

DEPARTMENT OF PHILOSOPHY  
UTRECHT UNIVERSITY

**$\theta$**  KNOWING AND DOING  **$\pi$**

**Niels VAN MILTENBURG**

*Student number: 0483931*

First Supervisor: Dr. T. MÜLLER  
Second Supervisor: Dr. A. KALIS  
Third Supervisor: Prof. Dr. J.M. VAN OPHUIJSEN

*This Master's Thesis is submitted as the final part of the research master Philosophy*



*Es ist nichts schrecklicher als eine tatige Unwissenheit.*  
— Johann Wolfgang von Goethe

# Contents

<b>Introduction</b>	<b>1</b>
<b>1 Anscombe on Practical Knowledge</b>	<b>6</b>
1.1 Setting the stage . . . . .	7
1.2 Knowledge without observation of your intentional actions . . . . .	8
1.3 Knowledge in intention . . . . .	9
1.4 Measure, not measured . . . . .	10
1.5 Inferred by practical reasoning . . . . .	12
1.6 Know how . . . . .	15
1.7 The cause of what it understands . . . . .	17
1.8 Practical knowledge . . . . .	19
<b>2 Writing Without Ink</b>	<b>20</b>
2.1 Ten carbon copies . . . . .	20
2.2 Understanding practical knowledge . . . . .	21
2.3 Fitting practical knowledge . . . . .	25
2.4 Action and aspect . . . . .	30
2.5 Advanced carbon copy cases . . . . .	32
<b>3 Implications for epistemology</b>	<b>36</b>
3.1 Non-practical knowledge in action . . . . .	37
3.2 Knowledge acquisition . . . . .	41
3.3 Causal knowledge and the ability to act . . . . .	42
3.4 Action in contemporary epistemology . . . . .	47
<b>4 Implications for Action Theory</b>	<b>50</b>
4.1 Neo-Wittgensteinian action theory . . . . .	51
4.2 Causal action theory . . . . .	53
4.3 Deviant causal chains . . . . .	55
4.4 Non-causal action theory . . . . .	58
4.4.1 Reasons and circumstances . . . . .	59
4.4.2 Decisions and concurrent intentions . . . . .	60
4.5 Practical knowledge and causation . . . . .	62
<b>Conclusion</b>	<b>68</b>
<b>References</b>	<b>72</b>

# Introduction

**A**BOUT to pour yourself a nice glass of Château de Chasseloir, you realize that you forgot to fetch a corkscrew. You walk to the kitchen, thinking about the hard day of work now left behind. Upon arriving in the kitchen you suddenly stop. ‘What am I doing here?’ It is only after a short moment of deliberation that you remember you were here for the corkscrew, and you’ll proceed by fetching it.

Such a scenario will be familiar to all. And while these moments are rare, they are also revealing.<sup>1</sup> The rarity of such moments indicates that we know what we are doing most of the time. The fact that we stop to catch our feet when we forget what we are up to, suggests that we might be unable to perform actions without such knowledge. In a short but extremely dense book Elisabeth Anscombe reflects on both the nature and the role of this knowledge of our own actions, or, as she calls it ‘practical knowledge’. *Intention* is widely regarded as one of the most important books on the philosophical explanation of action. As important as they find the book, most philosophers also regard it as fundamentally unsound. For some, however, “being incompatible with Anscombe is a lot like being incompatible with the facts.”<sup>2</sup>

I am surely of the latter persuasion. To me, Anscombe’s teaching on practical knowledge and intentional action seem both important and true. Nevertheless, a lot has happened in philosophy since *Intention* appeared in the late fifties. For one thing the philosophy of mind, of which action theory is a part, has matured into a substantial discipline with its own customs and terminology. Therefore, *Intention*, full of Wittgensteinian jargon, can easily look somewhat old fashioned or even outdated. However, a renewed interest in practical knowledge has recently emerged. David Velleman [1989] — with Setiya [2008] and Paul [forthcoming] in his wake — has tried to adapt the concept of practical knowledge to suit his contemporary views on the nature of action. As I will argue, much of the truth and importance of Anscombe’s notion is lost with this adaption.

---

<sup>1</sup>I’m borrowing Velleman’s [1989] phrase here.

<sup>2</sup>This phrase is from [Gibbons 2010].

Therefore the objective of this thesis is to reintroduce truly Anscombean territory in the contemporary landscape. I will try to reformulate Anscombe's doctrine of practical knowledge in contemporary terms and show that it is relevant to current philosophical debates about both knowledge and action.

Before we rush off into philosophical discussions about the merits and flaws of an Anscombean account of practical knowledge, it will be helpful to get some basic grasp of intentional actions themselves. We call my opening the door *intentional* if I do it *on purpose* — to let my cat in for instance. Had I *accidentally* opened the door by leaning against it, we would tend to say that I opened it unintentionally.<sup>3</sup> So, a rough characterization of intentional actions might be that they are actions done purposefully.<sup>4</sup> Non-intentional events, such as the toppling over of a building crane, are explained by citing a *cause*: 'Why did the crane topple over?' 'Because of a sudden gust of wind'. My unintentional opening of the door is also explained by means of such causal explanation. It is my leaning against the door that caused it to open. But since intentional actions have a purpose, they can be explained differently: by *rationalization*. If you were to ask me why I (intentionally) opened the door, I could specify my *reason for* doing it, i.e. 'to let my cat in.' Giving my reason shows why opening the door was the *rational* thing to do. Thus, intentional actions can be distinguished from other events in that they are rational.<sup>5</sup> However, actions still are events: happenings in the physical world. My opening the door is constituted by the movement of my body and the movement of the door. The fact that intentional actions are material processes that nevertheless happen for a reason makes them particularly interesting and problematic for philosophers.

Anscombe takes the double life of intentional actions to be their key feature. According to her, this feature prominently reveals itself in that actions can be known in two separate ways. The first way of knowing them is by observation. We can observe what happens in the world, and see that James is reading. Some would object that we can *only* see that James holds a newspaper and that his eyes move quickly from left to right. In order to *know* that he is reading, we should know what goes on in his mind, which we cannot observe. Perhaps James is an agent for MI6 who only pretends to be reading,

---

<sup>3</sup>Anscombe speaks of intentional and unintentional actions. Nowadays, it is more common to reserve the term 'action' for intentional actions, and talk about 'unintentional behaviour.'

<sup>4</sup>This is indeed a *rough* characterization, because many philosophers do not want to call the bodily movements of my cat intentional if she herself pushes the door open in order to get in. Let alone the target directed movement of a heat-seeking missile.

<sup>5</sup>Of course, this doesn't mean that agents cannot behave irrationally. It means that it makes sense to speak of intentional actions as being rational. This contrasts action with natural events in that we cannot say of an earthquake that it is rational nor that it is irrational.

while he is in fact spying on someone. Anscombe could reply to such an objection that it only shows that observation is fallible. But this is in no way peculiar to the observation of actions. In a shop with unusual lighting, observation can lead us to believe that a tie is green while in fact it is blue.<sup>6</sup> But that we can be wrong about the colour of objects does not show that it is impossible to come to know the colour of an object by observation. The same holds for actions. Most of the time, people are doing precisely what they appear to be doing. So in general we can find out what an agent is doing by observing what happens. But if we are unsure we can always probe deeper and ask the agent what it is that she is doing.

This gets us to the second way of knowing. Practical knowledge is exclusively available to the acting agent herself. According to Anscombe, an agent's knowledge of her own action is inseparable from her performance of it. This shows when we compare intentional with unintentional actions. As soon as we do something intentionally, we are able to answer questions about the rationale of the action. If I'm asked why I opened the door, I do not first have to observe or recognize that I am opening the door. I can immediately tell that I did it to let the cat in. But if was asked why I opened the door when I unintentionally did so by leaning against it, then I first need to realize that 'opening the door' is indeed a description of what I am currently (unintentionally) doing, before I can answer that it only happened by accident.

Anscombe believes that the knowledge an agent has of her own action is not only inseparable from it, but also constitutive of it. The corkscrew-fetching case already suggested such a constitutive role in that an agent stops acting when she forgets what she is doing. According to Anscombe, practical knowledge determines that an event is an intentional action. It gives what happens its rational structure. Therefore, understanding practical knowledge is paramount to understanding the dual character (rational and material) of intentional actions.

This formulation of the dependence of actions on practical knowledge is, of course, quite rash. I can imagine that it sounds mysterious and raises a host of questions. What is the nature of practical knowledge? How can it be constitutive of intentional action? In what respects does it differ from non-practical knowledge? How do we acquire it? etc. The answers to the majority of such questions can be found in the pages of *Intention*. However, most of them aren't in plain sight. They are buried under a thick layer of jargon and of intricate philosophical argument — sometimes arguments against positions long

---

<sup>6</sup>To borrow Sellars's [1997 [1956]] famous example.

since abandoned. The objective of chapter 1 is to do some excavation work. By scrutinizing Anscombe's own words, I hope to arrive at a clear conception of Anscombean practical knowledge.

The idea of practical knowledge might not only raise questions, it might also lead to some critical remarks. In fact, contemporary philosophers were sceptical to such an extent that their storm of criticism killed off the entire topic of practical knowledge for some decennia.<sup>7</sup> Their main issue is that we can easily fail to know what we are doing. This happens when the world does not cooperate: I extend my arm to open the door (all the while thinking that I am opening the door), but when I grasp the handle and pull it towards me I find out that it is bolted shut, and cannot be opened. It turns out that I wasn't opening the door after all, and hence I didn't know that I was opening it. According to the practical knowledge sceptic this shows that the connection between an agent's actions and her knowledge of them isn't as strong as Anscombe proposes. Therefore, it has become common to think that an agent's knowledge in action is only knowledge about her own intentions, and not about what is happening in the world. In chapter 2, I will attempt to rebut such criticism and show that we can deal with cases of unsuccessful action and still maintain that practical knowledge is knowledge about what is actually happening.

Together, chapter one and two try to establish that Anscombe is right about both the nature and the existence of practical knowledge. But as mentioned above, I do not only think that Anscombe's teachings on the matter are true, I also called them important. The final two chapters will establish this importance for the study of knowledge (epistemology) and contemporary action theory. In Chapter 3, I will argue that our ability to act knowingly allows us to understand the world in specific ways. It will be suggested that both causal as well as modal knowledge are dependent on action. This chapter will also look at the ways in which non-practical knowledge can enter into human action, and at the impact practical knowledge can have on contemporary philosophical studies of the concept of knowledge: analytic epistemology.

Chapter 4 will return to the study of the nature of actions. An influential and prominent contemporary theory of the nature of actions is Causal Action Theory. According to causal action theorists, rationalization actually is a species of causal explanation. To say that I opened the door in order to let the cat in, is to say that my desire to let the cat in *caused* my door opening movements. The difference between intentional and unin-

---

<sup>7</sup>This observation is from [Thompson, forthcoming].



tentional actions is that intentional actions are caused by reasons, whereas unintentional movements are not. But how can a reason (something that is in the mind of the agent) cause a *bodily* movement? Causal action theorists answer that reasons are also embodied, i.e. they are realized in our brains. Thus, the rational aspect of intentional action is reduced to the material aspect of being caused in a certain way. I will argue that causal action theory is, despite its popularity among the current generation of philosophers of mind, deeply flawed. I'm not the first to notice that there are problems with causal action theory, but I will give reasons that suggests it is beyond repair. Luckily, a philosophical account of actions based on practical knowledge does not suffer from the same flaws as causal action theory.

At the end of the thesis I will summarize my findings, indicate further research and offer some final thoughts in the conclusion.

# Chapter 1

## Anscombe on Practical Knowledge

*Anscombe's classic work [Intention] is the front from which all subsequent philosophical thought about agency flows.*

– Robert B. Brandom

**E**lisabeth Anscombe's monograph *Intention* is widely regarded as one of the most important treatments of intentional action. Many of the ideas which are now commonplace in the philosophy of action (e.g. the idea that an action is intentional under a description) were first expressed in it. But not all of Anscombe's teachings have found wide acceptance among contemporary philosophers of action. Among them is the doctrine of practical knowledge. This thesis will defend this doctrine (chapter 2) and show its importance for contemporary epistemology (chapter 3) and philosophy of action (chapter 4). But before I advance any arguments, a clear understanding of practical knowledge itself is imperative. Acquiring such an understanding may not be an easy task. *Intention* provides no less than six characterizations of practical knowledge in a style of writing that is "like the confection *panforte*, all fruit and nuts and no dough, very chewy and tough."<sup>1</sup> This may lead some philosophers to think that it is impossible to find *one* coherent conception of practical knowledge in Anscombe's book.<sup>2</sup> This chapter is an attempt to rebut this suggestion by proposing a univocal understanding of practical knowledge. As I will argue, Anscombe's six characterizations in fact point to six different, but cohering, aspects of the concept of practical knowledge. I will start in section 1.1 by sketching the background against which practical knowledge is introduced in *In-*

---

<sup>1</sup>As Anscombe's daughter puts it in the introduction to [Geach and Gormally 2006].

<sup>2</sup>Anselm Müller [2010], for instance, suggests that there are at least two different concepts of practical knowledge at play in Anscombe's book.

*tention*. The subsequent sections (1.2-1.7) will each be devoted to one of the six aspects of practical knowledge. The coherence between these aspects will become increasingly clear as we go along. The final section (1.8) will highlight the importance of some of the aspects of practical knowledge for the rest of the thesis.

## 1.1 Setting the stage

Intentional actions are done for *reasons*. When you see someone pressing a button on a stove, you can ask ‘Why did you press that button?’, and she may answer: ‘To turn on the stove’. You may continue your enquiry and ask ‘Why did you turn on the stove?’ and again there will be an answer that specifies a reason, e.g. ‘To make a risotto’. Each ‘why?’ question *describes* an action, and each answer brings out the reason for the action. It is in this way that the ‘why?’ question reveals the structure of *means* and *ends* contained in intentional actions: turning on the stove is a means for cooking rice, which is again a means for making a risotto.

That such a means-ends structure is typical for intentional actions becomes clear once we consider a case in which something is done unintentionally. Suppose someone accidentally presses the button on the stove because she was leaning against it, and you ask her ‘Why did you turn on the stove?’ Her reply would be something like ‘I wasn’t aware I was doing it’ and nothing like ‘To make a risotto’. In such a case of unintentional action we might say that the question ‘why?’ is not answered at all. It is rather denied application, because there is no reason or end the agent could specify.<sup>3</sup> For this reason, Anscombe defines intentional actions as “those to which a certain sense of the question why is given application” [Anscombe 1963, §5].<sup>4</sup>

But might a perfectly legitimate answer to the question ‘Why did you turn on the stove?’ not be: ‘because I was leaning against it’? Of course it is a good reply, but it is an answer that specifies a *cause* rather than a reason for acting. But Anscombe wants to find out what distinguishes intentional actions from unintentional happenings. Both have causes, but only intentional actions are done for reasons. So if we want to use the

---

<sup>3</sup>If we had described the agent’s movements differently, as ‘leaning against the stove’ for instance, we might get a clear answer, e.g., ‘to rest a bit’. So whether certain movements amount to an intentional action depends on the way in which we describe them. This is why Anscombe claims that actions are only intentional under a description. I may intentionally put a book on the table, but with the same movements unintentionally put the book on a puddle of ink. See [Anscombe 1979].

<sup>4</sup>In the rest of this thesis I will refer to sections of *Intention* with just a section number between brackets, e.g., [§21].

‘why?’ question to single out intentional actions, we need a ‘why?’ question that only accepts reasons as answers. Hence Anscombe’s phrase “*certain sense* of the question why [italics NvM]” in the definition above.

It would, however, not suffice to say that intentional actions are those action to which the question ‘why?’, in the sense that asks for reasons for acting, is given application. This definition would be circular in that the *definiendum* is presupposed in the *definiens*. To be more precise, the concept of intentional action (defined as acting for a reason) is already used to specify the sense of the question ‘why?’ that should single out the concept. What is required in order to avoid circularity is to describe the class of intentional actions “without using any notions like ‘intended’ or ‘willed’ or ‘voluntary’ and ‘involuntary’.” [§5]. Anscombe argues that we can describe this class with the help of practical knowledge. It is against this background that the first characterization of practical knowledge enters the scene.

## 1.2 Knowledge without observation of your intentional actions

Anscombe starts her quest for a non-circular isolation of intentional actions by pointing to the class of things *known without observation*. She illustrates this class by example of the knowledge we have of the position of our limbs, e.g. we can normally tell whether our left knee is bent or not, immediately and without looking. The knowledge of our own emotions, and Wittgenstein’s examples of knowing that one can continue a mathematical series without directly envisaging the exact next numbers [Wittgenstein 2001 [1953], §187], come to mind as well. This illustration has unfortunately misled many of Anscombe’s readers in thinking that practical knowledge is, or is a consequence of knowledge of the position of our limbs. Their idea is that I can deduce what I’m doing, when I recognize my own bodily movements, e.g. I recognize that my arms are going through certain movements and deduce that I am putting a golf ball.<sup>5</sup> Many who interpret Anscombe in this way, continue by dismissing practical knowledge as a pre-scientific fantasy. They argue that physiological research has taught us that have proprioceptive *senses* that register the positions of our limbs. It follows that knowledge without observation does not exist.

Such an argument is based on a misunderstanding of what the knowledge of our own

---

<sup>5</sup>For an example of this mistake see [Pickard 2004].

intentional actions consists in. When you make up your mind to cook a risotto, you already know that you are doing so when you are walking to the kitchen. If someone stops you and asks you what you are doing, you will not hesitate to reply: ‘I’m making a risotto’. But up until you arrive in the kitchen and take the rice out of the cupboard, the exact same movements could have contributed to the making of a pasta. If I had to deduce what I was doing from my own bodily movements, I would still have been in doubt whether I was making a pasta or a risotto before I took the rice or the spaghetti from the cupboard.

But if practical knowledge is quite different from the knowledge of the position of ones limbs (or of our own emotions etc.) why does Anscombe talk about these forms of non observational knowledge? Her answer is that:

[T]he class of things known without observation is of general interest to our enquiry because the class of intentional actions is a subclass of it. [§8]

So, intentional actions are known without observation, along with emotions or the position of ones limbs. However, as will become clearer in the course of this chapter, the knowledge we have of our intentional actions is quite a peculiar sort of knowledge. For one thing practical knowledge is different from other types of knowledge in that it is “knowledge in intention” [§32].

### 1.3 Knowledge in intention

When you are intentionally making a risotto, your knowledge of what you are doing is *inseparable* from your intention to do it. You cannot know what you are doing without intending to do it, just as you cannot do something intentionally without knowing that you are. Of course, we often act without consciously thinking about what we are doing, especially when habitual actions, like brushing one’s teeth, are concerned. But still, when we are asked what we are doing we can immediatly reply that we are brushing our teeth. This stands in contrast with unintentional behaviour. In the case of the unintentional turning on of the stove, the agent’s reply was: ‘I *wasn’t aware* I was doing it.’

Due to the inseparability of knowledge and intention, we might say that practical knowledge is knowledge that is “in intention” [§32]. This also explains, in part, why practical knowledge is non-observational. From a third-person perspective it can be unclear whether a bodily movement that is *seen* is intentional or not. Another issue is

that I can be sitting on my couch, reading a paper, while I'm intentionally baking a cake at the same time – I am only waiting for the dough to rise. An observer can only see me going through the movements of paper-reading (something I am doing intentionally as well), but only I know, even while I cannot see the dough rising in its bowl, that I am also baking a cake.

At this point it is important to note that Anscombe does not think that practical knowledge is knowledge *of* your intention, it is instead knowledge that is *in* intention but *of* the action. For some this can seem counter-intuitive, because it might be the case that, unbeknownst to me, my action fails. It might, for instance, be the case that I am trying to execute my intention to bake a cake, but that I unwittingly forgot to put in the yeast before I let the dough 'rise'. It is only after returning to the kitchen an hour later that I notice my mistake, by observation, and conclude that I actually wasn't baking a cake, I only thought that I was. This idea is the basis of the so-called 'two factor thesis', the thesis that one typically knows without observation what one *intends* to do, but that knowledge of what one is *actually* doing must stem from observation.<sup>6</sup> In chapter 2 I will argue that the two factor thesis is wrong, and that Anscombe's idea isn't counter-intuitive. For now my objective is merely to clarify the concept of practical knowledge. Anscombe's next characterization of it will help to show why practical knowledge in her sense must be knowledge of what we are actually doing.<sup>7</sup>

## 1.4 Measure, not measured

Anscombe contrasts practical knowledge with what we might call (following Aquinas) 'speculative knowledge'.<sup>8</sup> Speculative knowledge is the knowledge we may have of the world around us. It is the knowledge with which epistemologists are concerned, and the knowledge natural scientist wish to attain. An example of speculative knowledge is my knowledge that the hippopotamus eats about 65 kilogram of grass each night. The most prominent difference between practical and speculative knowledge is that they have a dissimilitude in 'direction of fit'.

---

<sup>6</sup>The term 'two factor thesis' is from [Falvey 2000] and recent statements of it can be found in [Mele 2009] and [Adams and Mele 1989].

<sup>7</sup>The accounts of practical knowledge developed by both Velleman [1989] and Setiya [2008] are therefore quite different from Anscombe's conception of it. I will return to this in section 2.3.

<sup>8</sup>I will use the terms speculative knowledge and theoretical knowledge interchangeably. Both terms have their shortcomings. It would go a bit far to call my knowledge that  $2 + 2 = 4$  speculative, but it is also strange to call my knowledge that my cat is on my lap theoretical.

Although the words ‘direction of fit’ themselves do not figure in *Intention*<sup>9</sup>, Anscombe’s example of the man who goes shopping has become the *locus classicus* of the idea of direction of fit:

Let us consider a man going round a town with a shopping list in his hand. Now it is clear that the relation of this list to the things he actually buys is one and the same whether his wife gave him the list or it is his own list; and that there is a different relation when a list is made by a detective following him about. If he made the list himself, it was an expression of intention; if his wife gave it him, it has the role of an order. What then is the identical relation to what happens, in the order and the intention, which is not shared by the record? It is precisely this: if the list and the things that the man actually buys do not agree, and if this and this alone constitutes a *mistake*, then the mistake is not in the list but in the man’s performance (if his wife were to say: ‘Look, it says butter and you have bought margarine’, he would hardly reply: ‘What a mistake! we must put that right’ and alter the word on the list to ‘margarine ’); whereas if the detective’s record and what the man actually buys do not agree, then the mistake is in the record. [§32]

We might say that the record the detective makes has to ‘fit the world’, whereas in the usual case of a shopping list, the world (what the man actually buys) has to ‘fit the list’ instead.

Now it is clear that speculative knowledge is the type of knowledge that has to fit the world. If I find out that hippos do not eat grass at all, I have to alter my beliefs. I won’t go up to the hippos to make them eat grass. In the case of practical knowledge, however, the world must fit the knowledge. Suppose you are stirring in your pan of rice, when someone asks you ‘What are you doing?’ You reply with ‘I’m making a risotto’, but she retorts ‘That cannot be! The stove isn’t even turned on’. In this case it would be strange if you would just keep on stirring and say ‘You are right, I only *thought* I was making a risotto but you made me realize that I in fact wasn’t doing so’. Such behaviour would make us doubt whether you were actually intending to make a risotto in the first place. If you were serious about making a risotto the natural thing to do would be to turn on the stove. The point that practical knowledge is the measure against which the world is measured, is important, delicate and often misunderstood or overlooked. But it is key to truly understanding practical knowledge. Something, I hope will become clearer during the course of this chapter. My defence of Anscombe’s position in chapter 2 will largely focus on this idea.

---

<sup>9</sup>The phrase was first coined in the context of speech act theory by Austin [1952]. Searle brought it into the philosophy of mind [Searle 1983].

The direction of fit of practical knowledge also explains why Anscombe writes:

[t]he contradiction of, 'I'm replenishing the house water-supply' is not 'You aren't, since there is a hole in the pipe', but 'Oh, no, you aren't' said by someone who thereupon sets out e.g. to make a hole in the pipe with a pick-axe. [§31]

On the one hand this is a clear and extremely vivid illustration. It nicely shows that a practical knowledge claim is about what one is actually doing, without conflicting with speculative knowledge claims about what happens. A speculative and a practical knowledge claim cannot gainsay each other because of their different directions of fit. My belief that there is war in the world can not be denied by my desire for peace; my desire for a beer isn't countered by my belief that the fridge is empty. But on the other hand it is very hard to see how a claim about what you are doing can be *contradicted* by the expression and execution of the intention to stop your action.<sup>10</sup> Contradictions typically hold between indicative statements and their negations, they typically occur in logic and *theoretical* reasoning. Anscombe, however, has an Aristotelean view of *practical* reasoning according to which the *conclusions* of bits of practical reasoning are actions. If an action can be a conclusion then might one action not also contradict another?

## 1.5 Inferred by practical reasoning

In *Intention* we read: "The notion of 'practical knowledge' can only be understood if we first understand 'practical reasoning'." [§33]. The book is, however, quite unclear about how the understanding of the one concept hinges on the understanding of the other. This section hopes to shed some light on this matter. I will not extensively discuss criticisms of Anscombe's Aristotelian idea of practical reasoning, but merely illustrate how practical reasoning elucidates practical knowledge.

Anscombe first construes practical reasoning very broadly as reasoning with "a view to action" [§33]. Suppose someone wants to make a risotto, she may then reason: 'In order to make a risotto I need to get some onions. They have onions at the grocery store, so I'll go there.' Here the words 'so I'll go there' stand proxy for the action. Notice that the action inferred (going to the store) is a means to an end (getting onions) which is

<sup>10</sup>Colloquially it isn't that strange to say of two actions that they contradict each other. Here is Douglas Adams: "He hoped and prayed that there wasn't an afterlife. Then he realized there was a contradiction involved here and merely hoped that there wasn't an afterlife." [Adams 1996 [1982], p.458].



itself a means to a higher end (making a risotto). Thus, according to Anscombe, practical reasoning is the deduction of means from ends. One has a certain end in mind and one tries to derive an action that is a means towards reaching it. Practical reasoning therefore reveals the same means-end structure that is also laid bare by the question ‘why?’ [§45]. Practical reasoning goes downward: I can get some onions *by means of* going to the grocery store; answers to the question ‘why?’ go upward: I’m going to the grocery store *in order to* get some onions.

This shows why practical reasoning is important to understanding intentional actions: it is the deduction of the structure that can be brought to the surface by application of the question ‘why?’. Not every intentional action, however, is always preceded by some specific deduction of means from ends. We don’t ponder the merits of brushing our teeth twice a day, although every time we brush our teeth we do it for a reason. Thus, practical reasoning can also be a way of bringing the already existing rational order of our intentional actions into view. To see why practical reasoning is especially important for understanding practical knowledge, we need to take a closer look at some of the peculiarities of Anscombe’s ideas about practical reason. Her view diverges greatly from what we might call the ‘Standard view’<sup>11</sup>. According to the Standard view practical reason has the following general form:

I want A  
B is a way of attaining A  
∴ It is good for me to do B

Practical reasoning proceeds, according to this model, from a premise about what I want and a premise that relates means to ends towards a judgement about what is good for me to do. But isn’t this simply a piece of theoretical reasoning about what one should do? It does not seem to be with a view to action, it is rather with a ‘view to judgement.’ Anscombe’s account of practical reasoning differs from such a standard account in two important ways: 1) She does not think that ‘wanting’ should be included in the premises, and 2) she thinks that the conclusion of a chain of practical reasoning is an action and not a judgement about what would be good to do.

Both points are related, but I will start with the first. Anscombe writes: “the role of ‘wanting’ in the practical syllogism is quite different from that of a premise. It is that

<sup>11</sup>I dub this view ‘standard’ because it is held by many philosophers and psychologists alike, see the influential humean theory of motivation [Smith 1987] and Bratman’s [1987] belief-desire-intention (BDI) model for practical reasoning.

whatever is described in the proposition that is the starting-point of the argument, must be wanted in order for the reasoning to lead to any action” [§35]. This becomes clearer once we see that the same considerations might culminate in different actions:

Strong alkaloids are deadly poison to humans.

Nicotine is a strong alkaloid.

What’s in this bottle is nicotine.

In practical reasoning, such reflection might “terminate either in careful avoidance of a lethal dose, or suicide by drinking the bottle” [Anscombe 1974, p. 382]. Which action is performed depends on what the agent wants. If she sees prolongation of her own life as good she will avoid drinking the bottle, if she wants to put an end to it she will drink the nicotine.<sup>12</sup> The considerations above might also occur in a piece of theoretical reasoning of which the conclusion might be the judgement that drinking the bottle’s contents would be deadly to humans. Anscombe expresses the difference between practical and theoretical reasoning as follows: “Truth is the object of judgement, and good the object of wanting” [§40]. The goal of theoretical reasoning is to arrive at a true judgement, and the attainment of this goal is the driving force behind theoretical reason. But practical reasoning is done with a view to action. Its driving force is something that is seen as good (something wanted), and whenever means of attaining the end are deduced, they are also seen as good (insofar as they contribute to the attainment of the end<sup>13</sup>). Therefore, practical reason concludes in action, i.e. in performing the means to achieve the end.<sup>14</sup> This explains the second difference between the Standard account and Anscombe’s view.

Now that we have a better understanding of practical reasoning, we can return to practical knowledge. Since the conclusion of practical reasoning is an action instead of a judgement, the knowledge derived from practical reasoning is knowledge of the action. But because practical reasoning is the deduction of means from ends, the knowledge derived is also knowledge of the *means-end structure* of the action. Practical knowledge is not merely knowledge that I am causing a certain event in the world, it is the knowledge that I am intentionally performing an action – which is more than just an event, it is a

<sup>12</sup>For Anscombe, wanting something is thinking it is good. You cannot simply want a saucer of mud without finding something about it desirable. See [§37].

<sup>13</sup>I add this proviso because we do not think of the destruction of the human race as good, although we might say that it is good insofar as it puts an end to all wars.

<sup>14</sup>As Müller suggests: there isn’t a large difference between the calculation steps in practical reasoning compared with theoretical reasoning about means and ends. The difference solely consists in the fact that the one reasoning is done with a view to action, while the other with a view to judgement. See [Müller 1979].

rationally structured process.<sup>15</sup>

Another point that seems to follow from the fact that practical reason does not terminate in judgements, is the difference in direction of fit between practical and speculative knowledge discussed in section 1.4. Theoretical reason aims at judgements about what is the case, therefore the knowledge derived from theoretical reasoning should fit that which is the case. But practical knowledge is derived from practical reasoning, it is the knowledge that I am doing something with a certain aim and that it is this aim that should be realized. In other words: it is the world that should fit my knowledge.<sup>16</sup>

## 1.6 Know how

“A man has practical knowledge who knows how to do things” [§48]. Indeed, in philosophy the term ‘practical knowledge’ is often used to refer to *know how*. Know how was distinguished from propositional knowledge (knowing that) by Ryle [2002 [1949]]. I can know *that* there is a bike in the shed and know *how* to ride on one. The concept of *know how* is quite broad. On the one hand we can say that Gil knows how to ride a bike because she can give a lecture on it, even though she has no legs. On the other hand, we can say of Bert that he knows how to ride a bike because he possesses the *skill* of bike-riding, even though he is a monkey and cannot exactly tell us how he is doing it.<sup>17</sup> Now where in the spectrum of know how can we fit Anscombe’s practical knowledge?

A man has practical knowledge who knows how to do things; but that is an insufficient description, for he might be said to know how to do things if he could give a lecture on it, though he was helpless when confronted with the task of doing them. When we ordinarily speak of practical knowledge we have in mind a certain sort of general capacity in a particular field; but if we hear of a capacity, it is reasonable to ask what constitutes an exercise of it. [...] In the case of practical knowledge the exercise of the capacity

<sup>15</sup>Here Anscombe clearly diverges from Davidson’s doctrine that actions are just a type of event (i.e. bodily movements) [1963]. This matter will be revisited extensively in chapter 4.

<sup>16</sup>As Haase [2010] argues, this difference in direction of fit is nicely reflected by the fact that we might reason from  $p$  and  $p \rightarrow q$  to  $q$  in theoretical reasoning, whereas in practical reason we proceed from wanting  $A$ , and  $B \rightarrow A$  (read as:  $B$  leads to/is a way of attaining  $A$ ) to doing  $B$ . But although the different direction of the arrows between these particular forms of practical and theoretical reasoning is a compelling feature, one should be wary of exaggerating the analogy between practical and theoretical reasoning. In practical reasoning, for instance, we do not say of *propositions* that they are good, while we do ascribe truth to propositions in theoretical reasoning. Also, that  $B$  is a way of attaining  $A$  is not itself good, whereas that  $q$  follows from  $p$  is itself a true proposition.

<sup>17</sup>In fact you don’t need to be a monkey to be in this position. Most of us suppose that we go round a corner by turning the handlebar, but in fact we lean into the corner in order to turn.

is nothing but the doing or supervising of the operations of which a man has practical knowledge; but this is not just the coming about of certain effects [...] for what he effects is formally characterised as subject to our question 'Why?' whose application displays the [means-ends] order which we discovered. [§48]

Thus the answer seems to be that it fits somewhere in between Gil's and Bert's types of knowledge. Know how is an insufficient description of practical knowledge, because practical knowledge is the knowledge we have of our own action, as we saw in section 1.2, and Gil is evidently not riding a bike while she is giving the lecture. She merely seems to possess speculative or theoretical knowledge of how to ride a bike.

It is more difficult to see why Bert's know how also might fail to attain the status of practical knowledge in Anscombe's sense. But I think we can gather some insight by considering the distinction between specific and general know how made by Sebastian Rödl [2010]. Specific know how contains reference to particulars (e.g., I know how to escape from *this* prison cell), as opposed to general know how (e.g., I know how to escape from prison cells). General knowledge how should be identified with skills, crafts, techniques, practices or even life.<sup>18</sup> According to Rödl, particular actions employ specific know how. This know how concerns specific means and ends: I can escape from this prison cell by using this key on this lock. I can get it by luring this dog with that key around his neck towards me... I can lure the dog towards me with this bone etc. It is specific know how that enters into practical reasoning.<sup>19</sup> As we saw in section 1.5, practical reasoning deduces the order of means and ends (also revealed by the 'why?' question) that is essential for every intentional action. The practical knowledge of our own intentional actions is practical knowledge of this structure. I know, while I am cutting these onions, that I am making a risotto – not just any one but precisely the one that will be the result of my making it. It is thus contained in my practical knowledge of my action that cutting these onions is a means to making the risotto, and thus that I *know how* to make the risotto (i.e., by cutting the onions).

Let's go back to Bert. Bert has a skill, he knows how to ride a bike. This is general know how that is retained even if Bert is not riding a bike right now. Therefore it is quite different from Anscombean practical knowledge, because this is knowledge of the current intentional action. Even if Bert is exercising his bike riding skill unintentionally

---

<sup>18</sup>It is true that the know how of bike riding is quite different from knowing how to live. The former constitutes a skill, while the latter is related to practical wisdom (*φρόνησις*).

<sup>19</sup>General know how may also enter into practical reasoning, but only in so far as it is the source of specific know how (e.g. I know how to ride this bike, because I know how to ride bikes).

(if such a thing is possible), we can still say that he has general know how of bike-riding, but that would not be the specific know how, the practical knowledge, of his intentional action. If Bert was however intentionally (if that is possible) riding his bike to get to that banana, he would have practical knowledge and specific know how. So, (specific) know how can amount to practical knowledge when it is exercised in “doing or supervising” [§48] the operations that are known.

The next section will focus on the active causal (but not merely efficient) role of practical knowledge in action. For now, I hope to have made clear that Anscombe’s practical knowledge is not to be confused with the broader concept of *know how* that pervades philosophy, but that practical knowledge nevertheless presupposes (general) know how. I cannot intentionally be making a risotto if I do not know how to make one, or at least if I do not how to find a recipe.

## 1.7 The cause of what it understands

The last of Anscombe’s characterizations I want to discuss can sound quite archaic: ‘practical knowledge is the cause of what it understands’ [§48]. Perhaps, this is because it is a translation from St. Thomas’s phrase: *qui est causa rerum interlectarum* [Summa Theologica, Ia IIae, Q. 3, art. 5, obj. 1]. In this section I will explain what is meant by it, and show its relation to the other characterizations.

Remember that intentional actions are special because they have a certain internal rational structure. Making a risotto, for instance, consists of: cutting onions, taking butter out of the fridge, cooking rice, adding wine, etc. Remember furthermore (from section 1.4), that the world has to fit the practical knowledge I have of my risotto making. I already (know that I) am making a risotto, when I start to cut the onions. In fact I cut the onions *because* I am making a risotto. It is because we know what we are doing and what the means are, that we perform the actions that constitute the means. As Michael Thompson [2008] explains: it is the knowledge of what I am doing that *unifies* the phases of what I am doing as being phases of a single intentional action. My knowledge that I am making a risotto makes my cutting the onions a part of risotto making. I could have made the exact same onion-cutting moves, but if my practical knowledge was different, I would not be making a risotto, I would be making pasta instead.

Calling practical knowledge the *cause* of what it understands “means more than that practical knowledge is observed to be a necessary condition of the production of vari-

ous results” [§48]. It does mean that “without it what happens does not come under the description — execution of intentions — whose characteristics we have been investigating.” [§48] When we are making a risotto, the description ‘making pasta’ does not apply to my onion cutting, while the description ‘making a risotto’ does, and this is caused by my practical knowledge. When I act intentionally, I am not merely producing “various results”, “or the events of my movements themselves, but rather the *fact* that I am doing this rather than that”, as Moran puts it [2004].<sup>20</sup> It is the description specified in my practical knowledge that makes it the case that *this* description counts as a description of an intentional action. Without the practical knowledge of my intentional action, it would not be an action at all.

It is clear that the causation Anscombe has in mind is not merely efficient causation. Perhaps we might call it formal causation in that the knowledge causes the form (the means-end structure) of what is happening. But it would go too far to claim that practical knowledge only causes the *form* of the event, while it isn’t productive of the actual *happening* of the event.<sup>21</sup> To repeat Anscombe: that knowledge is the cause of what it understands “[m]eans *more* than that practical knowledge is observed to be a necessary condition of the production of various results [emphasis NvM]” [§48]. That it means *more*, does not mean practical knowledge *isn’t* productive. Moreover, it does not mean that it *only* causes the form of the action.

It is, however, very difficult to fully understand how practical knowledge produces our actions. It comes close to the enigma of why human beings can act at all. But I think that the reflections on practical reasoning and know how from the previous sections can take us some steps towards understanding.<sup>22</sup> As we saw, the exercise of the capacity of practical knowledge is nothing but the doing or supervising of the operations of which a man has practical knowledge. How might this doing or supervising go about? Mathias Haase [2010] suggests the following picture: When I am getting an apple, I derive the specific phases of this action from the knowledge of what I’m doing. This derivation determines the specific shape of the action. I could have chosen to get the apple by climbing up the tree, but I have decided to rather hit it with a branch so it falls to my feet.

<sup>20</sup>Moran’s talk about facts reminds one of McDowell. The latter notoriously claims that our observation “does not fall short of the facts” [McDowell 1996]. Anscombe’s idea of intentional actions might be subsumed by the slogan: Our actions do not fall short of the facts either.

<sup>21</sup>This seems to be the suggestion of [Moran 2004] and [Teichmann 2008].

<sup>22</sup>For a full understanding I think we need a good account of what sort of things humans are, and we need a broader account of causation than the efficient event causation that is popular in twentieth century philosophy of science.

As the action proceeds towards its conclusion, through the phases I derive, my practical knowledge brings the action as a whole about. All this suggests that the practical knowledge of the phases of my action is playing a productive role in bringing about my intended results in precisely the way in which I produce them. Perhaps it also helps to remind oneself of the fact that the conclusion of practical reason is not a judgement. It is an action. Therefore, practical reason itself is productive. In chapter 4 I will return to this matter, by discussing how practical knowledge can cause actions.

## 1.8 Practical knowledge

This chapter aimed at arriving at a uniform understanding of what practical knowledge is, or at least of what practical knowledge might be. I think we have achieved that feat. We have taken quite a step because it is far from easy to see how *one* thing can be *in* the action, knowledge *of* the action, knowledge of the structure of the action, knowing how to do the action, and even the cause of the action. In trying to see how one thing could be all of this, it helped considerably to look at practical reasoning and the role practical knowledge plays in action. It also helped to contrast practical knowledge with speculative knowledge. As it turned out these types of knowledge are different because they have different directions of fit. As we will see in the subsequent chapters, a lot of confusion about Anscombe's position – but also about action at large – stems from a misguided picture of the concept of *practical knowledge*. I hope, now that we have a proper understanding of the concept, that the rest of the thesis can lift some of the fog that has descended in the last few decades. In chapter 2 I will argue that practical knowledge is often misunderstood by contemporary critics, because they don't fully appreciate the world to mind direction of fit of practical knowledge (as elucidated in §1.4). The contrast between speculative and practical knowledge will also be important in chapter 3. There we will see that practical knowledge isn't the only knowledge that plays a role in intentional action. Speculative knowledge is also required, especially since it can enter into our practical reasoning (see §1.5). In chapter 4 the idea of knowledge that is the cause of what it understands (see §1.7) takes center stage, as I will argue a causal theory of the nature of action can only work if it takes Anscombean practical knowledge into account.

# Chapter 2

## Writing Without Ink

### Reflection and interpretation

*Often quoted, sometimes read, rarely understood, Anscombe's Intention is nevertheless the defining moment in twentieth-century philosophy of action.*

– J. David Velleman

#### 2.1 Ten carbon copies

[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 it intentionally. [Davidson 1978, p. 129]

This example, “easily multiplied, convincingly demonstrate[s] that non-observational knowledge (or even belief) is not as coincidental with intentional action as Anscombe would have it” [Paul forthcoming]. In this example ten carbon copies are being made intentionally, but the agent has no knowledge that he is doing so. Therefore, Anscombe’s ‘overly strong’ claim that intentional actions are those “of which a man has practical knowledge” [§48] is flatly wrong. Such is the current opinion in the philosophy of action.

In this chapter I will formulate and defend an answer to this example, an answer that can already be found in the pages of *Intention*, and was recently brought back to attention by Thompson [2008]. The answer comes down to the claim that the person who is writing heavily on a stack of carbon paper, with the aim of making ten carbon copies, *does* know that he is making ten carbon copies. He knows this, even if his current



writing only results in eight copies. The act of making ten carbon copies may include checking how many you have after the first try, and making two extra if you have eight copies. Or throwing three copies away if you have thirteen. One simply does not have to believe that one's first try is successful in order to know that one is doing it.

Davidson's example is, I think, part of a larger class of examples that express a worry about practical knowledge: how can we know what we are doing when the world may be secretly uncooperative? Suppose a blind man is writing his name on a piece of paper. When he is writing he knows that he is doing so. But what if you can see that his pen has run out of ink, doesn't it turn out that he only thought that he was writing but actually wasn't? Anscombe's answer would be that he actually *did* know that he was writing, even though no letters appeared on the paper. It would be true to claim that he was *writing* his name, but it would also be true that he *did not write* it.<sup>1</sup> Some think that this claim is too ridiculous to take seriously, and that it is in fact so ridiculous that Anscombe could not have made it — although they would perhaps accept it from an audacious graduate student like myself. I will therefore start this chapter by defending my interpretation of Anscombe in section 2.2, before I set out to argue that the claim is systematically correct in §2.3 and §2.4. I will conclude by looking at a version of the carbon copies example that might be more troublesome for a theory of practical knowledge and suggest a way to deal with it (§2.5).

## 2.2 Understanding practical knowledge

The claim that one can know that one is writing even if no letters appear on paper, seems *prima facie* wrong. Knowledge is supposed to be factive, so how can you know that you are writing when in fact you aren't producing any letters? Anscombe was of course well aware of the factivity of knowledge, as is clear from the following passage:

Say I go over to the window and open it. Somebody who hears me moving calls: What are you doing making that noise? I reply 'Opening the window'. I have called such a statement knowledge all along; and precisely because in such a case what I say is true — I do open the window; and that means that the window is getting opened by the movements of my body out of whose mouth those words come. [§28]

---

<sup>1</sup>Note here that 'writing' is progressive, while 'dit not write' is perfective. We will linger on this difference in aspect later in §2.3.

But what if the handle breaks and the window isn't getting opened? Does Anscombe really want to insist that the agent *knew* she was opening the window? In his book on Anscombe, Teichmann tries to avoid attributing such a reckless position to her:

If doubt arises as to whether I am in fact doing what I say I am doing, then I will need to resort to observation, or to asking others. But doubt is the exception, not the rule, and lack of doubt is the default position. For others, if I say I am V-ing, then that counts as grounds for them to say that I am V-ing; but of course it may be perceptibly evident that I am not V-ing, and this will trump what I say. [Teichmann 2008, p. 26]

He finds support for the idea that the perceptually evident can trump the agent's practical knowledge in the following passage from *Intention*:

Sometimes, jokingly, we are pleased to say of a man 'He is doing such-and-such' when he manifestly is not. E.g. 'He is replenishing the water-supply', when this is not happening because, as we can see but he cannot, the water is pouring out of a hole in a pipe on the way to the cistern. [...] or again one might say of a madman 'He is leading his victorious armies.' [§23]

Now if we want to hold on to the factivity of knowledge, shouldn't we say that an agent only has practical *belief* in cases where what is happening does not correspond with what she claims to be doing?<sup>2</sup> If this is correct one might wonder why Anscombe does not speak about 'practical belief' at all. The claim made by both Teichmann and Velleman is that Anscombe doesn't talk about practical belief because she thinks that "the failure to execute intentions is necessarily the rare exception." [§48] Here is Velleman:

Anscombe was also, as I interpret her, a reliabilist about knowledge — in particular, about what is "known by being the content of intention". She thought that a reliable connection in general between what's intended and what's done is sufficient to confer the status of knowledge on a particular intention, provided that the connection holds up in the particular case. [Velleman 2004, p. 227]

Teichmann interprets Anscombe similarly, although he envisages an important role for language:

Linguistic training enables me to say straight off what I am doing. The training is successful to the extent that my utterances are generally in sync with my actions. I need no grounds for these utterances. [Teichmann 2008, p. 26]

---

<sup>2</sup>Teichmann isn't alone in suggesting this, see [Velleman 2004] and [Paul forthcoming].

As Velleman and Teichmann have it, the reliability of our mechanisms of intention execution justify our practical beliefs.<sup>3</sup> Therefore when my intention in action provides me with a true belief about what I am doing this belief is justified and thus knowledge. My beliefs about what I am doing only fail to amount to practical knowledge if they are wrong, but this is the exception.

In the remainder of this section I will argue that Anscombe does not have an implicit notion of practical belief, nor was she a reliabilist. Instead, as we will see, she *does* defend the rash claim that we *always* know what we are intentionally doing. The fallback or ‘charitable’ positions Velleman and Teichmann try to ascribe to Anscombe are in fundamental conflict with her theory of action. Whenever someone acts intentionally she has practical knowledge, regardless of whether she is successful or not.

Anscombe was well aware of the peculiarity of her claim that there is such a thing as knowledge in intention, distinct from observational knowledge of what is happening:

What can opening the window be except making such-and-such movements with such-and-such a result? And in that case what can *knowing* one is opening the window be except knowing that that is taking place? Now if there are two *ways* of knowing here, one of which I call knowledge of one’s intentional action and the other of which I call knowledge by observation of what takes place, then must there not be two objects of knowledge? [§29]

As it turns out Anscombe answers this last question in the negative. And *Intention* can be seen as an attempt to explain how it can be possible that there are two different knowledges (practical and speculative) of exactly the same thing. A first suggestion is that there are two descriptions of the same thing, both of which are known. For instance, I know that the woman on the television is Beatrix and that she is the queen of the Netherlands. But this cannot be what is going on in the case of intentional action. The description ‘opening the window’ is exactly the same whether one knows it by observation, or by it being one’s intentional action. Another suggestion is that we cannot see how there can be two knowledges of the same object because our picture of what kinds of knowledge exist is too narrow:

Can it be that there is something that modern philosophy has blankly misunderstood: namely what ancient and medieval philosophers meant by *practical knowledge*? Certainly in modern philosophy we have an incorrigibly contemplative conception of knowledge. Knowledge must be something that is judged as such by being in accordance with

---

<sup>3</sup>[Peacocke 2003] also interprets Anscombe in this way.

the facts. [...] And this is the explanation of the utter darkness in which we found ourselves. For if there are two knowledges — one by observation, the other in intention— then it looks as if there must be two objects of knowledge; but if one says the objects are the same, one looks hopelessly for the different *mode of contemplative knowledge* in acting, as if there were a very queer and special sort of seeing eye in the middle of the acting. [§32]

Thus, we failed to see how we can have two knowledges of the same object because our conception of knowledge was restricted to speculative knowledge that has to fit the world.<sup>4</sup> There is no “seeing eye” in action, because action, unlike observation, doesn’t yield speculative knowledge. This point is rarely understood and it seems that Teichmann’s and Velleman’s conception of practical knowledge exactly is just a different mode of contemplative knowledge. For them, practical knowledge is judged by the facts. If it does not correspond to the facts it is at most practical belief. On their account we are lucky because we do have practical knowledge most of the time. When we intend to do something, it is an exception if we do not succeed. But for Anscombe practical knowledge isn’t a mode of contemplative knowledge, it is entirely different in that it is not judged by the facts. The facts are judged by it. Therefore, whether an agent fails or succeeds does not impinge on his practical knowledge.

The latter point becomes clear when Anscombe discusses an example similar to the blind man who writes his name, the case of writing ‘I am a fool’ on the blackboard with the eyes closed:

That intention [to write ‘I am a fool’ on the blackboard] for example would not have been executed if something had gone wrong with the chalk or the surface, so that the words did not appear. And my knowledge would have been the same even if this had happened. If then my knowledge is independent of what actually happens, how can it be knowledge of what does happen? Someone might say that it was a funny sort of knowledge that was still knowledge even though what it was knowledge of was not the case! On the other hand Theophrastus’ remark holds good: ‘the mistake is in the performance, not in the judgment’. [§45]

Sometimes it might be difficult to assess whether someone who says that she is X-ing actually possesses practical knowledge that she is X-ing. An example of this is the case of the madman who claims to be leading his victorious armies. We might try to question him further, and ask things like: where is your army? How are you leading them? What

<sup>4</sup>See §1.4 for the difference in direction of fit between speculative and practical knowledge.

is your quest? Now if he answer these questions (My army consists of this table and the three chairs over there, I'm leading them by giving them orders via hand signals, and we are on a quest to overthrow the Roman empire) we might indeed say that he is intentionally 'leading his army'. Of course, it is quite clear that in saying this we mean quite a different thing from what we normally would mean by 'he is leading his victorious armies', but this is due to the peculiar state of mind of our madman, and not a problem specific to Anscombe's theory. When he fails to answer any further questions, he will probably not possess practical knowledge at all, and therefore we will say that he is not actually performing an intentional act. This once again shows the inseparability of practical knowledge from intentional actions. As Anscombe has it:

[W]ithout [practical knowledge] what happens does not come under the description — execution of intentions — whose characteristics we have been investigating. [§48]

## 2.3 Fitting practical knowledge

A reply to the argument brought forward in the last section might be that the claim that the blind man knows that he is writing his name despite the fact that his ink has run out still is hopelessly counterintuitive or even paradoxical.<sup>5</sup> If Anscombe actually held such a claim, then her theory of practical knowledge is even more obviously false than we previously thought.

There are two alternatives to Anscombe's theory of practical knowledge. The first is the two factor thesis mentioned in section 1.3. Proponents of this thesis hold that we can only know our own intentions without observation, but need observation to know what happens. The second alternative is the claim that we can have practical knowledge of what happens, but only when our actions are successful. When our actions fail, the knowledge is defeated because the description it specifies of what is done fails to be the description of what is happening.<sup>6</sup> In this section I will try to refute both alternatives. The next section will show why the consequence that one can know that one is X-ing, even when one fails to X, isn't that rash or counterintuitive after all.

Interestingly, the main problem Anscombe makes out for the two factor thesis precisely occurs when the execution of intention fails. She envisages a case where she intends to move her toe, but fails to do so:

---

<sup>5</sup>See [Grunbaum forthcoming].

<sup>6</sup>The position attributed to Anscombe by Teichman and Velleman, mentioned in the previous section, is a variant of this position.

nothing guarantees that my toe moves when I ‘move my toe’; so the only thing that does happen is my intention; but where is that to be found? I mean: what is its vehicle? [§29]

For contemporary proponents of the two factor thesis (like Al Mele) this is not a problem at all. Having an intention is being in a mental state. When the toe doesn’t move, this must be due to something that interfered with the causal relation between the mental state and the bodily movement. It would go beyond the scope of this thesis to take sides on the debate about mental states,<sup>7</sup> therefore I will provide some different reasons to doubt the two factor thesis.

A first problem is provided by the case we already encountered in chapter 1: I can be sitting on my couch, reading a paper while I’m intentionally baking a cake at the same time — I am only waiting for the dough to rise. In this case I know more than that I intend to bake a cake, I know that I am doing so even if I can’t see the bowl with the dough from where I’m sitting. Hence, contra the two factor thesis, I do not need observation to know what I am doing. It might be objected that I remember that I saw myself successfully making the dough and putting it in a bowl to rise. But what if I could use the exact same dough, both to make an apple pie as well as a cake? Since I only saw myself making the dough, the two factor theorist is entitled to claim that I only know that I have made dough and know that I intend to make a cake, but that I still don’t know that I am making a cake. But this fails to correspond to how we normally speak and think. There is nothing wrong with saying of someone who is rolling out the dough, that she is baking a cake, or to say that I am making tea when I’m putting water in the kettle. As Kevin Falvey puts it:

In relying as it does on a more or less radical distinction between intending and doing, the two factor thesis overlooks the continuity that exists between them. [Falvey 2000, p. 26]

Another issue for proponents of the two factor thesis is that we often do not doubt what we are doing when we observe that what is happening is inconsistent with our intention. Suppose I intend to drive from Utrecht to Amsterdam (northwards). After accidentally taking a wrong turn I get lost, and some time later I read a signpost that tells me I’m on the A2 driving in the direction of Eindhoven (southwards). According to the two factor theory I would now conclude that I only intended to go Amsterdam but that I am in fact going to Eindhoven. But in fact I would conclude no such thing, I am still going to Amsterdam and I know that I am going there. If someone were to call me and ask me

<sup>7</sup>Although I will refute causal action theory in chapter 4.

where I was going, I would confidently answer: ‘I’m going to Amsterdam’. Later in this thesis (§4.5) I will even suggest that it is my knowledge that I am going to Amsterdam that moves me to turn around at the next exit.

Although the above isn’t a knock down critique of the two factor thesis, it would certainly be nice to have a theory of the knowledge of our actions that can account for cases in which we know what we are doing without observation. But can such an alternative to the two factor thesis avoid Anscombe’s supposedly paradoxical result that we *always* know what we are doing? Velleman, Falvey, Setiya and Moran believe so.<sup>8</sup> They claim that we *almost always* know what we are doing. Only when we fail to be successful practical knowledge fails. Below I will show that such an alternative is either inconsistent, or has to collapse to a form of causal action theory (to be criticised in chapter 4).

According to Moran, Anscombe arrives at the paradoxical result because she

seems to confuse the requirement of truth for knowledge, which applies to any knowledge practical or speculative, with the question of whether it is the action or the putative knowledge that is to be corrected in the case of disparity [...] The distinction between practical and speculative knowledge does not concern the requirement of truth, but the question of what is to be corrected in cases of failure of fit. [Moran 2004, p. 60-61]

Thus according to Moran, there are two separate issues: first, does an agent believe that she is doing such and such, and second, what must be done if her belief turns out not to be true? If she correctly believes that she is doing A, this might be knowledge. Whether the knowledge is speculative or practical then depends on what will be altered, if the belief somehow turns out to be false. This picture of the direction of fit of practical knowledge corresponds to the contemporary view on the direction of fit of *propositional attitudes*. As Searle [1985, 2003] argued, we can take several different attitudes towards a proposition *p*. Each attitude comes with a direction of fit.<sup>9</sup> If I believe that ‘Niels’ is written on this page, I will change my belief if it turns out to be wrong. If I want that ‘Niels’ is written on this page, and find out that it isn’t, I will proceed by writing my name on this page. But the proposition ‘Niels’ is written on this page’ is true or false, independent of my attitude towards it.

But if this contemporary conception tells us what direction of fit is, what Anscombe had in mind surely wasn’t direction of fit after all. Remember her example of the two

<sup>8</sup>See [Velleman 1989], [Falvey 2000], [Setiya 2007] and [Moran 2004].

<sup>9</sup>Also see [Velleman 1992].

shopping lists (see §1.4). In the case where the man bought margarine although the list he got from his wife mentioned ‘butter’, Moran has to say that the list was incorrect because it did not correspond to what the man actually bought. He will then account for the word to world direction of fit of the list by claiming that the list should not be corrected by crossing off ‘butter’ and replacing it with ‘margarine’. It should rather be corrected by the man returning to the grocery store to return the margarine and by butter instead. But this fails to capture the difference in nature of the two lists. Imagine the man telling his wife: ‘Ok I will return to the shop, but nevertheless *your* list is false’ and their subsequent argument! Anscombe’s point was that it is wrong to think of the list as being incorrect when the man buys margarine instead of butter. Rather, the list is the *norm*, against which the man’s performance is judged. In its status of being the norm, practical knowledge corresponds to the list:

[W]e do *not* say: What you *said* was a mistake, because it was supposed to describe what you did and did not describe it, but: What you *did* was a mistake, because it was not in accordance with what you said. [§32]

Despite the fact that Moran (wrongly) thinks that Anscombe is confused about direction of fit, he still wants to follow her in claiming that having practical knowledge is constitutive of intentional action. Since practical knowledge specifies the description under which what happens is intentional, it is what makes a mere bodily movement into an intentional doing:

[A] failure of practical knowledge means that the event in question no longer counts as a thing of a certain kind (an intentional action, as described in those terms) [...] An object of speculative knowledge, on the other hand, is independent, both formally and materially, of being known. [...] Because of this, a failure of some claim of speculative knowledge does not make any difference to the character of the object of that failed claim. [Moran 2004, p. 67]

Now recall that Moran wanted to claim that an agent has practical knowledge when what happens corresponds to what she thinks she is doing, and only has practical belief when it fails to correspond. But if what is special about *practical* knowledge is that it is the criterium for what counts and does not count as an intentional action, it seems impossible to have a mere belief that is practical in the same sense. That is it seems impossible to have a *practical belief*. Practical belief should determine, by virtue of its being practical, that what happens counts as the intentional action it is believed to be. If the blind man has practical belief that he is writing his name, what happens comes under



the description ‘writing one’s name’ precisely because he has that belief. In other words, because the belief determines that what is happening is truly described as ‘writing ones name’, the blind man’s belief that he is writing his name must be true. Practical belief would be somewhat like a selffulfilling prophecy: it is always true because it determines that the world is precisely as it is believed to be. Since it always has to be true (and not only reliably true), it is impossible to come across something that is practical belief but isn’t practical knowledge. Hence it makes no sense to introduce a separate category of practical belief apart from practical knowledge.

It turns out that Moran’s theory is incoherent: on the one hand he agrees with Anscombe that if practical knowledge fails, the event in question does not count as an intentional action. But on the other hand he also thinks that when the action in question fails (because the ink has run out) the practical belief does not count as knowledge — which cannot be the case if the practical belief truly is the norm for what happens. Thus, Moran is unclear about whether it is an agent’s practical belief/knowledge, or what happens in the world, on which the (un)intentional character of an event depends. At this point every theory of practical knowledge has to make a choice, either to go fully Anscombean by picking the former, or to go for the latter and abandon the idea that practical knowledge has a special role in deciding whether an event is an intentional action. In this thesis I am making an attempt to defend the first choice, but the second option is favoured by the majority of contemporary philosophers of action.

Although Moran is unclear on his choice, both Velleman and Setiya go with the majority and decide that the world is the norm for the agent’s knowledge. Therefore, they have to find another way to explain the inseparability of her knowledge of what she is doing and her doing it intentionally.<sup>10</sup> Their solution is that an event only is an intentional action when it is (efficiently) caused by the agent’s knowledge of it. This places them in the category of causal action theorists, who try to reduce the intentionality of actions to their being caused in a certain way. In chapter 4 I will argue that causal action theory is misguided and should be abandoned in favour of an Anscombean view. For now it is important to stress that what Setiya and Velleman call ‘practical knowledge’ is quite different from what Anscombe expresses with the term. Only her theory of practical knowledge leads to the seemingly paradoxical result that the blind man knows that he is writing his name despite the fact that his ink has run out. The next section will attempt to show that this result isn’t paradoxical at all, and is in fact quite natural.

---

<sup>10</sup>See §1.3

## 2.4 Action and aspect

Michael Thompson argues that a difficulty in comprehending Anscombe's teaching arises from an emphasis on certain types of examples in the action theoretic literature since Davidson. Davidson's, and subsequent, examples of actions are almost always formulated under the perfective aspect: I *flipped* the switch; Brutus *kissed* Ceasar; Jones *did* it slowly, deliberately, in the bathroom with a knife at midnight (and what he *did* was butter a piece of toast). These perfective descriptions entail the existence of a completed event: a flipping, a kissing or a buttering. But usually such actions take time, and hence they can be interrupted. This is colourfully illustrated by Thompson:

[I]t might be that at the fateful moment when I flipped the switch, and there was a switch-flipping . . . at that moment Jones *was buttering* the toast. But the wiring I triggered was faulty, and so I thereby electrocuted the poor devil, who was leaning against a copper water-pipe. Toast and butter were only a half-centimeter apart, when the knife and toast fell to the ground, and with them Jones, quite dead. [Thompson forthcoming]

If this episode happened in the past, we will have to say that Jones *did not* butter the toast, i.e., there was no completed toast-buttering. Nevertheless the past *imperfective* phrase 'Jones was buttering a toast when he died' seems clearly true. Interestingly, the possibility of failure is in no way peculiar to events that are intentional actions. We can, for instance, say of a tree that it *was falling* over but *did not fall*, because something stopped it. Thompson argues that, due to the focus on examples formulated perfectly, post-Davidsonian action theory tends to forget that events are processes that take time and can fail:

*There is thus something that Davidson's doctrine of events or of things that happened is missing.* —Namely, not to put too fine a point on it, the things that *didn't* happen . . . but were happening. [Thompson forthcoming]

Interestingly enough, Anscombe's examples of intentional actions are typically formulated in the progressive: fetching my camera, replenishing the water supply, opening the window, going to the Hereford market. As we saw in §1.3, practical knowledge is the knowledge one has during one's performance of the action, hence practical knowledge claims are also formulated progressively: I *am making* a risotto. The claim that I *did* make a risotto, by contrast, involves checking whether I was successful, it is thus a claim to speculative knowledge.<sup>11</sup> It is the difference in aspect that Anscombe had in mind

<sup>11</sup>As we will see in chapter 3 such speculative knowledge of the success of one's action can play an important role in the performance of certain actions. E.g., my eyes can aid my writing in making sure that

when she wrote (using the present tense):

A man can be *doing* something which he nevertheless does not *do*. [§23]<sup>12</sup>

With this difference between the progressive and perfective aspect in mind, it doesn't seem problematic (let alone paradoxical) to say that Jones knows that he is buttering a toast when he is moving the knife towards the toast, even though his premature death prevents his action from being successful. The case of writing without ink may seem a little bit more problematic because the hand movements made by the blind man do not currently contribute to the emergence of letters, whereas Jones's moving of the knife towards the toast was contributing towards the buttering of the toast before disaster struck. However, Thompson argues that:

it is a mistake to look, at each moment at which a progressive proposition is true, for something in which the progress might be supposed to consist; indeed, it is often in the nature of an event or process-form that there should be times at which nothing of the sort can be found, as any piece of music is likely to contain silences. [Thompson 2008, p. 141]

Hence, it is perfectly fine to say that I *am baking* a cake, while I am reading a newspaper and the dough is resting in a bowl in the kitchen. Both in this case and the writing blind man's case, the agent's current bodily movements are not contributing much to the completion of the event, but that does not mean that there is no baking or no writing. As soon as the blind man finds out that his ink has run out, he will undertake an action that will put him back on track, like grabbing a new pen or refilling his current one.

Doesn't this mean that one can claim to be doing anything at any time, and be right? Can I for instance claim that I am currently solving all the problems of philosophy? Of course not — in fact I do not think that my thesis is a means towards solving all philosophical problems, nor do I intend to solve all philosophical problems. But even if I intended to solve all problems, then this would at most be a further intention with which I would write the thesis, but not my intention *in* writing it.<sup>13</sup> So while I can say that I am currently writing my thesis, or even that I am elucidating practical knowledge (even though I might fail in the end), it would go way too far to say that I'm solving all philosophical problems. Sometimes, however, the boundary between what we can and

---

I don't go over the edges of the paper. Or indeed, I can see that my pen has run out of ink.

<sup>12</sup>It is exactly this difference in aspect that explains why this thesis is called *Knowing and Doing*.

<sup>13</sup>For a clear explanation of the difference between 'intention with which' and 'intention in action', see [Moran and Stone 2008].

cannot claim to be doing is less clear. Anscombe argues that there isn't much to choose between 'She is making tea' and 'She is putting on the kettle in order to make tea' [§23], but what about 'she is buying tea leaves in order to make tea'? If she is just doing her weekly grocery shopping, we would be inclined to say that she isn't making tea while she is shopping. But what if she just told her friend that she will make some tea for her, finds out that she is out of tea leaves, and runs off to the store? Here it wouldn't feel strange to say that she is making tea (and knows this) while she is running to the store. And perhaps it wouldn't be entirely wrong to say of Wittgenstein that he was solving all the problems of philosophy, when he was working on the *Tractatus*. At least, he explicitly intended to.

In any case, what is clear is that if the movements an agent is currently making do not contribute to a goal she claims to have, this does not show they do not belong to a wider process that is aimed at reaching her goal. Hence we can *truly* say of our blind man that he is writing while his ink has run out and of our office clerk that he is making ten carbon copies when his current movements only produce eight of them.<sup>14</sup> The practical knowledge they have during the performance of their action still is undeniably factive—even the sudden onset of the apocalypse cannot change that.

## 2.5 Advanced carbon copy cases

In this section I will discuss two variants of the carbon copy case that some think provide a counterexample to practical knowledge, and are immune to Micheal Thompson's strategy. The first one was presented by Tim Henning on a recent conference on practical knowledge in Basel.<sup>15</sup> I think we can brush it aside quite easily, based on what we already established in this chapter. The second example is mentioned by Thompson

---

<sup>14</sup>It is interesting to note that Davidson himself does not use his carbon copy example to explicitly attack Anscombe. He uses it to argue for a negative answer to the question: is intending to act a belief that one will? Anscombe would certainly agree with him that the answer to this question is 'no'. Having practical knowledge is quite different from believing that you will perform some action. I believe that I will change the acoustic properties of any room I enter, but I do not intend such a change. As we will see in chapter 4, Davidson is also wrongly thought to argue explicitly against Anscombe in his famous article 'Actions, Reasons and Causes' [1963]. It seems to me that Davidson actually understood Anscombe very well. I think he could not accept her view entirely because of his Quinean ontological preferences. For Davidson there can be no such thing as a rationally structured event. Events are just occupiers of space-time (see [Lepore and Brian 1985, 162-157]). In his articles on action he actually did a tremendous job to preserve many of Anscombe's insights—like the idea that something is intentional under a description—despite his meager ontology.

<sup>15</sup>Practical Knowledge II; A workshop on the philosophy of action, July 9-11, 2010.

himself, but he doesn't spill much ink over it. I will say a bit more.

*Example 1:* Suppose that X hates his neighbour, and therefore wants to wreck his neighbour's car. Today X is on holiday in a foreign country and suddenly forms a random and irrational belief that he has his neighbour's car in front of him. Now assume that by an incredible coincidence, it really happens to be his neighbour's car. X starts wrecking the car. Now, he believes that he is wrecking his neighbour's car. And it is true that he is wrecking his neighbour's car. And finally, the belief brings it about that it is true that he is wrecking his neighbour's car. (If he did not believe that this was his neighbour's car, he would not be wrecking it in the first place.) Now, does X *know* that he is wrecking his neighbour's car? Tim Henning thinks that the man obviously fails to know that he is doing so.

Henning's thought supposedly is that knowledge does not only require true belief, it also requires justification, and X certainly has no justification for believing that it is his neighbour's car in front of him. Furthermore, it might seem that Thompson's reply cannot work because when X really believes that he has wrecked his neighbour's car once he is finished, he will not check whether he actually succeeded (unlike the guy who only made nine carbon copies where he intended to make ten), but will just go about his business.<sup>16</sup>

I think that this counterexample confuses practical with speculative knowledge.<sup>17</sup> I would agree with Henning that X does not (speculatively) know that his neighbour's car is being wrecked, because his belief that this is the case lacks justification. But this does not mean that X fails to practically know that he is wrecking his neighbour's car. Practical knowledge does not need justification through evidence, the world has to fit it. As we saw above we can know that we are writing our name, even if all that currently happens is that an inkless pen is moving over a piece of paper. An analogy to Henning's case would be that the name miraculously appears on the paper, even though the pen has ran out of ink — due to some strange properties of the paper's surface perhaps.

*Example 2:* Suppose our office clerk is under some strange mafia threat. He must succeed to make ten copies in one try, no checking, and he is dead if he doesn't manage to do so. Now we cannot reply that this clerk still has practical knowledge that he is making ten copies because his first try to make ten copies was only a phase of his action. There simply are no second or third tries. And indeed Thompson's reply to this example is that the office clerk lacks practical knowledge:

---

<sup>16</sup>And even if he checks, he will find out it is his neighbour's car, and he will just go about his business.

<sup>17</sup>Or at least it confuses Anscombean practical knowledge with what Velleman and Setiya take it to be.

for him, the making of the inscription is like the buying of a lottery ticket. You can say he made ten copies intentionally if you like, but it will not be an illustration of the topic of Anscombe's book, any more than lottery-winning is when you bought the ticket with that aim. [Thompson forthcoming]

What Thompson suggests is that doing something with a certain aim, is not enough for describing what happens as the intentional action that is the attainment of that aim (e.g. describing buying a lottery ticket as 'lottery-winning'). Now what is the difference between these cases and the ones in which is perfectly fine to e.g. describe the onion cutting as 'risotto making'?

As we saw at the end of section 2.4 there is a (not always definite) boundary between what we can and cannot know to be doing. One of the criteria for truly claiming to be doing A by doing B, is that one sees B as a means that promotes A. Onion cutting clearly is a means towards making a risotto, but buying a lottery ticket isn't a means to winning. Of course it is a condition for winning the lottery, but it isn't a (reliable) means to an end, that is unless you are buying all the existing lottery tickets. But in the normal case 'winning the lottery' isn't something we intentionally do, it rather is something that happens to one. Whether an agent sees something as a means through which he can intentionally achieve an end, depends on at least two factors. The first is whether the means could promote the end at all, the second is the agent's competence in producing the end via the means. If you see me hitting a golf ball with a putter at a distance of two meters from the hole, I would not claim to be putting the ball by hitting it with the putter— even if the ball ends up in the hole.<sup>18</sup> Tiger Woods, however, might confidently and correctly claim that he is putting the ball— even if he somehow fails. If our confidence that we can achieve a goal by means of our action is below a certain threshold, we express this by saying that we are *trying* to do such and such. I would say that I am trying to put the ball, and if the ball were twenty meters from the hole even Tiger would have to resort to such language.

I would argue that the break between what we can and cannot (practically) know to be doing precisely coincides with the threshold for claiming that one is doing it instead of trying to do it. If you were to ask me what I was doing while I am filling out a lottery ticket, and I non-jokingly replied 'I am winning the lottery', you would suspect that I was cheating. If I somehow already knew the lucky numbers I would definitely see filling them out on a lottery ticket as a means towards winning. On the other side of the

<sup>18</sup>I might, however, say that I am putting it if were granted a high number of tries.

spectrum is a case where you ask me what I am doing and I reply ‘I’m trying to move my finger’, this would lead you to suspect that I am recovering from a paralysed hand or the like.

Let us return to the office clerk who has only one go at making ten copies. If we were to ask him what he was doing he would probably say that he was *trying* to produce ten copies. He does not know that he is making ten copies, but only that he is trying to do so by pressing very heavily on the paper. We could however imagine a clerk who is very proficient at making carbon copies, every day he makes hundreds of them. He cannot only make ten copies in one go, he can easily make twenty. If this clerk were under a mafia threat to make ten copies, there would be nothing wrong with saying that he does have practical knowledge that he is making ten copies.

I think we can now see that the example of the clerk under a mafia threat isn’t a problem for Anscombe’s doctrine of practical knowledge; it might even illustrate it. The less proficient clerk does not know that he is making ten copies, and we therefore will not say that he is intentionally making ten copies, but at most that he is trying to do so. The proficient clerk however, does know that he is making ten copies, and there is no problem whatsoever in saying that he is doing so— with or without a mafia threat.

The knowledge we have of what we can and cannot succeed in is speculative knowledge. I can practice very hard at rifle shuffling a deck of cards and observe whether I am successful. When I have achieved a high success rate I will confidently claim that I can rifle shuffle. Now rifle shuffling becomes a means for me to change the order of the cards in a deck. This means that I can have practical knowledge of the fact that I am rifle shuffling in order to mix up the cards. In this way speculative knowledge of our skills can enter into our intentional actions. The next chapter will focus on this function and other roles speculative knowledge plays in action.

## Chapter 3

# Implications for epistemology

*[Hume] confidently challenges us to “produce some instance, wherein the efficacy is plainly discoverable to the mind, and its operations obvious to our consciousness or sensation”. Nothing easier: is cutting, is drinking, is purring not ‘efficacy’?*

– Elisabeth Anscombe

**I**N the previous chapters I defended the view that intentional actions are essentially constituted by a peculiar kind of knowledge: practical knowledge. Practical knowledge of one’s own action has the distinguishing feature that the world should fit it, whereas speculative or theoretical knowledge should fit the world. Hence we cannot acquire practical knowledge by observing the world as it is. This brings us to Anscombe’s famous claim that we know what we are doing without observation. One objection to this idea might run as follows: Anscombe cannot be right because performing almost all of our actions depends on sensory input. Take her example of writing your name with your eyes closed. Here the subject still has tactile, proprioceptive and kinaesthetic information, and such information is needed for writing with one’s eyes closed. As Kevin Falvey puts it:

Surely vision or touch or proprioception is necessary for one to be writing at all, not just for one to be writing legibly. It is difficult to imagine a person doing anything intentionally, even raising his arm above his head, if he is deprived of all forms of sensory input. The targeting and tracking of any bodily movement depends on a continual flow of sensory information of some kind or other. Action and perception are inextricably linked. [Falvey 2000, p. 30]

Falvey concludes that such knowledge, along with knowledge of one’s own skills and abilities and knowledge of various background conditions, is “necessary for one to know



that one is in fact doing what one intends to be doing” [Falvey 2000, p. 31]. Thus the knowledge of our own intentional actions cannot be entirely non-observational and Anscombe must be wrong.

I think that Falvey is right in stressing the importance of the role of perception, background knowledge, and knowledge of one’s skills in intentional action. They are indeed necessary for many of our actions. He is, however, wrong in thinking that the need to have such knowledge in order to act at all, constitutes a problem for Anscombe’s doctrine of practical knowledge. It is the aim of this chapter to fit these non-practical kinds of knowledge into a roughly Anscombean framework, and to show some of the consequences this has for our more general epistemological views. I will start section 3.1 by showing why Falvey’s objection fails, and continue by sketching some of the roles of (non-practical) knowledge in action. The picture that will emerge has many implications for our philosophical understanding of knowledge acquisition. I will defend in section 3.2 that actions can play an important role in the obtaining of knowledge, in fact it will be shown that actions are necessary for the acquisition of so-called causal knowledge (§3.3). While the acquisition of knowledge and the role of knowledge in action are indeed epistemological topics, most of the current debates in analytic epistemology focus on the analysis of the concept of knowledge. Section 3.4 will focus on the debate about the dependence of speculative knowledge on the practical context of the observer. I will show how a clear understanding of the ties between knowledge and action can contribute to that debate.

### 3.1 Non-practical knowledge in action

Consider Falvey’s argument against the non-observational character of knowledge of our own actions. He stresses that all of our bodily movements depend on sensory information. The senses are necessary to perform actions, so we cannot *know* that we are doing anything without observations, since we cannot *do* anything without observation. Falvey’s talk about knowing “that one is in fact doing what one intends to be doing”, shows that he is committed to the gap between what we intend and what we are actually doing. But as we saw in the preceding chapters, Anscombe does not think there is such a gap; “we do what happens” [§29]. Remember the blind man whose ink ran out, he was still *writing* his name and knew this, although he did not succeed in producing his name on the paper. “[H]ere is where the eyes are useful; but the essential thing he

does, namely write such-and-such is done without the eyes” [§29]. Falvey’s claim is that Anscombe “cannot be right” [Falvey 2000] because the “essential thing he does”, the writing, *essentially* involves perception of some sort. But it is a mistake to think that if an action requires certain knowledge for its successful execution, the same knowledge is required for knowing that one is in the process of conducting the act. Making a risotto requires knowing how long the rice should cook, but when I am still cutting the onions I already know that I am making a risotto, even if I have not yet consulted my cookbook to look up the rice’s cooking time. Of course Falvey is right in claiming that our actions depend on perception, someone who knows by observation that her pen has run out of ink cannot write her name, and hence she cannot have the practical knowledge that she is writing her name. Practical knowledge and intentional actions are inseparable, you cannot not have one without the other. So, since actions depend on sensory input, having practical knowledge also depends on it. But this does not show that the practical knowledge itself is observational knowledge. Practical knowledge cannot be observational, since it specifies the description under which an action is intentional. Such a description cannot be gathered from perceiving the world, because it is the world that should fit the description.

It is important to stress that this claim that practical knowledge is entirely non-observational does not entail that sensory information isn’t required for action. The varieties of non-practical knowledge employed by acting human beings are many. And the boundaries between them are vague, so it is hard to categorize them all. I do think however, that it is very helpful to consider some of the ways in which theoretical knowledge enters into intentional action. From Anscombe and Falvey we can gather that intentional action may require theoretical knowledge for at least two purposes. Falvey stresses that action presupposes certain knowledge in order to exist at all. Anscombe shows that we may require knowledge to guide the action as it proceeds. We might call the former kind of knowledge ‘enabling knowledge’ and the latter kind ‘guiding knowledge’.

According to Anscombe, observation can guide one’s action as it proceeds. She talks, for instance, about observation as an “aid in writing” [§29] which prevents one from going over the edge of the paper, or overwriting already written lines. The category of guiding knowledge seems to be very broad. It ranges from unconscious know-how that constitutes skills (e.g. my right foot is in front of my left foot, so I move my left foot in order to walk down the street), to very specific procedural knowledge applied in complicated actions (e.g. I should adjust the trajectory of the proton in the large hadron collider in the way my computer data specifies, in order to create a collision). It is a topic

that is studied by psychologists, but it is of interest to philosophers as well. Philosophers of science who think about the way in which scientists set up experiments, for instance, should take note of the role guiding knowledge has in our actions.

Falvey focuses is on the enabling role of speculative knowledge in speaking about vision or touch or proprioception as “necessary for one to be writing at all”. It is clear that we need some idea of where our pen-holding hands are in order to go and write at all — although we might not need to entertain such an idea consciously. Something similar holds for almost all intentional actions. Maintaining balance when standing upright, for instance, requires the vestibular system and at least two of the three senses of proprioception, touch and vision. Hence the Romberg-test for drunken driving: test subjects are asked to stand up straight and close their eyes, blocking vision. Sober individuals should be able to stand still because they still feel their feet touching the ground and can sense the position of their limbs. Drunken subjects will, however sway or even topple, because alcohol is known to impede proprioception.<sup>1</sup> But the category of enabling knowledge is broader than just the knowledge of the position of one's own body or the objects around one. One needs to know, for instance, that one is playing baseball in order to do such a thing as stealing a base.<sup>2</sup> Another kind of knowledge that is constitutive to action is the knowledge that the world will remain relatively stable. Here is an example from [von Wright 1974, p. 42]: I am confident that I can turn over the piece of paper in front of me. In a windy room I would be less confident. But in a universe where pieces of paper flip over on their own all the time, there would be no such action as ‘turning over a piece of paper’.

It is clear that the picture sketched above of the role of speculative knowledge for intentional action needs more development. But for now it is enough to have some grasp of the matter. What remains unclear, however, is *how* speculative knowledge can enter into our actions. I suggest that the theory of practical reasoning, explained in section 1.5, might contribute to clarification of this issue. Practical reasoning is reasoning about means and ends with a view to action. When something is deduced to be a means to an end that is seen as good (wanted) the goodness is, *ceteris paribus*, transferred to the means. Practical reasoning terminates in an action that is done to produce the means deduced from the end. But the judgement that something is a means to an end is, or is

---

<sup>1</sup>See [Nieschalk et al. 1999].

<sup>2</sup>Interestingly enough, the knowledge that I am playing baseball right now is most probably practical knowledge of what I am doing. Which suggests even further entanglement of the different types of knowledge with each other and with action. More on the enabling or causal role of practical knowledge will follow in 4.5.

based on, speculative knowledge about the world.<sup>3</sup> It is a matter of fact about risotto that cutting onions is a step towards cooking it. I need to know where my camera is in order to conclude that going upstairs is a means to fetch it.

Speculative knowledge about means and ends, employed in practical reasoning, can have the role of both enabling as well as guiding knowledge. If I want to read other people's minds, I would have no clue whatsoever as to how I should start learning this, i.e., I cannot deduce any means that could help me accomplish it. Therefore I cannot even start to learn mind reading. But as soon as I can think of any means that could get me closer towards my goal, I can start doing it. If I want to prove Goldbach's conjecture for instance, it might be a good idea to go and study mathematics — even though it might be a bit far fetched to say of a first year undergraduate who is doing her assignments that she is proving that every even integer greater than two is the sum of two primes. On the other hand, when I find out more about mathematics, I will perhaps find new means that can guide my action towards a successful completion, and thus my studying mathematics indeed is a means towards the end of proving Goldbach's conjecture.<sup>4</sup>

Since putative speculative knowledge is indeed *speculative* it can fail to fit the facts. It might turn out that it wasn't knowledge after all, it was mere belief. What happens when a piece of speculative knowledge, which is employed in action, fails? I can clearly be wrong in thinking that my camera is upstairs, and therefore in thinking that going upstairs is a means to fetching it. The first thing to note is that when my knowledge of means and ends fails, my action is likely to be unsuccessful. If I think my camera is upstairs, while it is actually in the basement, I will fail to fetch it. But if someone claims that she is going upstairs to get her camera and he fails to come down with it, it would be wrong to say that she was lying. We can exploit the difference in aspect the English language possesses to say that our agent was *fetching* the camera when she went upstairs, but actually failed to *fetch* it. The case of writing without ink, discussed in chapter 2, is similar to this one: there the blind man's judgement that moving his pen in a special manner was a means to writing his name simply was wrong. Even in a case where someone really believes that she can learn to read minds by massaging her temples in a certain manner, it wouldn't be wrong to say that she is learning to read minds, although she will never have learned mind reading. Thus the following seems to

---

<sup>3</sup>The difference between practical reasoning and theoretical reasoning, lies in the fact that practical reasoning is done with a view to action. Without that, it would simply be theoretical reasoning about means and ends.

<sup>4</sup>It was a major step for Andrew Wiles when he realised that proving the Taniyama-Shimura conjecture for semi stable elliptic curves was a means to proving Fermat's last theorem.

hold: we usually need to be correct in our judgments about certain means-end relations in order to perform *successful* actions, but we don't need to be correct in order to engage in action.

## 3.2 Knowledge acquisition

The previous section showed some of the ties between speculative knowledge and our capacity to act. But can action also play a role in the *acquisition* of such knowledge? The empiricist doctrine is that all knowledge must be obtained through the senses.<sup>5</sup> This section wants to provide some reasons to doubt such a strict claim. The next section will illustrate the role actions can play in the acquisition of a special kind of knowledge: causal knowledge.

In section 3.1 we saw that when speculative knowledge fails my action will be unsuccessful. It is enough that I believe (with a fair degree of certainty) that my camera is upstairs, in order to start fetching it, but I will not come down with it if I was wrong. If all this is true, it is tempting to draw the following conclusion: When an action is successful, the means-end beliefs employed in the action were correct. Moreover, the success of our actions also *justifies* our beliefs about the relevant means and ends. By using the crude justified true belief model of knowledge, we might say that the successful completion of actions can turn belief into knowledge. Hence, successful action yields knowledge.

Unfortunately things aren't this easy. I might be entirely wrong in my means-ends judgements but still achieve my ends, if only by sheer luck. It would be wrong however, to conclude from this that successful action never helps in validating our beliefs. That would be plain scepticism, similar to scepticism about our (peckishly misleading) senses. I would like to weaken the idea that successful action yields knowledge. My suggestion simply is that successful actions can provide justification for certain beliefs, without claiming that all knowledge stems from successful action, or that successful action always yields knowledge.

The above suggests a quite static and overly simple picture of agents that justify the beliefs they already possess with the help of intentional action. I believe that actions might also play a role in the formation of new beliefs, and that what really goes on when humans acquire new knowledge is far more complex. What often seems to take place

---

<sup>5</sup>Quine famously states that "whatever evidence there is for science is sensory evidence" remains to be an unassailable tenet of empiricism [Quine 1969, p. 73].

is an interaction between agentive and observational powers. I play around with some objects, and look what happens: I confirm what I saw by repeating whatever I did, or just by performing a different action etc. What I have said is of course still a long way from a full fledged account of the role of actions in knowledge acquisition. The development of such an account would take a lot more philosophical and scientific research. I do hope, however, to have shown that the close connection between speculative knowledge and (successful) intentional action provides a reason to think that it is *possible* for actions to play a role in the acquisition of knowledge. A truly persevering empiricist might stress that this section does not show that we acquire knowledge by acting, it shows that we gain knowledge by *observing* whether our actions are successful or not. If such an empiricist is right, then it might turn out that action isn't that important in the acquisition of knowledge after all. But actions will be important for knowledge acquisition if it turns out that there are cases where action itself yields knowledge. The next section will argue that we can find such a case in the acquisition of causal knowledge.

### 3.3 Causal knowledge and the ability to act

At one point in *Intention* Anscombe wonders whether the knowledge of means and ends employed in practical reasoning is causal knowledge. Here is Anscombe:

[D]oes this mean that people must have notions of cause and effect in order to have intentions in acting? Consider the question 'Why are you going upstairs?' answered by, 'To get my camera'. My going upstairs is not a cause from which anyone could deduce the effect that I get my camera. And yet isn't it a future state of affairs which is going to be brought about by my going upstairs? But who can say that it is going to be brought about? Only I myself, in this case. It is not that going upstairs usually produces the fetching of cameras, even if there is a camera upstairs [...] In order to make sense of 'I do P with a view to Q' we must see how the future state of affairs Q is supposed to be a possible later stage in proceedings of which the action P is an earlier stage. [§22].

So our means-end judgements aren't judgements of *general* causality (i.e. of the form "going upstairs causes camera fetching"), but we might say that they are judgements of *singular* causality (i.e. "this particular going up the stairs brings about my fetching of the camera"). Both Hume and Davidson, however, have suggested that singular causality implies general causality. Hume writes that:

we may define a cause to be an object, followed by another, and where all the objects

similar to the first, are followed by objects similar to the second. [Hume 2000 [1748], sec. VII, part II]

Davidson notes that it might not be apparent which generalisation is instantiated by an episode of singular causation, but that there nevertheless must be a covering law, although we might not be able to specify it [Davidson 1967]. In her inaugural lecture at Cambridge University, Anscombe [1971] points to the lack of arguments Hume and Davidson have for supposing that singular causality must be analysable in terms of general causal statements:

It is over and over again assumed that any singular causal proposition implies a universal statement running ‘Always when this, then that’ [...] Even a philosopher acute [...] as Davidson, will say, without offering any reason at all for saying it, that a singular causal statement implies that there is such a true universal proposition— though perhaps we can never have knowledge of it. Such a thesis needs some reason for believing it! [Anscombe 1971, p. 104]

Anscombe’s own suggestion is that we can only understand general causation if we first understand singular causal processes.<sup>6</sup> She argues that we can only have the concept “cause” in a language in which there already exist many other causal concepts, like “scrape”, “push” and “knock over”. To learn the concept “cause” we have to see many examples of particular situations in which a causal process is going on: someone knocks over a glass of milk, a boulder is pushed over the edge, etc. It is very interesting for the purpose of this chapter that the verbs like “scrape” and “push” are *action verbs*. Anscombe’s suggests that we first need to understand what it is to push something over ourselves, before we can apply this to nature (i.e. the wind’s pushing something over). This thought forms a basis for the increasingly popular interventionist account of causation.<sup>7</sup>

Von Wright was the first philosopher to develop a well worked out manipulability account of causation.<sup>8</sup> I will explain his position with the help of an example. Imagine the following setup. We have a magnifying glass, which is aimed at a piece of paper. If we observe what happens several times we find that we first see a brown spot emerging

---

<sup>6</sup>This is not the same as claiming that every statement of general causality should be analysed in terms of singular causal statements. In fact there are good reasons to think that general and singular causality aren’t reducible to each other in any way, see [Hitchcock 1995].

<sup>7</sup>Anscombe also disagrees with Hume that we are unable to see causal relations, see [Anscombe 1971]. Recent evidence from the field of psychology suggests that Anscombe is right and Hume was wrong about this, see [Dullstein 2007].

<sup>8</sup>See, [von Wright 1971] and [von Wright 1974].

and after that we see a hole. Because we see this regularity we might wonder: does the brown spot cause the hole? Observing the same experiment again and again will not help us to solve this question, so maybe we need to interfere with what is going on. We can for instance push over the magnifying glass at a moment at which the brown spot has already formed but when there is no hole yet. In this way we can use action to discover that the brown spot does not cause the hole.<sup>9</sup>

Someone who wants to play down the role of actions might suggest that we could get lucky if we run the experiment a few more times. The wind for instance might blow over the magnifying glass, which would also allow us to conclude that the brown spot does not cause the hole. Interestingly, von Wright would object to this. His problem is that we cannot make sure that the wind is not also a counteracting cause to the formation of holes in papers. Perhaps brown spots do cause holes, unless wind blows on them. But how can we be sure that the same does not hold for our actions? Pushing over nearby magnifying glasses might also be a counteracting cause to hole formation. What makes actions so special? Von Wright's answer is that actions are special because they are conceptually connected to causality.

Let me explain. One of the problems with Hume's definition of causation as "an object, followed by another, [...] where all the objects similar to the first, are followed by objects similar to the second" is that it fails to discriminate between causation and mere regularity. High tide always follows low tide, but it isn't caused by it. One of the differences between A regularly following B and A causing B is that in the former case B might still occur if A doesn't whereas in the latter case prevention of A will prevent B from occurring. Causation thus contains a counterfactual element.<sup>10</sup> Saying that wind caused the magnifying glass to topple over implies that if there hadn't been any wind the magnifying glass would still be standing. According to von Wright actions also contain an element of counterfactuality. We can only act when we suppose that the world would have remained the same— in certain respects— if we hadn't acted. Remember the example of the flipping over of pieces of paper: we can only flip over a piece of paper if we are confident that it does not do so by itself. To put it differently: we might never know for sure whether a counterfactual claim like "*p* would have remained

---

<sup>9</sup>Recently evidence has emerged that interventions indeed play a large role in causal learning, see [Gopnik and Schulz 2007].

<sup>10</sup>Hume already noticed this: "[when an object causes another, then] if the first object had not been, the second never had existed" [Hume 2000 [1748], sec. VII, part II]. It is, however, hard to see how he could account for such a tight connection if causality consists of nothing but one object regularly following another.



absent had I not produced it” is true, but we can have trust in such a counterfactual claim and only when we confide in it are we able to act.

If we return now to our magnifying glass example we can see why there is a difference between the cases of human and non-human intervention. In the case of non-human intervention we know nothing about the possible causal side effects. But in the case of human intervention, von Wright has it, we know more. We know — or are at least confident in the fact — that the world would have remained stable if we hadn’t acted. From this we might conclude that ‘*p* causes *q*’ means that we could produce *q* by bringing about *p*. And von Wright is often thought to have explicitly defended this conclusion.<sup>11</sup> The problem with this conclusion is that it is hopelessly anthropocentric. We cannot produce gravitational collapse in white giants in order to create black holes, it would however be wrong to say that gravitational collapse cannot cause the formation of black holes.

That von Wright’s point about the conceptual connection between ‘action’ and ‘causality’ was more subtle shows from his reaction to charges of anthropocentrism. His reply was that he did not mean to suggest that causal processes would not occur if humans did not exist, but only that there wouldn’t be a concept of causality without humans [von Wright 1974, lecture 4]. I think von Wright’s idea can be made clearer if we look at the later sections of *Intention*. There Anscombe writes:

The term ‘intentional’ has reference to a *form* of description of events. [...] Events are typically described in this form when ‘in order to’ or ‘because’ (in one sense) is attached to their descriptions. [§47]

It is very interesting that we can use terms that primarily occur in the descriptions of intentional actions to describe natural events. We can speak of the wind as picking things up and putting them down again and of a cleft in the rocks as holding something. ‘Picking up’, ‘putting down’ and ‘holding’, all get their sense from usage in descriptions of intentional actions. In fact some descriptions of natural events are *formally* descriptions of executed intentions: ‘the oak drops its acorns in order to produce offspring’. While Anscombe only talks about applying forms of description to events in the world, it seems clear that the fact that we can apply such descriptions allows us to understand the world in a particular way. It especially allows us to understand the natural teleology of living

---

<sup>11</sup> Although I defended in [van Miltenburg 2008] that von Wright really is defending the much weaker claim that we can distinguish between causal relations and accidental regularities by using the notion of action.

organisms.<sup>12</sup> I believe that von Wright is making a similar point about causation. We can only understand events as causes and effects of each other because we ourselves can produce changes in the world through action.

If we return to the example of the magnifying glass that was blown over by the wind, I think we might relax our standards a bit. We can allow that the wind blowing over the magnifying glass *does* allow us to conclude that the hole wasn't caused by the brown spot, but only because we can see the wind as an intervention on what would naturally happen. And we can only understand what an intervention is because *we* can interfere with the course of nature. The notion of 'natural intervention' plays a central role in Woodward's [2003] increasingly popular interventionist account of causality. It is often suggested that Woodward has rid the manipulability theory from its anthropocentrism. I agree, but want to add that his theory of natural interventions is only possible because we can understand the world in a typically human way. Some anthropocentrism has to remain, and rightly so.

The thought that our ability to act allows us to understand the world in a particular way, pertains to much more than just causality. We can, for instance, acquire modal knowledge of the possibilities and impossibilities of objects surrounding us by actively manipulating them. Just like actions can help us to test causal hypotheses, we can test modal hypotheses by means of acting. If I wonder whether this chameleon can change color, I can find an answer by placing it on several different colored surfaces. The counterfactual element in action that secured its conceptual connection with causality, might also secure a special connection between modality and action. I am well aware that more analysis of both modality and action is required before such a claim can be bolstered. But for the purposes of this thesis it is enough to note a certain resemblance between causal and modal knowledge acquisition. The similarity between the two isn't limited to the fact that knowledge of both singular causality and of concrete possibility can be gained via manipulation of objects in our environments: it has recently been suggested that knowledge of such concrete (or *real*) possibilities is epistemically prior to knowledge of more abstract modalities (like metaphysical or natural modality) [Müller, T. 2010], much like knowledge of singular causation is prior to an understanding of general causation. It seems that our ability to act again plays a role in the ways in which we can understand the world. It allows us to view events in the world as causally related, and it might also allow us to discover the possibilities and impossibilities of the objects and substances

---

<sup>12</sup>This is worked out in [Thompson 2008, part I].

we engage with. But the implications of the capacity of intentional action for our understanding might go even further. It seems to me that our ability to act might also enter into our ability to discriminate and individuate objects and in our ability to establish the existence of hypothetical objects (cf. Ian Hacking's defence of realism about electrons: "If you can spray them, then they are real" [Hacking 1983, p. 23]).

### 3.4 Action in contemporary epistemology

Where the previous sections highlighted the role speculative knowledge plays in action, and the role actions play in the acquisition of (special kinds of) knowledge, most discussions in contemporary analytic epistemology focus on the analysis of the concept of (speculative) knowledge. This section aims at showing that a thorough understanding of the relations between knowledge and action might also have an impact on these discussions. Familiar questions in contemporary epistemology are: 'What sort of justification does knowledge require?', 'Can someone know  $p$  and  $p \rightarrow q$ , without knowing  $q$ ?', or 'Are knowledge ascriptions relative to conversational contexts?' Recently a new question has surfaced: Does whether or not someone knows that  $p$  when she truly believes that  $p$  only depend on her justification for her belief (her evidence, whether she used reliable methods etc.)? Or does it depend also on her *practical* situation (what is she currently doing, how high the stakes are, etc.)? Traditionally all epistemologists would answer this question in the negative, but recently subject sensitive invariantists have argued for the dependence of knowledge on action-contexts. Fantl and McGrath [2002], Hawthorne [2004] and Stanley [2005] all argue along different lines for such a conclusion. But the three accounts share two features: 1) they focus on scenarios in which we have (or are supposed to have) the intuition that we lose or gain knowledge as our practical contexts shift, and 2) they offer remarkably little discussion of the nature of action, or of the conceptual connection between action and knowledge, which is supposed to underlie our intuitions.

Both Stanley and Hawthorne summarise the conceptual connection between knowledge and action in the following slogan: 'one should act only upon what one knows.' This slogan is supposed to be bolstered by the fact that ordinary speakers justify their actions by means of knowledge claims all the time. When someone asks 'Why did you go left rather than straight?' I can reply with: 'Because I knew it was the shortest direction to the restaurant.' The dependence of knowledge on practical context follows

from a formalized version of the slogan:  $K(p) \rightarrow A(p)$ , read as: If one knows  $p$  then it is appropriate to act on  $p$ . If we contrapose this we get:  $\neg A(p) \rightarrow \neg K(p)$  i.e., if it isn't appropriate to act on  $p$  one doesn't know  $p$ . Whether it is appropriate to act on  $p$  varies according to the practical environment of the subject and hence, the subject's knowledge also depends on it. Hawthorne illustrates this with the example of an agent who justifiably believes that she is going to Blackpool next year. When this agent is in a bookshop it might be appropriate for her to act on her belief and buy a travel guide that highlights Blackpool's main attractions. In such a case we should say that she knows that she is going to Blackpool next year. But what about the following scenario: our agent is offered life insurance, and she declines based on the belief that she will not die before next year, because she will be going to Blackpool then. In this scenario it is clearly inappropriate that the agent act on the belief that she will be going to Blackpool next year, and therefore we shouldn't say that she knows that she will be going to Blackpool next year. The conclusion that our knowledge may shift along with our practical environment seems inevitable.

In the context of this thesis it seems that Hawthorne's agent possesses *practical* knowledge of the fact that she is going to Blackpool in both scenarios. There is no contradiction between being unsure about whether one will die before next year and possessing such knowledge. This has to do, of course, with the difference in direction of fit between practical and speculative knowledge. The judgement that I will not die before next year can turn out to be wrong, but if I have practical knowledge that I will go to Blackpool and I fail to arrive there because of my premature death, the mistake is in the performance. With this in mind the problem with the agent's declination of life insurance is not that she inappropriately acted on her belief that she will be going to Blackpool. It rather is that she acted on a *speculative* belief (that she would not die) that was deduced from her *practical* knowledge. Such a switch from practical to speculative also requires a switch in justification. Practical knowledge is justified by reasons for action, whereas speculative knowledge needs to be justified by evidence. There is nothing wrong with believing that a certain thing will happen, because you intend to make it happen. Such belief might even amount to knowledge if you are justified in thinking that you will be successful (the task might be easy, or you might be really good at it), and if you in fact will make it happen.<sup>13</sup> The problem is that one should not conclude that *nothing* can stop

<sup>13</sup>So Falvey, Velleman and Paul are right that there is a link between intentions and speculative knowledge. The problem is that they think that this link explains the only way in which we can acquire knowledge about our own actions, where I argue that Anscombean practical knowledge gives us direct access to

you from making it happen, because you think you will be successful. This is the error our agent made. Once it is clear that the agent's knowledge that she will go to Blackpool is practical knowledge, the conclusion that a knowledge shift occurs when the agent goes from the bookstore to the insurance broker is no longer inevitable.

What I have said is of course far from a refutation of subject sensitive invariantism—which would fall outside the scope of this thesis. Hawthorne also offers different examples that are not as easily analysed in terms of practical knowledge, and Stanley's focus is on cases where an agent's action is inappropriate not because of a wrong deduction from current knowledge, but simply because the stakes are too high. Nevertheless, I hope to have shown that a greater awareness of the intricacies of the relation between knowledge (both of the speculative and the practical variety) and action, can have a profound impact on arguments put forward in favour of subject sensitive invariantism. I therefore want to plea for restraint in drawing grotesque conclusions about the impact of action-contexts on knowledge, before one has a well developed account of the connections involved. I hope that this chapter has contributed to taking the first steps in the construction of such an account.

## Chapter 4

# Implications for Action Theory

*[Davidson] speaks of the possibility of 'wrong' or 'freak' causal connexions. I say that any recognizable causal connexions would be 'wrong', and that he can do no more than postulate a 'right' causal connexion in the happy security that none such can be found.*

– Elisabeth Anscombe

**W**HERE the previous chapter suggested that practical knowledge might play a role in epistemology, this chapter argues for the reintroduction of the concept of practical knowledge in contemporary action theory. All action theory aims at the explanation of intentional actions. This is a very hard nut to crack because of the peculiar dual status of intentional actions. On the one hand actions are rationalizable entities (like decisions), we can ask what reasons someone has for performing an action. On the other hand they are happenings in the world (like earthquakes), we can ask how they come about. A full explanation of intentional actions should account for both aspects. I will argue that most contemporary action theories (in fact all the accounts I know of), focus on one of the two aspects and fail to explain the other. Causal action theory can perfectly explain how actions come about, but has difficulty in rationalising actions, whereas non-causal action theory cannot explain the happening of actions at all. My suggestion will be that an action theory based on practical knowledge might not suffer from this problem because practical knowledge can explain both aspects of intentional actions in that it is the *cause* of what it *understands*. This chapter is set up chronologically: I will start by discussing neo-Wittgensteinian theories of actions and their shortcomings (§4.1), and continue by explaining how causal action theory tries to overcome these shortcomings (§4.2). In section 4.3 I will argue that causal action theory is fundamentally flawed,

because it explains the causation and the rationality of actions separately and therefore cannot overcome problems of causal deviance. Recently some contemporary non-causal action theories have been developed, but unfortunately they also fail to fully explain actions. This will be discussed in §4.4. In the final section I will argue that the outlook for a theory of action based on practical knowledge is bright since it might not suffer from the flaws of theories discussed earlier.

## 4.1 Neo-Wittgensteinian action theory

Wittgenstein notoriously and rhetorically asked: “[W]hat is left over if I subtract the fact that my arm goes up from the fact that I raise my arm?” [Wittgenstein 2001 [1953], §621] This question is supposed to bring out how tempting it is to think of the intentionality of an action as something extra (a psychological occurrence like an ‘act of the will’) that needs to be present on top of the action. Wittgenstein argues that the view that actions require an act of the will to be intentional is absurd. If acts of the will exist how do I perform them? “It comes when it comes, and I cannot bring it about. [...] I can’t will willing; that is, it makes no sense to speak of willing willing.” [Wittgenstein 2001 [1953], §611-613]. Anscombe perfectly expresses the absurdity as follows: “the only sense I can give to ‘willing’ is that in which I might stare at something and will it to move”[§29]. There is no intermediate process that needs to be completed before I can perform an action. “I don’t need to wait for my arm to go up—I can raise it” [Wittgenstein 2001 [1953], §621].

In the wake of such remarks, and Wittgenstein’s general scepticism about the internal mind, many so-called neo-Wittgensteinian philosophers tried to explain actions without referring to mental phenomena.<sup>1</sup> Their idea is that when we want an explanation of a certain action, we want to be provided with an interpretation. When X’s behaviour seems strange, we want to know what reason X has for such behaviour. We should not look for X’s reasons by enquiring into the content of his mind, instead we can bring out his reasons by redescribing X’s actions in a certain way. Let’s look at an example from Melden [1967]. A man driving a car raises his arm. Puzzled by this we might wonder what reason he has for raising his arm. This puzzlement might be alleviated when we redescribe the raising of the arm as an act of signalling. The man thus raised his arm

---

<sup>1</sup>Examples of such neo-Wittgensteinian thought are found in: Melden [1967] and Ryle [2002 [1949]]. I do not include Anscombe among *these* neo-Wittgensteinians, as is often done, because she *does* point to something cognitive (albeit not a mental happening) that causes the action, viz. practical knowledge.

in order to signal. Perhaps one is inclined to continue the investigation and wonder why the man was signalling. This can be explained by considering the circumstances in which the signalling happened. Supposing he was approaching a turn and his directional indicators did not work, it becomes quite obvious why he signalled.

Often, however, it does seem that we reflect on the contents of a person's mind in order to explain her action. We explain why X did A by saying that she *believed* that such and such was the case or by pointing out that X had a certain *desire*. Ryle famously argued that such mental conduct verbs do not refer to something hidden in the mind of the actor, instead they are tools to describe the "capacities, skills, habits, liabilities and bents" [Ryle 2002 [1949], p. 45], that constitute the dispositions of which our intentional actions are actualisations. Therefore, even if we seem to describe a person's mentality in order to explain his actions, we are not actually scrutizing the content of the agent's mind in search for something that accompanies intentional actions but not mere bodily movements.

There are two major problems with such a neo-Wittgensteinian approach to the explanation of actions. First, it seems that Ryle and Melden can merely show how certain reasons *justify* (or make sense of) the performance of some action, while they have no explanation of what it is to act *for* a specific reason. The second problem is that they cannot account for the first person authority that agents have over their reasons for action. I will consider both in turn.

Davidson [1963] considers Melden's example of the driver who raises his hand. He agrees with Melden that if the driver was approaching a turn he had reason to signal, and if his indicator lights did not work he had reason to raise his arm. What however has not been shown is that this reason for raising his arm was *the* reason why he did it. He might also have raised his arm in order to greet a friend who was walking on the side walk. Both greeting and signalling justify the arm-raising, but since Melden and Ryle do not want to enquire into the mind of the agent, they have nothing that can help them to adjudicate between the two cases. The upcoming turn and the friend on the sidewalk might both be in the context of the agent's action, and the agent might be both disposed to signal for corners and to greet his friends. Therefore Davidson argues that Melden and Ryle cannot provide a *full* explanation of intentional action.

As we saw in §1.3, when an agent performs an intentional action she cannot fail to know that she is doing so. And when she knows what she is doing she cannot claim that it was unintentional. As Wittgenstein puts it: "voluntary movement is marked by the absence of surprise" [Wittgenstein 2001 [1953], §628]. Next to the knowledge of



her own action, an agent is able to say what her reasons for her intentional action are as well. In other words: one has authority over what one is doing and why one is doing it. According to Crispin Wright the real problem for any Wittgensteinian psychology is to account for such first person authority [Wright 2001, p. 128].<sup>2</sup> We simply do not need to look at our circumstances in order to know why we raise our arm; similarly, we do not need to look at our own behaviour to find out what we believe. Our knowledge of such things seems to be direct, and neither Ryle nor Melden can deal with that. In the next section we will see how Davidson's causal account deals with both problems.

## 4.2 Causal action theory

In his 'Actions Reasons and Causes' [1963] Davidson argues that rationalization (explanation in terms of reasons) and causal explanation are not as different as they *prima facie* might seem. Rationalization is in fact a species of ordinary causal explanation.

According to Davidson explaining an action comes down to explaining a bodily movement that takes place at a certain spatio-temporal location. In order to *causally explain* such an event we have to look for the event (or the events) that is (are) its cause. But what kind of events are the causes of actions? We will find an answer if we look at rationalization, Davidson thinks. To rationalize someone's action we must cite or indicate her *primary reason*. A primary reason consists of a belief and a pro attitude. Davidson generally speaks of pro attitudes as wantings (in a very broad sense of wanting). Here is an example of how a primary reason can rationalize an action: X wants to turn on the light (she has a pro attitude to turning the light on). She believes that flipping the switch will turn on the light. Therefore, X decides to flip the switch and does, in fact, flip the switch. We can answer the question 'Why does X flip the switch?' with: 'because she wants to turn on the light'. Davidson stresses that we should take the 'cause' in 'because' seriously. Having a primary reason to turn on the light causes X to flip the switch. Or more generally, an agent acts *for* a certain reason when the reason causes her to act.

Let us return to the example about the driver in the last paragraph. It may be that the driver has a pro attitude towards greeting his friend and the accompanying belief that he could do so by raising his arm and also had a pro-attitude towards signalling and the accompanying belief that he can signal by raising his arm. In other words: he has

---

<sup>2</sup>Wright's discussion isn't limited to action theory but pertains to self-knowledge (via introspection) in general.

both reason to signal and to greet. Neo-Wittgenstinian action theory could go no further than this, but Davidson can fully explain the action by pointing out for which reason our driver acted, i.e. what pro-attitude-belief pair caused the arm to move. Davidson of course agrees that reasons justify actions — his point is that they do not *merely* do so, the true force of the ‘because’ in rationalizing explanations only comes out when we interpret it causally. When a reason causes the action, the causal role the reason has, fully explains why the bodily movement that is the action takes place. It is in this way that rationalization derives its explanatory force from causal explanation.

One might object that, since causal explanation works by pointing at *events* that are the causes of the *explanandum*, reasons must be events in order to carry any causal burden. It seems however that reasons like ‘wanting to express your views’ are quite different from events like ‘the toppling over of a crane’. The anomalous monism presented in ‘Mental Events’ [Davidson 1969] addresses this objection head on by claiming that although reasons and events *seem* different, reasons actually *are* events, mental events to be precise. What makes an event a mental or physical event is, according to Davidson, dependent on the vocabulary we describe it in. When we describe Jerry’s belief that he is a man as ‘Jerry’s belief that he is a man’ we describe a mental event, but we can also (very roughly at least) describe it as a physical event, e.g., ‘a complex of neuron firings in Jerry’s brain’. Davidson’s doctrine is that every mental event also is, or is supervenient on, a physical event (hence, anomalous *monism*). Reasons are complex mental events consisting of mental events like wantings and beliefs. From this it follows that reasons can cause physical bodily movements, because reasons themselves are also (neuro-)physical events.

Doesn’t this make a mystery of our mental powers? We are free-willed human beings that can decide for themselves what they want to do, and on what reasons they want to act. If Davidson is right that the mental events in our mind are describable as physical events, and that it are these physical events that do the actual causing, then there is the threat of treating our beliefs, desires and intentions as mere epiphenomena. In other words if we can reduce our mental events to physical events that cause and are caused by other physical events, then there is no room for a free will. It shows Davidson’s brilliance that he also could overcome this objection. To ensure that the mind stays anomalous (that mental events do not follow strict laws), Davidson argues that there is only token-token correspondence between mental and physical events. So the thought ‘I would like some cake’ sometimes corresponds to a brain-event of type *X* and sometimes to a brain-event of an altogether different type. There is no law that determines the

correspondence between mental and physical events (hence, *anomalous* monism). So for every occurrence of a mental event there is always a corresponding occurrence of a physical event, but it is not always the same *type* of physical event that corresponds to the same *type* of mental event. In this way Davidson can explain that an action is done *for* a reason if it is also caused by the reason.

The second problem sketched in §4.1 was that an action theory should account for the authoritative self-knowledge an agent has of his own actions and reasons. Davidson can deal with this because reasons are realized in the brain of the agent. An agent has authority because she is the only one who can introspectively access it.

Whether anomalous monism is entirely successful is still heavily debated. Nevertheless, most philosophers of mind — at least the ones that take human agency seriously — have adopted theories that diverge only slightly from Davidson’s anomalous monism. Some diverge from Davidson in their view on the nature of the causal relation between mental and physical. Fodor, for instance, abandons Davidson’s strict law requirement on causation, by allowing laws with *ceteris paribus* hedges [Fodor 1992], and Searle talks about a subspecies of efficient causation: ‘intentional causation’ [Searle 2003, pp. 40–45]. Others digress from Davidson’s supervenience relation between mental and physical events. Thus Kim wants to “reductively identify [a] particular mental state with its neuro/physical realizer” [Kim 2007], and Horgan abandons Davidson’s talk about events altogether and talks instead about “physical properties that “realise” mental properties on particular occasions of instantiation” [Horgan 1989]. Yet other philosophers deviate from Davidson in the types of mental states that are responsible for causing actions. Bratman [1987] thinks that intentions themselves cause actions, and as we already saw in §2.3, David Velleman [1989] thinks that the agent’s knowledge of her action efficiently causes it. But despite these points of divergence the idea that intentional actions are bodily movements that are efficiently caused by certain mental states remains an unassailable doctrine for almost all contemporary philosophers of mind and action.<sup>3</sup> Ruth Millikan [1993] even called this doctrine one of the few achievements of analytic philosophy.

### 4.3 Deviant causal chains

A problem for almost any causal account of anything is the problem of deviant causal chains. Here is an example of a deviant causal chain when it comes to action: I despise X,

<sup>3</sup>Other noteworthy proponents of causal action theory are Goldman [1979] and Dretske [1989].

which provides me with a reason to kill him. Suddenly I see X crossing the road in front of my car. I now think: ‘This is my chance to finally get rid of the bugger.’ The sudden entertaining of such a murderous thought startles me, causing my foot to accidentally hit the gas pedal. Here my despisal of X causes his death. But would we really say that I killed him intentionally? In this case, contra causal action theory, the reason causing the action is not equal to the action being done *for* the reason. Davidson himself was already aware of this problem and tried to solve it by claiming that the reason should not only cause the action. It should cause it *in the right way*.<sup>4</sup> But what could such a right way consist in? Is it possible to specify the right way in a non-circular manner, i.e. without saying that the right way is precisely the way in which reasons cause actions? In this section I will show that the problem of deviant causal chains isn’t solvable by adhering to the idea of a right way of causing. The possibility of deviant chains will turn out to be a mere symptom of a much more fundamental flaw in causal action theory.

But before I can bring out this flaw we will need to take a closer look at the way in which actions are explained by causal action theory. The occurrence of the event that is the action is explained by the existence of a causal relation between the mental state and the physical happening, and the rationality of the action is explained by the content of the mental state that causes the action. Here is an example: we possess a mental state with a certain *intentional content* (e.g. I’d like some milk and there is milk in the fridge) that *causes* a bodily movement (e.g. me walking to the fridge). The content of the mental state gives the reason for the bodily movement; wanting milk and knowing that it is in the fridge justifies a trip to the fridge. And the causal relation makes sure that the reason provided by the mental state is *the* reason on which is acted.

In a recent paper David Horst [2010] made a few interesting points about the assumptions causal action theory adopts with regards to the concepts of causality and intentional content. Causal action theory claims that there is no difference between the kind of causality exhibited in intentional action, and the kind of causality to be found within inanimate material processes. The word ‘because’ refers to the same type of causal connection in both ‘I made a risotto, because I was hungry’ and ‘The crane fell, because it was blown over’. It is this *uniform* understanding of causality that allows causal action theory to explain the happening of actions: the mind is able to cause material bodily movements, because mental states are themselves material processes. It follows that the causal relation between the mental state and the bodily movement is independent of the

---

<sup>4</sup>See [Davidson 1978].

fact that this state's content also *justifies* the performance of the bodily movement.

Causal action theory also has a *uniform* understanding of intentional content. All mental states are propositional attitudes that are to be analysed into two separate elements: 1) a content that is given by a proposition and 2) a causal role indicated by the relevant attitude verb. A belief, for instance, is caused by what it represents (the milk in the fridge causes my belief that there is milk in the fridge), a desire causes what it represents (wanting the milk in the fridge causes me to go and put it there).<sup>5</sup> This view allows causal action theorists to assume material processes that correspond to certain types of mental state, on the basis of causal relations, i.e. it allows them to claim that reasons indeed somehow are materially realized. From this it follows that it can make no difference to the nature of the content of the mental state on which I act, that this same state also causes the bodily movement that is my action. Thus causation and content are in principle only externally related and logically independent,<sup>6</sup> something that is nicely reflected in Davidson's observation that both relations are "very different" [Davidson 1963].

Let us return to the problem of deviant causal chains. A deviant causal chain occurs, when we have a bodily movement that is both in accord with and caused by a particular mental state, but where the relevant movement is obviously not an intentional action. In our driver case the driver's despisal of the pedestrian was in accord with his killing him, and it also caused the foot to slip on the gas pedal. Yet we don't want to say that the driver intentionally killed the pedestrian. As mentioned above, proponents of causal action theory respond to such cases by claiming that mental states should cause movements in the right way. But what can this right way be? Perhaps the problem is that the driver pressed the gas pedal only by accident, and in the case of intentional actions we don't want our mental states to cause accidents, we want them to non-accidentally cause the bodily movements that are in accord with them.<sup>7</sup> I couldn't agree more, but the problem is that the accordance of content and the causal relation are logically independent. This means that *every single time* a mental state causes a bodily movement that is in accord with it, it does so accidentally. Another suggested solution is that we do not need cau-

---

<sup>5</sup>Searle [2003] calls this the direction of causation and links the concept to his idea of directions of fit discussed in §2.3.

<sup>6</sup>There is no *logical* impossibility in a desire to drink a beer that causes an opening of the window.

<sup>7</sup>One strategy to prevent accidentality is to stress that the mental state should not only trigger the movement, but guide it through the different phases of its occurrence. I don't think that this would help much because it must be possible to cook up a more devilish deviant causal story where the mental state accidentally guides a bodily movement that is accidentally in accord with it.

sation in the right way, but *accordance in the right way*: the content of the mental state should contain an action plan that spells out the way in which an action should proceed. This action plan should make sure that the bodily movements are not merely accidentally in accord with the mental states that cause them. This strategy of course does not work for the same reason for which the causation in the right way-strategy did not work: accordance of content and causal relation are logically independent.

Horst uses the unearthed assumptions to show that the problem of deviant causal chains is inescapable for causal action theorists. I would like to go one step further: if the co-occurrence of justification and causation *can* only be accidental, due to the logical independence of causality and content, then every intentional action is in a sense deviant. Or better: the problem of deviant causal chains is only a sub-problem of a wider problem in causal action theory. This problem is that when one acts *for* a reason one acts in order to achieve an end and on grounds of the belief, but this cannot be accounted for by the existence of accidental causal relations. Anscombe writes:

“[Davidson] speaks of the possibility of ‘wrong’ or ‘freak’ causal connexions. I say that any recognizable causal connexions would be ‘wrong’, and that he can do no more than postulate a ‘right’ causal connexion in the happy security that none such can be found. If a causal connexion were found we could always still ask: ‘But was the act done for the sake of the end and in view of the thing believed?’” [Anscombe 1974, pp. 110–111]

It seems impossible to evade this problem. A Davidsonian cannot simply come up with a version of causal action theory in which content and causality are somehow dependent, because it is precisely this independence that allowed Davidson to explain the physical happening of actions, i.e. them being caused by materially realized mental events. It seems that if a theory of action wants to be successful, it should drop the logical independence between the causal relation and the relation of accordance of intentional content. Causal action theory failed because it tried to explain the material occurrence and the rationality of intentional actions separately, but in intentional actions both aspects are united. Intentional actions wouldn't be intentional if they were not rational, and they wouldn't exist if they didn't happen.

## 4.4 Non-causal action theory

Due to the problem of deviancy and other problems for causal action theory, some philosophers have attempted to come up with contemporary versions of non-causal ac-

tion theory, i.e. a theory that wants to explain actions without pointing to their causation. In this section I will discuss and criticize two types of non-causal theory. I will start by discussing Tanney's modern neo-Wittgensteinian theory in §4.4.1, and follow by considering the alternatives of McCann and Ginet in §4.4.2.

#### 4.4.1 Reasons and circumstances

Julia Tanney disagrees with Davidson about the nature of rationalization. According to her rationalization isn't causal explanation, it rather is context placing explanation. Tanney notes that an action is often easily rationalized once we get to know more about the circumstances in which the action took place. 'Why did the woman run out of the building?' 'The building was on fire.' According to Tanney this fully explains the action. We need no further inquiry into something inner or hidden in the agent, something that might be called the reason that caused the agent's action. The question whether the agent was motivated or moved to act by noting that the building was on fire is a question that normally does not arise. It may be objected that learning more about the circumstances does not always relieve our puzzlement about why a person acted as she did. Sometimes we will have to mention the agent's conception of the circumstances (e.g. when the agent fled because she thought the building was on fire while in fact it was not.) And sometimes we may even have to ask her why she did it. Tanney stresses however that when we ask an agent what she does, this answer is also constrained by the circumstances. We wouldn't accept the woman's answer that she ran from the building because it was on fire, if it turns out that she was a firefighter on the job. Furthermore Tanney argues that in appealing to how an agent conceives of her situation we are *not* homing in on something inner or hidden. Talk about what the agent knows, desires, believes or realizes is simply a reconstruction of how it makes sense for her to have acted, it is not supposed to refer to mental or brain states. Tanney concludes that: "the explanatory function of reason concepts may be fully discharged by the placement of the action to be explained within the appropriate circumstances or wider context" [Tanney 2009]. There is no need to look further for the psychological states that moved a person to act. There is simply nothing missing from ordinary reason explanation. Therefore, there is no philosophically interesting question about mental causation that remains to be asked.

Tanney's account of the explanation of actions is problematic in several respects. First of all, there is nothing in Tanney's theory that explains where the knowledge of our own actions and reasons for acting comes from. We do not need to look at the

wider context of our own actions in order to be able to say what we are doing and why we are doing it; we can say what we are doing and why without looking anywhere. Another issue mirrors the problem of deviant causal chains and might perhaps be called the problem of ‘deviant contextual circumstances’. Recall the driver who despised the pedestrian. We are perfectly able to rationalize his bodily movements by a look at his circumstances. He drove over the pedestrian because he despised him (attributing hate to the driver here has, of course, got nothing to do with what goes on in the driver’s head, it is simply a reconstruction of how it makes sense for the driver to act). Even if the driver were to state, in a court of law perhaps, that his foot accidentally slipped onto the gas pedal, I sincerely doubt his account would be publicly accepted (in Tanney’s sense) as an explanation. This shows that we can use contextual circumstances to rationalize events that are not intentional actions (like the jerk of the driver’s foot) and thus stand in no need of rationalization.<sup>8</sup>

But even in the case of intentional actions context-placing explanations can lead us astray. This happens, for instance, when the circumstances in which an agent acts allow us to rationalize her actions in several different ways. This can lead to several different rationalizations all of which we might be equally keen to accept. It gets even worse when rationalization *A* might be more acceptable than rationalization *B*, but where *B* nevertheless is in fact the right rationalization. The woman might have run out of the burning building because she wanted to go and grab a hamburger. There can simply be a gap between the reasons that justify someone’s actions and the actual reasons for which she acted. It seems that context-placing explanation can do nothing more than rendering an agent’s action acceptable in a certain situation, thus it provides justificatory reasons for the action without telling us anything about the actual reason for which the agent acted. We must conclude that Tanney’s theory cannot cope with Davidson’s original criticisms of neo-Wittgensteinians: mere justification is not enough to fully explain actions.

#### 4.4.2 Decisions and concurrent intentions

Two other alternatives to causal action theory are sketched by Ginet [1990] and McCann [1998]. Both take up the challenge of explaining how an action can be done *for* (and not merely justified by) a reason. Unlike Tanney they do not hesitate about invoking the

---

<sup>8</sup>This is the parallel with the deviant causal chain problem; intentions can cause events that are not intentional actions.



private contents of an agent's mind in the explanation of her actions.<sup>9</sup> I shortly discuss both alternatives and show why they are unsatisfactory.

According to McCann agents act for a reason when they have *decided* to act for that reason. He thinks that decision making is a type of information processing: it is to transform what is presented in our reasons, so that the scenario of action they embody comes to be intended. Therefore we can never just decide, we always decide on certain information.<sup>10</sup> If we decide to act, this information forms the reason for which we act. Thus we can explain why an agent acted for a specific reason, by pointing to the decision that formed her intention to act. Now it seems that we can very easily cook up a scenario of deviance for McCann's story. Again our driver is waiting for the pedestrian and recognizes that it is the despicable X. In this case the driver is a bit more malicious so upon seeing X he decides to run him over. But before he actually executes his intention to kill X, something else causes his foot to accidentally press the gas pedal. The bodily movement happened and the decision was in place, so on McCann's account the agent acted for a reason. But we know that the bodily movement only occurred accidentally, so there was no intentional action at all.

Ginet argues that an agent acts for a reason if she acts in such a way as her desire dictates and also has a concurrent intention during her action that by it she fulfills her desire.<sup>11</sup> Here is an example: suppose Sarah wants her glasses that she believes are in the bedroom. If Sarah does enter the bedroom we can say that she did it to get her glasses only if she concurrently intended of her bedroom-entering that by it she would get her glasses. Thus it is the content of the concurrent intention that specifies the reason *for* which the action was done. To me this account seems hopelessly complex. I am for instance never aware of my concurrent intention that I intend my finger movements to produce the thesis I intend to write. I just intend to write a thesis and therefore I type it. But the real problem with Ginet's account is that it also allows for scenarios of deviance. Clarke [2006] suggests such a scenario. Suppose that Sarah not only wants her glasses, she also has a desire for her friend Ralph's company. Ralph just came home after a hard day of work and is taking a nap in the bedroom. Sarah knows this and therefore she also wants to let him sleep. Now suppose Sarah enters room with the concurrent intention that by it she would get her glasses. However as it turns out, her desire to get her glasses

---

<sup>9</sup>Therefore they can explain an agent's authority over what she is doing and why she is doing it in the same way as causal action theory: only the agent has introspective access to her own mind, and therefore the agent has authority over what she does.

<sup>10</sup>See [McCann 1998, chapter 8].

<sup>11</sup>See [Ginet 1990].

plays no role at all in causing her entry, while her desire to wake Ralph, of which she is fully aware when she acts, does play such a role. Is the content of her concurrent intention authoritative about the reason for which Sarah acts? Or is she fooling herself? In a recent paper Ginet answers Clarke's challenge:

I have to admit that, since it is stipulated that Sarah's desire to wake Ralph plays a role in causing her act of entering, it will certainly seem right to say that she entered *because* she had a desire to wake Ralph and even that her desire to wake Ralph was one of her reasons for entering (assuming that the causation was not deviant). [Ginet 2008, p. 232]

According to Ginet this shows that the having of a concurrent intention isn't (unlike he previously defended) necessary for establishing the truth of reason-explanations. He still holds, however, that there are cases where it is *sufficient*. Therefore he argues, contra causal action theory, that we do not always have to cash out rationalization in causal terms [Ginet 2008, p. 236]. I think Ginet is right about this, just like Tanney was right that rationalizations often point to circumstances instead of causes. But Davidson was also right when he stressed that there is more to actions than their rationality: they also are material happenings. Because the non-causal accounts focus entirely on the rational aspect of intentional actions, they can never account for their *happening* in the world.<sup>12</sup>

## 4.5 Practical knowledge and causation

On the bases of our discussion of both causal and non-causal action theories we must conclude that we are faced with quite a challenge. A definitive theory of action should be able to explain both why actions are rational and why they happen. And not only should it explain both of these aspects, it should explain them conjointly in order to avoid that the rationality and the happening of intentional actions are only accidentally related. Thus we need an account that drops the assumptions of the causal action theorist and unifies the relation of accord with the causal relation; an account that truly explains how an intentional action and the reason for which it is performed relate to each other. An account based on Anscombean practical knowledge might be able to do just that.

The first major difference between Anscombe's account and the earlier discussed accounts of intentional action is that Anscombe doesn't take actions as 'given' events.<sup>13</sup> The causal (§4.2) and non causal (§4.4) accounts discussed above saw actions as bodily

<sup>12</sup>That is unless the decisions or concurrent intentions that account for the rationality also produce the action. Precisely at this point, however, the action theory ceases to be a *non-causal* account.

<sup>13</sup>I'm borrowing Vogler's terminology here [Vogler 2001].

movements. Of each given bodily movement in can be asked whether it is or isn't an intentional action. This question is to be answered by looking at the way in which the event was caused (Davidson), or on whether it could be rationalized in the circumstances (Tanney), or on whether the appropriate mental processes were preceding (McCann) or concurrent (Ginet) with the movement. For Anscombe intentional actions are different from other events not because of their circumstances or causation, they are *intrinsically* different from mere happenings in that they possess a rational structure. An action can often be described in different ways, e.g. as 'moving your hand in a certain manner' as 'cutting an onion' and as 'making a risotto'. These different descriptions stand in a means-end order: moving your hand is a means to cutting an onion, and cutting an onion is a means to making a risotto. Now while the descriptions are different they all are descriptions of the same bodily movement.<sup>14</sup> According to Anscombe we need to have a grip on the internal means-end structure in order to single out something as an intentional action. This is why she defines intentional actions as those movements to which a certain sense of the question 'why?' can be given application. The 'why?' question she has in mind is precisely the one that brings out the means-end structure of the action. Thus movements to which this question can be given applications are those that bear a rational structure.<sup>15</sup>

If we view intentional actions as not just events like any other except for the fact that they are caused in a certain way or can be rationalized in certain circumstances, we can easily overcome the problem of deviant causal chains.<sup>16</sup> When we ask our driver why his foot slipped on the gas pedal, he will answer: '*because* I was scared by my thoughts about killing that bugger.' Of course he might also reply 'because of my desire to kill him', but here his 'because' would be causal as in 'the crane toppled over because of a gust of wind', it would not point to a means-end structure. Therefore, on Anscombe's account the slipping of the foot comes out as what it is: an unintentional movement, more like a peristaltic movement of the gut or like the jerk one's body sometimes gives when one is falling asleep. This strategy of denying upfront that a certain bodily movement is intentional is open to neither causalists nor non-causalists, because they *define* intentional actions as events or movements with certain causal histories or contextual circumstances or preceding/concurrent mental processes.

---

<sup>14</sup>Of course moving your hand in a certain way only is a part of onion-cutting which is only a part of risotto-making, but it is nevertheless correct to say of the agent who currently moving her hand in a certain way, that she *is* cutting onions or that she *is* making a risotto.

<sup>15</sup>See chapter 1, especially sections 1.1 and 1.5.

<sup>16</sup>As well as the scenarios of deviance for the theories discussed in §4.4.

As I argued in section 4.3 the deviant causal chains were only a symptom of a deeper problem for causal action theory: the accidentality of the accordance in content between a mental state and the movement it causes. My claim is that an Anscombean theory of action cannot only overcome the symptom, but it can overcome the disease. Its cure is practical knowledge which is the cause of what it understands.<sup>17</sup> In causal action theory we saw that a reason had a content that represented the action, which was independent from its causal relation with the bodily movement. A reason in causal action theory might thus be called the ‘cause of what it represents’. As Haase suggests [2010] every event or object can be represented. In the deviant case the driver’s thought represented (and caused) the running down of the pedestrian. However, only events or objects with a rational structure can be understood. We can understand essays, but merely represent stones; we can understand intentional actions, but merely represent startle responses. Because practical knowledge is the cause of what it *understands*, we can have practical knowledge of our own (rationally structured) intentional actions, but not of our unintentional bodily movements. Having practical knowledge thus isn’t merely having an attitude towards a proposition (believing it, desiring it) it is understanding a structured process.<sup>18</sup> Thus an action theory based on practical knowledge drops the causalist assumption of the uniformity of intentional content.

But how can practical knowledge cause actions? And especially, how can it non-accidentally cause the actions it understands? In order to answer these questions we have to recall (from §2.4) that actions aren’t pointlike events. They often consist of several different phases and it takes time for actions to reach completion. Making a risotto, for instance, consists of: cutting onions, taking butter out of the fridge, grating cheese, and adding wine. It is *because* I know that I am making a risotto that I perform all these subactions. My knowledge of what I am doing brings these phases about and makes them into phases of one risotto making. Without practical knowledge the cutting of the onions and the adding of the wine would not be part of a risotto-making process, i.e. the description ‘making a risotto’ would not apply to them. In other words, the structure of the intentional action of risotto-making is *caused* by my practical knowledge. It is, moreover, the understanding of the different phases involved in the action that causes this same structure in the actual action. Therefore practical knowledge is not *accidentally* both the understanding and causing of the intentional action. It is rather the

---

<sup>17</sup>See §1.7.

<sup>18</sup>As we saw in §1.6 having practical knowledge involves knowing how to do something, i.e. knowing which phases of the action must be performed in which order.

understanding of the action *that* causes it. It is in this way that practical knowledge is able to unify the causal relation between the thought and the action with the relation of accordance.

But how does practical knowledge cause what happens? Is having practical knowledge a mental event that causes the bodily movement? If this were to be the case we would be back at causal action theory. It has therefore been suggested that we should think about practical knowledge as the formal cause of the action.<sup>19</sup> The knowledge causes the process in the world to have a certain form, i.e. a means-end structure. But that is not to say that without practical knowledge the bodily movements would have happened anyway, and that it *merely* characterizes these events as intentional. Practical knowledge also is productive: it brings about the intentional action.<sup>20</sup> But how can it do so if practical knowledge isn't a material event?

I think that we might find an answer if we look at the notion of causality itself. In a recent paper Jonathan Lowe suggests that there is no room for human rationality, freedom and intentional action because of the almost universal presumption that all causation is event-causation:

Once it is assumed that all causation is a matter of one or more events causing another, it is difficult to resist the thought that human decisions and actions are, after all, just events and as such woven into the universal web of event causation, the only alternative being that they are random outcomes of chance [Lowe 2009, p. 340].

But what is the alternative? One suggestion is that human agents are special in that they can intentionally cause events to happen. In the natural world events cause other events via event causation, but humans possess the special power of agent-causation: the power to produce natural events out of thin air. But ascribing human agents their own peculiar causal powers seems both *ad hoc* and mysterious. How can an agent just cause an event? What makes them so special? Lowe answers these problems by stressing that we aren't special in our capacity to produce events, many other *substances* in the world can do so as well:

A magnetized piece of iron *acts* upon some nearby iron filings to make them move towards it: it 'attracts' them. Some sulphuric acid *acts* upon some bits of copper to make them dissolve and transform their constituent atoms into copper ions: it 'oxidizes' them. [emphasis NvM] [Lowe 2009, p. 341]

<sup>19</sup>see [Hursthouse 2000], [Teichmann 2008], and [Moran 2004].

<sup>20</sup>See *Intention* [§48].

Thus Lowe argues that it is wrong to think of causality as a relation between events (regardless of whether these events have a human or non-human origin). “Causality fundamentally is a matter of substances exercising their causal powers to act upon other substances possessing suitable causal liabilities” [Lowe 2009, p. 341]. This is not to say that all talk of events causing other events is misguided. It is, however, misguided to believe that it ultimately are events that bring about changes. According to Lowe events as such are utterly powerless. They are mere changes in things and not the source of those changes.

A full defence of substance-causation would be a nice topic for another thesis, but it is far beyond the scope of this one. There are, however, some considerations in this thesis that might turn out to favour substance causation over event causation. In chapter 3, we encountered Anscombe’s criticism of Davidson’s (and Hume’s) idea that all causal connections between particular events must instantiate general laws. She was amazed by the fact that Davidson simply assumes that every causal process requires a universal law, without offering any reason for it. I think that Davidson is forced into this assumption by his own ontology of events. According to Davidson events are fundamental, irreducible entities to which properties can be ascribed. By assigning them such a fundamental role there remains nothing that accounts for them. Davidson’s ontology contains no substances or things in which events are grounded. But as Lowe argued events themselves are powerless in that they are the changes themselves and not the sources of change. Therefore Davidson has to analyze causality in terms of generality to make any sense of it. There is simply no other option open to him. A defender of substance causality, however, does not have to make such a step because he can analyse the causes of any event in terms of the substance(s) that produce them. Another consideration in favour of substance causation is that it seems better suited to explain how we acquire causal knowledge. In §3.3 I suggested that causal knowledge is gathered by experimenting, intervening and playing with the objects in one’s environment. Now events aren’t the best test subjects in that they only occur once, never to return to our laboratories (or playgrounds) again. Objects however, persist: we can do something with them and if what we did wasn’t too destructive we can repeat or try something else. Of course the proponent of event causation might retort that while events do not occur more than once, different events can be of the same *type*. While this seems a perfectly fine reply, it does give them one problem that defenders of substance causation do not have to face: how to specify when two events are of the same type.

Let us return to intentional actions. The challenge was to explain how they can

be material happenings in the world. When event causation is presupposed, we must explain the material happening in terms of another happening that caused it. This was the presupposition of causal action theory that construed reasons as material events in order to allow reasons to cause actions. As we saw in §4.3 this explanation is flawed. Substance causation allows us a way out of our troubles, because in order to explain the material happening that constitutes an action we can point to the substance that is its source. In the case of intentional actions that means human beings. On this account explaining how humans can act is similar to explaining how magnetized iron can attract iron filings. We have to look at the properties and capacities of iron in order to explain its attracting powers, similarly we have to look at the properties and capacities of human beings in order to explain intentional actions. One of the capacities a human being has is the capacity for practical knowledge. As Anscombe explains, exercising this capacity is “nothing but the doing or supervising of the operations of which a man has practical knowledge” [§48].

The ultimate explanation of intentional actions, thus has to go hand in hand with the explanation of the nature of human beings. It might be objected that this isn't satisfactory because we precisely wanted to know how human beings can bring about intentional actions. Merely saying that human beings have the capacity to act doesn't sound like much of an explanation. However, the answer this thesis provides isn't vacuous in that I have offered an analysis of what this capacity consists in. As it turns out, the capacity to act intentionally precisely *is* the capacity for practical knowledge that has been the topic of our enquiry. That acting is a capacity of human beings is something a theory of action cannot get rid of. It was precisely the attempt to fully analyse intentional actions in non anthropocentric terms that has led much of the twentieth century philosophy of action astray.

# Conclusion

**T**HIS thesis had two objectives: 1) a defence of the Anscombean concept of practical knowledge and 2) a plea for the importance of practical knowledge to contemporary philosophical debates concerning both knowledge and action. Below I will summarize my findings per chapter and suggest some possibilities for further research.

The first chapter attempted to achieve a clear understanding of the nature of practical knowledge on the basis of what Anscombe related to us in *Intention*. Where some have suggested that there were multiple concepts of practical knowledge in play in Anscombe's monograph, I have attempted to show that we can extract one uniform conception of practical knowledge from it. Practical knowledge is the knowledge an agent has of her own intentional action. While it is aptly called knowledge in the sense that someone who possesses it can truly say what she is doing, it turned out to be quite different from what we would ordinarily call knowledge, i.e. speculative knowledge. First of all practical knowledge cannot be acquired from the senses, it is in the intention of the acting agent. It is also different in that it isn't supposed to fit how the world is, rather the world has to fit it. I tried to further explain practical knowledge by following Anscombe's suggestion that it can only be understood if we first understand practical reasoning. Practical reasoning deduces means from ends and is productive of the action. Practical knowledge understands the means-end structure of intentional actions and through this understanding it can cause the action to come about. Therefore having the capacity to act intentionally is having the capacity for practical knowledge.

Many have found the concept of practical knowledge ultimately problematic. Their main issue stems from the fact that an agent can intend to perform an action and think that she is doing so, while unbeknownst to her something goes wrong. Therefore her belief that she is doing X does not amount to knowledge. However she still made intentional bodily movements, ergo intentional movement doesn't require practical knowledge. Chapter 2



has attempted to rebut this criticism. The first step towards a rebuttal was to show what it means that practical knowledge has a mind to world direction of fit: it means that if something goes wrong the mistake was not in the knowledge, it was in the performance. The second step was to bring out that actions take time to complete. By using the imperfective aspect we can describe them while they are happening. We can use the perfective to refer to completed actions. I could, for instance, be writing my thesis right now, even if disaster strikes and I never manage to have written it. Similarly we can say of the blind man who has run out of ink that he is writing his name even if at that moment no letters appear on the page. Because of this we can say that he knows he is writing, even though he might be unaware of the fact that he will not succeed by merely continuing to move the pen over the paper. If we would however point this out to him, he would probably refill his pen and continue with writing his name. Thus objective 1 was reached.

Chapter 3 attempted to show how actions as they are understood by Anscombe can play a role in the study of knowledge: epistemology. I started with an inventarisation of the ways in which non-practical knowledge can enter into the capacity to perform intentional actions. It can enable us to act, and also guide our actions as they proceed. Unlike some have thought, the fact that speculative knowledge plays an important role in our performance of intentional actions is no threat to the doctrine of practical knowledge. In fact we can partly explain the connections between speculative knowledge and actions by looking at practical reasoning of which productive practical knowledge is the result. The chapter proceeded by suggesting that, due to the close connection between speculative knowledge and action, actions may play an important role in our acquisition of knowledge about the world around us. This was illustrated by showing how our actions can allow us to understand causal processes. The final part of this chapter showed that the contemporary epistemological discussion about the relation between actions and practical knowledge is misguided because it doesn't aptly distinguish between practical and theoretical knowledge.

In the final chapter (chapter 4) we saw that action theory since Wittgenstein has struggled to explain the duality of intentional action. On the one hand intentional actions are done for a reason and therefore they are rational entities. But on the other hand actions also are material processes that happen in the world. Neo-Wittgensteinians were only focussed on explaining actions via rationalization, therefore they were open to Davidson's criticism that they merely explained how reasons could justify actions, whereas a

complete account of intentional action should also explain how they come about. Causal action theory is an attempt to explain both aspects of actions. According to that approach reasons are mental states/events, the content of these entities rationalizes intentional actions and the physical realization of them can materially cause the actions to happen. Unfortunately causal action theory is flawed because it views intentional content and causality as logically independent. Therefore it cannot make sure that an action that is caused by a reason also is done on grounds of the reason. The contemporary non-causal action theories considered also turned out to be unsatisfactory. Not only did they have problems with deviancy of their own, they also generally failed to account for the material happening of the events that are intentional actions. The duality of actions can, however, be explained by an action theory that is based on practical knowledge. This knowledge understands the rational structure of the actions and is also able to bring actions with such a rational structure about. But how does this ‘bringing about’ work? To solve this I argued that we should abandon the assumption that all causation is fundamentally event causation. When substance causation is given primacy, the question of how actions are brought about becomes a question about the properties and capacities of the substance that brings them about, i.e., rational human beings. The human capacity to act intentionally precisely is the capacity for practical knowledge that this thesis has analyzed.

I believe that the final two chapters have indeed established the importance of the concept of practical knowledge in both epistemology and action theory (objective 2). But of course there still is a lot of work that needs to be done. The thought expressed in chapter 3 that actions play a large part in our understanding of the world should be developed and undergirded. Psychologists, for instance, should study how knowledge acquisition through manipulating one’s environment works. Philosophers of science can look at experimental design to see how scientists can test their hypotheses through intervention. Perhaps a better understanding of human interventions can even be applied to create stronger experiments. The notion of substance causation discussed in chapter 4 also requires a lot of development. One concrete project would be to formulate an account of interventions in terms of substance causation.<sup>21</sup> As I suggested the substance causation might be more suited to an interventionist approach than event causation, in that we normally manipulate objects and it is only through them that we interfere with

---

<sup>21</sup>cf. [Woodward 2003], which focuses entirely on event causation.

events.

This thesis has provided an analysis of the human capacity for practical knowledge. We have discovered that this capacity is inextricably linked to our ability to act intentionally. Despite the analysis I offered, one might continue to marvel at the fact that we can perform actions. It is, in a sense, amazing that we can just make stuff happen. For now, this will be our end station.

Perhaps we should call the fact that men can perform actions a ‘mystery’ – in the sense that it is something basic which defies explanation. At least, I cannot imagine what an ‘explanation’ would conceivably look like. We can wonder at this mystery – as we may wonder at the fact that man can know things or that there is an external world. It is interesting to note that wonder of this type can be both the starting point and the end station of philosophical inquiry. [von Wright 1989]

# References

- ADAMS, D., 1996 [1982]. *The ultimate hitchhiker's guide*. Random House Inc. New York.
- ADAMS, F. and MELE, A., 1989. 'The role of intention in intentional action'. *Canadian Journal of Philosophy*, 19(4): pp. 511–531.
- ANSCOMBE, G.E.M., 1963. *Intention*. Harvard University Press.
- ANSCOMBE, G.E.M., 1971. 'Causality and determination'. Page references to the excerpt reprinted in Sosa and Tooley [1993].
- ANSCOMBE, G.E.M., 1974. 'Von Wright on practical inference'. pp. 377–404. In: [Schilpp and Hahn 1989].
- ANSCOMBE, G.E.M., 1979. 'Under a description'. *Nous*, 13(2): pp. 219–233.
- AUSTIN, J., 1952. 'How to talk — Some simple ways'. In 'Proceedings of the Aristotelian Society', volume 53, pp. 227–246.
- BRATMAN, M., 1987. *Intention, plans, and practical reason*. Harvard University Press Cambridge, MA.
- CLARKE, R., 2006. *Libertarian accounts of free will*. Oxford University Press.
- DAVIDSON, D., 1967. 'Causal relations'. *The Journal of Philosophy*, 64(21): pp. 691–703.
- DAVIDSON, D., 1963. 'Actions, reasons and causes'. *Journal of Philosophy*, 60: pp. 685–700.
- DAVIDSON, D., 1969. 'Mental events'. Chapter 5 in: [Lepore and Ludwig 2006].
- DAVIDSON, D., 1978. 'Intending'. Chapter 6 in: [Lepore and Ludwig 2006].
- DIAMOND, C. and TEICHMAN, J., 1979. *Intention and intentionality: essays in honour of GEM Anscombe*. Cornell University Press.
- DRETSKE, F., 1989. 'Reasons and causes'. *Philosophical Perspectives*, 3: pp. 1–15.
- DULLSTEIN, M., 2007. 'Do we perceive causings?' Unpublished. Paper presented at the 'Deutsch-italienisches Treffen für analytische Philosophie' in Parma.
- FALVEY, K., 2000. 'Knowledge in intention'. *Philosophical Studies*, 99(1): pp. 21–44.

- FANTL, J. and McGRATH, M., 2002. 'Evidence, pragmatics, and justification'. *Philosophical Review*, 111: pp. 67–94.
- FODOR, J., 1992. 'Making mind matter more'. *The Philosophy of Mind: classical problems/contemporary issues*, p. 151.
- GEACH, M. and GORMALLY, L., eds., 2006. *Human life, action and ethics; essays by GEM Anscombe*. Exeter: Imprint Academic.
- GIBBONS, J., 2010. *Oxford studies in epistemology*, chapter Seeing what you're doing, pp. 63–85. Clarendon Press Oxford.
- GINET, C., 1990. *On action*. Cambridge University Press.
- GINET, C., 2008. 'In defense of a non-causal account of reasons explanations'. *The Journal of Ethics*, 12(3): pp. 229–237.
- GOLDMAN, A., 1979. 'Action, causation, and unity'. *Nous*, 13(2): pp. 261–270.
- GOPNIK, A. and SCHULZ, L., eds., 2007. *Causal learning: Psychology, philosophy, and computation*. Oxford University Press, USA.
- GRUNBAUM, T., forthcoming. 'Anscombe and practical knowledge of what is happening'. *Grazer Philosophische Studien*, 78.
- HAASE, M., 2010. 'Realizing a concept'. Unpublished. Paper presented at workshop 'Practical Knowledge II' in Basel.
- HACKING, I., 1983. *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge University Press.
- HAWTHORNE, J., 2004. *Knowledge and lotteries*. Oxford University Press.
- HITCHCOCK, C., 1995. 'The mishap at Reichenbach fall: Singular vs. general causation'. *Philosophical studies*, 78(3): pp. 257–291.
- HORGAN, T., 1989. 'Mental quausation'. *Philosophical Perspectives*, 3: pp. 47–76.
- HORST, D., 2010. 'Practical knowledge as an account of intentional action'. Unpublished. Paper presented at workshop 'Practical Knowledge II' in Basel.
- HUME, D., 2000 [1748]. *An enquiry concerning human understanding*. Oxford University Press, USA.
- HURSTHOUSE, R., 2000. 'Intention'. *Royal Institute of Philosophy Supplement*, 46: pp. 83–105.
- KIM, J., 2007. 'Causation and mental causation'. *Contemporary Debates in the Philosophy of Mind*, pp. 227–242.
- LEPORE, E. and BRIAN, M., eds., 1985. *Actions and events; perspectives on the philosophy of Donald Davidson*. Basil Blackwell.

- LEPORE, E. and LUDWIG, K., eds., 2006. *The essential Davidson*. Oxford: Clarendon Press.
- LOWE, E., 2009. 'Free agency, causation and action explanation'. Chapter 11 in [Sandis 2009].
- MCCANN, H., 1998. *The works of agency: On human action, will, and freedom*. Cornell University Press.
- MCDOWELL, J., 1996. *Mind and world*. Harvard University Press.
- MELDEN, A., 1967. *Free action*. Routledge & Kegan Paul London.
- MELE, A., 2009. *Effective intentions: The power of conscious will*. Oxford University Press, USA.
- MILLIKAN, R., 1993. *White queen psychology and other essays for Alice*. MIT Press.
- MÜLLER, A.W., 1979. 'How theoretical is practical reason'. In [Diamond and Teichman 1979].
- MÜLLER, A., 2010. 'How many kinds of practical knowledge?' Unpublished. Paper presented at workshop 'Practical Knowledge II' in Basel.
- MÜLLER, T., 2010. 'Branching in the landscape of possibilities.' Unpublished. Paper presented at workshop 'Branching and indeterminacy' in Bristol.
- MORAN, R. and STONE, M., 2008. 'Anscombe on expression of intention'. *New Essays on the Explanation of Action, Basingstoke*.
- MORAN, R., 2004. 'Anscombe on "practical knowledge"'. *Philosophy*, 55 (Supp): pp. 43–68.
- NIESCHALK, M., ORTMANN, C., WEST, A., SCHMÄL, F., STOLL, W., and FECHNER, G., 1999. 'Effects of alcohol on body-sway patterns in human subjects'. *International journal of legal medicine*, 112(4): pp. 253–260.
- PAUL, S.K., forthcoming. 'How we know what we're doing'. Philosophers' Imprint.
- PEACOCKE, C., 2003. 'Action: Awareness, ownership and knowledge'. In 'Agency and self-awareness', Clarendon Press Oxford.
- PICKARD, H., 2004. 'Knowledge of action without observation'. *Proceedings of the Aristotelian Society*, 104: pp. 205–230.
- QUINE, W., 1969. 'Epistemology naturalized'. In 'Ontological Relativity and Other Essays', pp. 68–90. New York: Columbia University Press.
- RÖDL, S., 2010. 'Two forms of practical knowledge and their unity'. Unpublished. Paper presented at workshop 'Practical Knowledge II' in Basel.
- RYLE, G., 2002 [1949]. *The concept of mind*. University of Chicago Press.
- SANDIS, C., 2009. *New essays on the explanation of action*. Palgrave Macmillan.

- SCHILPP, P.A. and HAHN, L.E., eds., 1989. *The philosophy of G. H. von Wright*. La Salle, Ill.: Open Court.
- SEARLE, J., 1983. *Intentionality, an essay in the philosophy of mind*. Cambridge University Press.
- SEARLE, J., 1985. *Expression and meaning: Studies in the theory of speech acts*. Cambridge University Press.
- SEARLE, J., 2003. *Rationality in action*. MIT Press.
- SELLARS, W., 1997 [1956]. *Empiricism and the philosophy of mind*. Harvard University Press.
- SETIYA, K., 2007. *Reasons without rationalism*. Princeton University Press.
- SETIYA, K., 2008. 'Practical knowledge'. *Ethics*, 118(3): pp. 388–409.
- SMITH, M., 1987. 'The Humean theory of motivation'. *Mind*, 96(381).
- SOSA, E. and TOOLEY, M., eds., 1993. *Causation*. Oxford University Press. With an introduction by Sosa and Tooley, pp. 1-32.
- STANLEY, J., 2005. *Knowledge and practical interests*. Oxford University Press.
- TANNEY, J., 2009. 'Reasons as non-causal, context-placing explanations'. In [Sandis 2009].
- TEICHMANN, R., 2008. *The philosophy of Elizabeth Anscombe*. Oxford University Press.
- THOMAS, ST., 1274. 'Summa theologiae'.
- THOMPSON, M., 2008. *Life and action*. Harvard University Press.
- THOMPSON, M., forthcoming. *Essays on Anscombe's Intention*, chapter Anscombe's intention and practical knowledge. Harvard University Press.
- VAN MILTENBURG, N., 2008. 'The (W)right view on causation and action; The agency account of causation and its limits'. Unpublished BA-Thesis, Utrecht University.
- VELLEMAN, D.J., 2004. 'Précis of the possibility of practical reason'. *Philosophical Studies*, 121: pp. 225–238.
- VELLEMAN, D.J., 1989. *Practical reflection*. Princeton University Press.
- VELLEMAN, J., 1992. 'The guise of the good'. *Nous*, 26(1): pp. 3–26.
- VOGLER, C., 2001. *Varieties of practical reasoning*, chapter Anscombe on Practical Inference. MIT Press.
- VON WRIGHT, G.H., 1971. *Explanation and understanding*. Routledge & Kegan Paul, London.
- VON WRIGHT, G.H., 1974. *Causality and determinism*. Columbia University Press, New York.
- VON WRIGHT, G.H., 1989. 'Reply to my critics'. pp. 731–888. In: [Schilpp and Hahn 1989].

WITTGENSTEIN, L., 2001 [1953]. *Philosophical investigations*. Wiley-Blackwell. Translated by G.E.M. Anscombe.

WOODWARD, J., 2003. *Making things happen*. Oxford University Press.

WRIGHT, C., 2001. *Rails to infinity: essays on themes from Wittgenstein's philosophical investigations*. Harvard University Press.



*(This page intentionally left blank)*<sup>22</sup>

---

<sup>22</sup>Or in fact I was intentionally leaving it blank, but unfortunately I did not successfully do so because I wrote the sentence to which this footnote is attached. . . and this footnote.