of Intention." *Philosophical Studies* 71, no. 3 (1993): 223–65.
- Vogler, Candace. "Anscombe on Practical Inference." In *Varieties of Practical Reasoning*, 437–64. XXXX: MIT Press, 2001.
- Wallace, John. "Propositional Attitudes and Identity." *The Journal of Philosophy* LXVI, no. 6 (1969): 145–52.
- White, Alan R. *The Nature of Knowledge*. Maryland Studies in Public Philosophy. USA: Rowman & Littlefield, 1982.

- Williamson, Timothy. *Knowledge and Its Limits*. Oxford: Oxford University Press, 2000.
- Wittgenstein, Ludwig. *On Certainty*. Translated by G. E. M. Anscombe and G. H. von Wright. Oxford: Blackwell Publishers, 1975.
- . *Philosophical Investigations*. Translated by G. E. M. Anscombe. 3rd ed. Oxford: Blackwell Publishing, 2001.
- Woozley, A. D. “Knowing and Not Knowing.” *Proceedings of the Aristotelian Society* 53 (March 1952): 151–72.