

PRACTICAL KNOWLEDGE WITHIN OUR MEANS

by

Benjamin Schulz

BA/BSc (Hons) Philosophy, Linguistics & Cognitive Science, UNSW, 2006

Submitted to the Graduate Faculty of

The Kenneth P. Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2018

UNIVERSITY OF PITTSBURGH

The Kenneth P. Dietrich School of Arts and Sciences

This dissertation was presented

by

Benjamin Schulz

It was defended on

September 11, 2018

and approved by

Edouard Machery, Distinguished Professor, History & Philosophy of Science

Kieran Setiya, Professor, Philosophy (MIT)

Michael Thompson, Professor, Philosophy

Karl Schafer, Professor, Philosophy (UC Irvine)

Dissertation Advisor: John McDowell, Distinguished University Professor of Philosophy

Copyright © by Benjamin Schulz

2018

# PRACTICAL KNOWLEDGE WITHIN OUR MEANS

Benjamin Schulz, PhD

University of Pittsburgh, 2018

Is there a capacity for knowledge? I argue that it is. More precisely, I argue that we possess a distinctive kind of ungrounded self-knowledge of our intentional actions. After defending this epistemological claim against putative counterexamples, I offer a metaphysical argument in support of its truth. On my view, the possession of such knowledge follows from the *nature* of an intentional action, from what an intentional action *is*. To show why this is so, I develop and defend an account of intentional action upon which such actions essentially involve our capacity for practical reasoning. By my lights, when we act intentionally, we are always in a position to offer a certain type of reasons explanation of what we are up to. It turns out, though, to be a condition on the possibility of such explanations that we also have knowledge of our reasons, and this in turn presupposes that we know what we are intentionally doing. In this way, I suggest, we can see why there is a necessary link between intentional action and knowledge. Aside from explaining how such a link holds between agency and cognition, I use this picture of intentional action to illuminate other issues in contemporary practical philosophy. First, I show how this picture helps to herald a distinctive rapprochement between so-called ‘causalists’ and ‘anti-causalists’ about reasons explanations. Second, I show how it lets us better grasp what G.E.M. Anscombe had in mind when she gnominically described our knowledge of our intentional actions as ‘the cause of what it understands’.

## TABLE OF CONTENTS

PREFACE.....	VII
I. ANSCOMBE'S PRINCIPLE.....	1
1. Introduction .....	1
2. The (Putative) Counterexamples.....	5
3. The Presuppositions of Practical Knowledge.....	12
4. The Role of Progressive Aspect.....	18
5. Two Aspects of Non-Accidentality .....	24
II. ACTING IN THE LIGHT OF MY END .....	31
1. Introduction .....	31
2. The Belief Condition.....	33
3. The Availability of Naïve Explanation.....	40
4. Acting in the Light of a Fact.....	42
5. Why We Always Know What We're Doing.....	48
III. REASONS EXPLANATIONS AS CAUSAL EXPLANATIONS .....	55
1. Introduction .....	55
2. Causalism, Anti-Causalism, and Causal Explanation .....	58
3. Knowledge of Causality.....	64
4. An Analogy with <i>Intention</i> .....	73
5. Causal Pluralism.....	75

6.	Causal Explanation by Intention.....	80
7.	Other Forms of Causal Explanation .....	86
8.	The Explanatory Limits of Causalism.....	88
IV. ANSCOMBE ON THE NATURE OF PRACTICAL KNOWLEDGE.....		92
1.	Introduction .....	92
2.	The Form of Descriptions of Intentional Actions.....	95
3.	Setiya's Reading of Anscombe's Argument.....	96
4.	Setiya on Condition ( <i>A</i> ).....	99
5.	Human-Action Descriptions .....	103
6.	Setiya on Condition ( <i>B</i> ).....	109
7.	Sliding on the Ice Like A Boss .....	112
8.	The Formal Dependence of Action and Knowledge .....	117
9.	The Neglected Category of 'Merely' Voluntary Action.....	123
APPENDIX.....		136
BIBLIOGRAPHY .....		145

## PREFACE

To the memory of Marianne Backhouse

Philippa Foot said that working on a philosophical problem made her feel like a geologist tapping away at a huge cliff face with a tiny hammer. That's an image of the activity of philosophy which I've always liked. It captures both the patience that the work can require and the humility it can demand, while also hinting at the kind of intellectual reward which it can deliver. For if your little hammer hits a fault-line, you get to experience the exhilaration of a large philosophical edifice suddenly falling away, maybe even an edifice that you have helped carve out. Well, I chipped away for a good while on what follows, though I don't think I ended up hitting any deep fault-lines. But maybe I've at least knocked off a few interesting shards. In any case, I know I couldn't have done even that much without the help of many people to whom I should record a debt of gratitude.

I'm grateful for the unflagging support of my director, John McDowell. Aside from helping me develop my inchoate ideas and offering me detailed feedback on my drafts, I'm especially thankful to John for his many acts of patience, charity and benevolence in helping me to make it through when life threw up some difficulties. I'm also very grateful to the other members of my committee — Michael Thompson, Kieran Setiya, Karl Schafer and Edouard Machery — for their inspiration, feedback, and assistance over the years. Outside of my committee, I want to thank Bob Batterman, Anil Gupta, and Mark Wilson for giving me aid and support when I needed it most, and I want to thank Collie

Henderson, Kathy Rivett and Katie Labuda, the department's fantastic administrators, all of whom have worked in various ways in the background over these years to keep my show on the road.

I'm lucky to have had friends across a few generations of graduate students at Pitt. Not least, I want to thank Ori Beck, Laura Davis, Casey Doyle, Kim Frost, Jenn Jhun, Kathryn Lindeman, Gus Law, Tom Marré, Stephen Makin, Adam Marushak, Tyke Nunez, Eric Palmer, Shawn Standefer, Greg Strom, Jonathan Surovell and Laura Tomlinson for lots of good times and good memories. But I also want to offer thanks for some more specific things that I won't forget about my time here. So, thanks to: Jason Carnell, for being my partner in coffee and crime; to Rachael Driver, for looking out for me like a big sister; to Dan Addison, for sustaining my provincialism, and for passing on some great hand-me-down shirts after he had outgrown them (at the bicep); to Joe Milburn, for possessing a seemingly bottomless well of good cheer; to David de Bruijn, for scanning book chapters and emailing them to me when I was nowhere near an academic library for a long stretch; and last but by no means least, thanks to Cathy Carroll, for her good sense, her moral support, and for urging me to 加油 at just the right time.

My family back home in Australia has been there for me throughout, as always. Now that I'm a father, I know why my parents keep insisting that it that is silly to thank them for their help. But I will anyway. Thanks also to my extended family in New Zealand for being there with us during hard times, especially to Tony Backhouse, Julie Sommerville, Carlos Gelling, and Sven and Nooka.

For Cristy, my soul mate and my best mate, I don't have the words to express my gratitude. I'm hoping you'll settle for actions instead, starting maybe with a hug and a good movie. For my little daughter Nina, I also don't have the words, but conveniently I don't need them, since at this point you can only sight-read 'cat' and 'dog'. And anyway I *know* you'd prefer actions to words, so long as penguins are somehow involved.



## I. ANSCOMBE'S PRINCIPLE

### 1. INTRODUCTION

Do we always know what we're doing? In the widest possible sense of the word 'doing', the answer is clearly "No". I may be standing on your foot but not know that I am doing so. But if we take the question to be restricted to what we are *intentionally* doing, then a number of philosophers will answer "Yes". That is, they accept the following principle, which I will name after the philosopher with whom it is perhaps most closely associated:

*Anscombe's Principle.* If an agent  $X$  is intentionally doing  $A$ ,  $X$  knows that she is doing  $A$ .

I'll call this principle *AP* for short. In *Intention*, Anscombe famously defined intentional actions as "those actions to which a certain sense of the question 'Why?' is given application, [...] in which the answer, if positive, gives a reason for acting" (1963a, p.9). And she also says that one way in which the question can be "refused application" is if the agent says "I was not aware I was doing that" (1963a, p.11). That much suggests she accepts *AP*. Elsewhere, she is more explicit. In papers written contemporaneously with *Intention*, she says that "one cannot intentionally ... do something without knowing that one is doing it" (1963c, p.7, n.2), and that "it is clear that for any deed  $X$ , you cannot have intentionally done  $X$  unless you know you are doing  $X$  ... [so if] you are to be charged with a certain intentional action it is an adequate defense that you did not know you were doing that" (2008a, p.104).

If *AP* is true, that would be an interesting fact in itself, and certainly one that cries out for explanation or elucidation. The reason that the principle is worth considering is precisely that the ground of its truth may lie in the nature or essence of its subject matters, and that this is something which we can come to understand. In this respect, Anscombe's own approach to the topic does not disappoint, at least in its promise. For her, the knowledge involved in *AP* is what she calls *practical knowledge*, and this kind of knowledge, according to Anscombe, is a form of non-observational, non-inferential 'groundless' knowledge, of a piece with our *self-knowledge* of our inner mental states and mental events, at least on a familiar (if not uncontentious) understanding of such knowledge. At the same time, practical knowledge purports to be a form of self-knowledge whose objects are intentional actions, and so encompasses observable events or happenings.<sup>1</sup> Moreover, Anscombe (following Aquinas) says that such knowledge is 'the cause of what it understands', which suggests, however darkly, that our grasp of this knowledge is somehow essentially bound up with its subject matter, that it somehow lies within the nature of our intentional doings that we have knowledge of them.

As interesting as these claims may initially be, many philosophers think that on close inspection the idea that we possess such practical self-knowledge is epistemically implausible. As in the case of the more familiar (putative) forms of knowledge of our own minds, no one denies there is *something* for philosophers to say about the character of our knowledge of our intentional actions, but it is a common thought that we can capture what is distinctive of such knowledge by seeing it as grounded in inference or observation (or through some kind of justificatory framework that combines elements of both

---

<sup>1</sup> For Anscombe, the object of practical knowledge is characteristically an observable event or process, but it may also include knowledge in intention which has not yet been realized in such a doing. She also allows for intentional omissions (of which, I presume, she would say we can have practical knowledge); and she may well also allow for practical knowledge of 'purely' mental events, though so far as I can tell she doesn't engage the latter topic directly. It is also worth remarking in this connection that when practical knowledge does concern an 'outer' physical event, it is by no means clear how far such knowledge 'extends' (i.e. whether it stops at the boundary of our bodily movements or encompasses what happens to the objects and materials upon which we act). For some interesting discussion of the matter inspired by Anscombe's views, see Hornsby (2011) and Ford (2016).

inferential and observational models), as opposed to the kind of ungrounded form of practical knowledge which Anscombe offers.<sup>2</sup> On such views, there may be groundless knowledge in play in our knowledge of action, but its province will be restricted in scope so that it doesn't directly encompass what we're doing under ordinary action-descriptions like 'typing a dissertation' — that is, descriptions which characterize the agent as acting in (and interacting with) the physical world — and only takes in our knowledge or purely 'inner' items, such as the fact that I *intend* to be writing a dissertation, or *want* to, or am *trying* to (when such notions are construed as purely 'inner'). That much, say the advocates of such inferential or observational models of practical knowledge, is enough to capture what is 'special' about it, without requiring us to take on the epistemic burden on showing how such knowledge could be *altogether* epistemically ungrounded.

Now I don't think such accounts of practical knowledge really can really capture what is distinctive of it. But I don't plan to argue for that negative claim here. My own sympathies lie with a story of practical self-knowledge of the sort that Anscombe proposes, and in the next chapter I'll offer a direct argument for the truth of *AP*. (Though, as we'll see, I'll will be focusing on what I take to be the central phenomenon: intentional actions done *with* or *through* intentions.) In this chapter, I want to take up and respond to one possible reason for moving away from a story like Anscombe's which I think has been quite influential, which is that it has seemed to many that *AP* is easily refutable by counterexample. An advocate of an inferential or observational theory of practical knowledge may not *necessarily* deny the truth of *AP*, but it will certainly make such accounts look more natural.<sup>3</sup> So my first goal here is to defend *AP* against this line. To be clear, I don't think an advocate of *AP* need defend the principle tooth

---

<sup>2</sup> For views that take seriously that practical knowledge is in some fashion grounded in our observational capacities, see Pickard (2004), O'Shaughnessy (2003), O'Brien (2003), and Peacocke (2003); for views that instead offer inferential models, see Grice (1971) and Paul (2009a, 2009b). For a view that seems to mix the two together, see Donnellan (1963).

<sup>3</sup> This is because if the connection between action and knowledge can break down, then we shouldn't be surprised if we epistemic story about how the two relate involves, say, a contingent inferential process.

and nail against every possible hard case. For when a philosopher say that a principle like *AP* is true, what they may mean — and what I mean — is that it is true in the ‘standard’ or ‘ordinary’ or ‘normal’ case; that is, it is true for cases which exhibit the proper functioning of mature human agency and mature human knowledge. From this perspective, we can investigate a principle like *AP* without necessarily stopping to keep track of how things look when (say) there are repressed or subconscious intentions on the scene, or when we are dealing with an agent hampered by illness or working under some type of pathological condition.<sup>4</sup> But the problem, as we’ll see in §2, is that the putative counterexamples that have been offered against *AP* in fact arise in perfectly quotidian contexts, and so the restriction of our attention to how things look in the ‘ordinary’ case may not seem to help.<sup>5</sup>

The burden of this chapter will be to show that in this respect the appearances are, in a way, misleading. Although it is clear that it is possible to imagine putative counterexamples of the relevant sort as taking place in ordinary *settings*, I’ll argue that does not mean that in the final analysis we should say that they properly fall under the ordinary *case*. My strategy to defend setting such cases aside will be a familiar one: I’ll suggest that the relevant cases draw whatever force they have from being under-described, and when properly fleshed out, their threat to *AP* dissipates. Of course, there is always a concern that such a response to *prima facie* troubling counterexamples will ultimately be grounded in mere

---

<sup>4</sup> Although it isn’t discussed in the commentary of her work, Anscombe herself in fact allows for an exception of the first sort just mentioned. The last remark I cited from her above endorsing *AP* is qualified: Anscombe says that we always know what we’re doing “except in a psychoanalytic sense in which there can be unconscious intentions” (2008a, p.104). Note that I don’t at all mean to deny that the study of human agency in its ‘non-standard’ guises can offer no insight, including philosophical insight. I just mean that I won’t approach the topic that way here. If (some) actual argument is wanted in defense of fixing the subject matter of the discussion in this way, I find helpful the views sketched, from different perspectives, by Neta (2009), Boyle (2012), and Williams (2015), though none of these authors is specifically concerned with the explanation of intentional action.

<sup>5</sup> See Setiya (2008; 2009), Paul (2009a, 2009b), and Gibbons (2010); for less recent skepticism about *AP*, see the references in Velleman (2007, pp.18-9, *n.6, n.7*; p.102, *n.23*).

stipulation. I will try to show, to the contrary, that my handling of these hard cases is not at all arbitrary but rather tracks the distinct ways in which advocates of *AP* and the advocates of the putative counterexamples to it may conceive of the will as an epistemic capacity. In doing so, I'll argue that when the putative counterexamples are left in their original under-described state, they can be seen to involve an element (or several elements) of *accidentality* that mean they are not properly classed as intentional actions. To that end, in §3 I'll draw on some of Anscombe's remarks in *Intention* about what is presupposed in intentional action and practical knowledge, and I'll supplement that discussion in §4 with some considerations about the crucial role of aspectual and temporal distinctions among verbs in our grasp of intentional action. The intended upshot, canvassed in §5, is to show that it is questionable that the putative counterexamples are anything more than that, and that a refusal to take them seriously can have a principled ground.

## 2. THE (PUTATIVE) COUNTEREXAMPLES

The most famous putative counterexample to *AP* is due to Donald Davidson. He asks us to imagine a person writing through a large stack of papers interleaved with sheets of carbon paper:

[A] man may even be doing something intentionally and not know that he is [...] (A man may be making ten carbon copies as he writes, and this may be intentional; yet he may not know that he is; all he knows is that he is trying.) (Davidson 1971, p.50)

A few years later he went a little further, or made his position a little more explicit:

It is a mistake to suppose that if an agent is doing something intentionally, he must know that he is doing it. [...] [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.92)<sup>6</sup>

We can extract from this example a pattern to generate more of the same: all we need is a set up where (i) the agent has some reason to doubt that they're successfully executing an intention; (ii) they are in fact successfully executing that intention; but (iii) their success is fundamentally occluded from the agent in some way. I'll call any scenario that fits (i)–(iii) a 'carbon-copy case'.

In this sense, carbon-copy cases seem pretty easy to come by. Here is another example. Your car is out of gas, so you go to refuel it at a gas station. You put the fuel hose in your fuel tank, pull the lever tight, and put in place that little lock that conveniently keeps the handle in place as the gas pumps, leaving you to be enthralled by the gripping advertisements playing on the fuel dispenser. You then notice that the gauge isn't moving from 0.00 gallons or 0.00 dollars. So you go inside to get the attendant. But when you come back, you find the little lever has switched off, and your gas tank is full. It turns out the gauge was broken, though the pump worked fine; you just couldn't hear it over the noise of the advertisements.

If such cases are indeed cases where an agent acts intentionally without knowledge that they are doing so, *AP* is false. But is that all that would follow? Already in Davidson we see not only the suggestion that his example shows that *AP* is wrong, but also that a weaker principle fails, one that would connect action only with belief with a high credence.<sup>7</sup> Some philosophers follow Davidson in this

---

<sup>6</sup> In these remarks Davidson is as much concerned with putative connections between the mental state of intention and cognition as he is with connections between intentional action and cognition. But I'll set aside consideration of how intention relates to cognition here. For a small sample of the large literature on this topic, see Hampshire & Hart (1958), Anscombe (1963a, §52), Bratman (1999, Ch.8), Ludwig (1992), Vermazen (1993), and Holton (2009, Ch.2).

<sup>7</sup> As Davidson is sometimes misrepresented on this score, it is perhaps worth noting that he did not deny that an intentional action must be known to an agent under some description, nor did he think carbon-copy cases show that there is no entailment between intention and belief. His main target in presenting his carbon-copier seems to have been Grice (1971), who thought intention required certainty in success. But Davidson is still a 'cognitivist' about intention, as

respect. Others suggest go yet further and suggest that we can devise further cases that show principles linking action to credence of any degree would be false. For example, Kirk Ludwig (1992, p.262) asks us to imagine a case where he has accidentally left his car parked in his neighbor's yard, and he realizes the battery is dead. The next morning, when he tries to tell his neighbor that he doesn't think the car will start — he believes there is “zero probability” that it will start — his neighbor starts to threaten him. Knowing his neighbor has a violent past, Ludwig decides to put on a show of starting the car to appease him. As it happens, the car with the dead battery is actually in *another* neighbor's yard, and *this* car has a charged battery. So when he tries to start the car, he succeeds. For Ludwig, this is a clear case of intentional action, even though he (the agent) believes there is no chance of success. Philosophers who go in for cases like Ludwig's will then presumably need to revise *AP* even more dramatically, perhaps by switching to a principle that does not connect intentional action *sans phrase* with knowledge or belief, but with actions that are, as it were, ‘closer’ to the agent, such as basic actions (one performed ‘straight-off’, or without taking any means) or bodily movements, or attempts (construed as a purely mental event).

The literature offers a wide range of options on this score.<sup>8</sup> That there *is* a wide range of views in this respect ought, I think, be of comfort to the advocate of *AP*, as it at least suggests that opposition to the principle is not necessarily based in some monolithic immobile intuition, but perhaps instead shows that there is some substantive theory lurking behind the counterexamples.<sup>9,10</sup> In any case, that is the

---

that term is sometimes used. See Davidson's (1979, p.91) for the requirement of knowledge on action, and also Hunter (2015) for some interesting further discussion on that score. For his allowance of entailments between intention and belief, see Davidson (1985).

<sup>8</sup> For discussion of various views along such lines, see Donnellan (1963), Danto (1963), O'Shaughnessy (2003), Paul (2009a, 2009b), O'Brien (2003), Peacocke (2003), Pickard (2004), Setiya (2008; 2009), and Gibbons (2010).

<sup>9</sup> I should note that other philosophers also present carbon-copy cases of a more *recherché* flavor, employing various scenes of pathology or impairment and such. (See, for example, Kieran Setiya's (2008, p.390-1) anaesthetized arm case, or Sarah Paul's (2009b, p.553) example involving an imaginary case of Alien Hand Syndrome.) I find it far less plausible

possibility that I want to explore. Or, more precisely, the question I want to consider is this: What am I, *qua* advocate of *AP*, supposed to believe about intentional action and practical knowledge such that I should be worried about the possibility of carbon-copy cases? The answer, I think, is that the advocates of *AP* must view the will, when construed as an epistemic capacity, as a substantively *independent* capacity, one whose exercise can issue in knowledge with very little assistance from other capacities, and in particular from our observational capacities. And if we think that practical knowledge must come very easily, then we will say the same about its object, intentional action. This, I think, is why so little can seem to go into a carbon-copy case: the only elements we need to find at work in an agent in the typical carbon-copy case for the agent to be intentionally doing *A* is just (say) an intention or desire to do *A*, and a successful attempt at *A*-ing. And this, in turn, is why the examples may seem so easy to come by, for our ability to know what we're doing is not quite so *awesome* as to be able to keep up. Of course, this a bit vague; but the question I want to take up is whether an advocate of *AP* needs to view the will in this way, or whether instead she might think that intentional action and practical knowledge is a humbler affair, in a sense we can make more precise as we go.

We can begin on bringing out the contrast between these two pictures if we consider a common interpretation of the main source of *AP* — Anscombe's *Intention*. The sort of view of the will which I'm describing as 'independent' or 'awesome' is one that is often attributed to Anscombe, as she is routinely

---

that such cases involve intentional action to begin with, but the response I sketch below to the more familiar types of carbon-copy cases will still apply to such cases, though I won't pursue the details.

<sup>10</sup> It is also worth remarking that the casual slide in the discussion of these cases between whether the agent knows what they're doing and whether they believe they are, or have a certain credence that are, may itself also betrays a very substantive assumption about the character of practical knowledge: namely, that it is the kind of knowledge that could be viewed as a species of belief in the familiar way. My own view is that this is false: practical knowledge may entail belief, but is not properly viewed as a species of belief at all. I don't think that this is irrelevant to the discussion of carbon-copy cases, and a fuller treatment than the one I'm offering might pursue the extent to which this assumption is at work in such examples. But I won't pursue this angle here, and I think I can take the sting out of carbon-copy cases without it.



billed as offering an account of practical knowledge which, in effect, can require no relevant non-practical knowledge to support it.<sup>11</sup> The attribution is understandable, because Anscombe gives examples in *Intention* which may seem to suggest as much. In one example she discusses, someone is writing with their eyes closed; in the other, a person is directing the construction of a building without seeing or hearing anything about its progress (1963a, pp.53, 81-2). In both cases, Anscombe thinks the agents lack any knowledge of the sort which she dubs “speculative”, that is, knowledge which is “derived” from its object which may be relevant to the task at hand. And by this expression — which I’ll also use — she means to include not just observational knowledge but also knowledge garnered through inference or testimony. If we generalize from such cases, it may seem that Anscombe’s view is that “even when [empirical or speculative] feedback [concerning the matter about which one has practical knowledge] is unavailable the practical *thought* of such an agent suffices on its own to be practical *knowledge* of what is happening” (Schwenkler 2015, p.17). And so practical knowledge, by Anscombe’s lights, appears to be the product of a free-floating capacity, even one with something of the divine about it.<sup>12</sup> This is the picture of the will which I suggest lies behind the carbon-copy cases.

Did Anscombe really hold such a view across the board? For my part, I’m not entirely sure what she is up to with these two particular examples. In the case of the person writing with their eyes closed, Anscombe says “the essential thing he does, namely to write such and such, is done without the eyes” (1963a, p.53).<sup>13</sup> But she also muddies the water by distinguishing between whether the agent is *writing* and

---

<sup>11</sup> See Velleman (2007, pp.18-9), O’Brien (2007, p.168), Paul (2009, p.2), and Gibbons (2010, p.77).

<sup>12</sup> For this allusion, see Velleman (2007, p.xv) and Paul (2009, p.1).

<sup>13</sup> Of course for many philosophers, such a setup does not itself show, as Anscombe thinks it does, that if the agent knows that he is writing then it is not through observation. This is because it will be protested that such an agent will still rely upon other forms of observational knowledge, such as kinesthesia and proprioception. Anscombe has a view about such matters, but this is not the place to pursue it. See McDowell (2011) for an important discussion of Anscombe’s view.

*writing legibly*, with perception serving as an “aid” to the agent succeeding in doing the latter but not the former, which makes it harder to know just what the object of practical knowledge is supposed to be in this case. More importantly, when she returns to this sort of example later in her book, she notoriously appears to take seriously the idea that practical knowledge of writing one’s name may be “the same” whether the words are appearing or not (1963a, p.82), and we still lack, I think, an interpretation of the relevant passages that renders them coherent.<sup>14</sup> And her other example of the project manager involves a different kind of problem. The setup here is that the director gives orders to his underlings about how to erect the building, though he can’t himself see or hear about the construction as it proceeds. However, Anscombe also says that the director is working with a plan that settles every detail of the construction in a “right” order. And although the director has no relevant speculative knowledge about the construction, he is described as possessing a “superhuman” imagination which “takes the place of the perception that would ordinarily be employed by the director of such a project” (1963a, p.82). So should we then say that the director has practical knowledge? Well, I’m inclined to say, sure, why not? You may disagree. Either way, it isn’t clear what the answer tell us about a creature stuck with plain old human powers, as we all are.

Now a more patient discussion of these examples and their wider context would, I think, bring out that Anscombe’s purpose in imagining cases where we strip out the normal role of speculative knowledge in agency is to bring out the distinctive features of practical knowledge that may be obscured from view when we consider ordinary instances of practical knowledge where such knowledge is at work

---

<sup>14</sup> For different takes on this difficult passage, see Haddock (2011), McDowell (2010), and Özaltun (2016). Of course, the text may just be incoherent. The reason why there is a temptation to think not is that Anscombe is very obviously aware of how her remarks may *seem* to deny that knowing that such-and-such must entail that such-and-such is the case. In so far as we don’t grasp what she is up to here, we may not really grasp what her account of practical knowledge involves. For my part, I make no attempt in chapter 4 — where I discuss Anscombe’s conception of practical knowledge in more detail — to accommodate what Anscombe says in this passage, though I believe it is in fact relevant; and so my account there is incomplete too.

in tandem with other forms of knowledge, whether speculative or not. But although I don't want to take a stand on Anscombe's particular examples, I also don't want to rule out the possibility of analogous cases. That is, I want to allow that it may be possible for an agent to have practical knowledge of what they are doing even if they can't (say) see what they are doing right at that moment. What I want to reject is the thought that we should generalize from the possibility of such cases. *Contra* the reading of her view just aired, this is not the right way to read Anscombe, because she thinks the typical case is one wherein an agent in possession of practical knowledge "makes use of his senses or of reports given him, the whole time" (1963a, p.88). And even if Anscombe did think practical knowledge was characteristically independent of speculative knowledge, then at that point we should go our own way. For even if we allow that cases of 'independent' practical knowledge are *possible*, as I want to do, such cases must, I think, involve special circumstances, and be intelligible only because the agent is in an otherwise favorable epistemic position vis-à-vis their actions. We'll consider what such cases may involve later on. For now, the important point is that by Anscombe's lights the standard case isn't at all like this, but is rather one wherein the agent makes use of speculative knowledge "the whole time".

If that is how things are in the standard case, then the picture of the will which I claim is at work in the counterexamples to *AP* is wrong. But we need to do more work to join the dots here to see just how that picture is at work in carbon-copy cases. Thus far, I've suggested that an advocate of *AP* need not view the will as independent of speculative knowledge in the manner suggested by Anscombe's two examples. But the *advocate* of such examples may not see the significance of the point of setting aside this (faulty) picture of the will, for she may well be quite happy to say that intentional action and practical knowledge *are* dependent upon various other forms of knowledge. However, there is still room to move between requiring *some* assistance and requiring *incredibly little* or *none*, and my aim in the remainder of the chapter is to show that in so far as the counterexamples appear convincing, it is because they exploit this space. To see why that is so, we will need to consider in more detail just what needs to be in place for an agent to act or have knowledge of what they are doing.

### 3. THE PRESUPPOSITIONS OF PRACTICAL KNOWLEDGE

In this respect, Anscombe also provides a useful starting point. In §28 of *Intention*, she introduces an interlocutor who finds it reasonable enough that we have groundless knowledge of the position of our limbs or even of our bodily movements, but who is nevertheless baffled by the idea that such knowledge could encompass our actions under ordinary descriptions like ‘painting a wall yellow’. She says in response to this bafflement:

[T]he topic of an intention may be matter on which there is knowledge or opinion based on observation, inference, hearsay, superstition or anything that knowledge or opinion ever are based on; or again matter on which an opinion is held without any foundation at all. When knowledge or opinion are present concerning what is the case, and what can happen — say *Z* — if one does certain things, say *ABC*, then it is possible to have the intention of doing *Z* in doing *ABC*; and if the case is one of knowledge or if the opinion is correct, then doing or causing *Z* is an intentional action, and it is not by observation that one knows one is doing *Z*; or in so far as one is observing, inferring etc. that *Z* is actually taking place, one’s knowledge is not the knowledge that a man has of his intentional actions. By the knowledge that a man has of his intentional actions I mean the knowledge that one denies having if when asked e.g. ‘Why are you ringing that bell?’ one replies ‘Good heavens! I didn’t know *I* was ringing it!’ (Anscombe 1963a, pp.50-1)

Anscombe is here concerned, in the first place, with acting with an intention and, in the second place, with succeeding in one’s intention. Or, as I’ll say for short, the topic here is doing one thing in doing another, in the (slightly peculiar) sense in which what the ‘things’ the agent does may be parts or phases of the one action, as when I write my name by holding and moving a pen in a certain way. And from here on in, this is how I will understand ‘intentional action’, as involving doing one thing in doing another.

Now I take it that Anscombe’s overriding point in this passage is that although a condition of an agent having practical knowledge of her doing *Z* — or, more fully, doing *Z* in doing *ABC*, where ‘*Z*’ may

be a description that goes ‘beyond’ a bodily movement — is that she has knowledge or belief<sup>15</sup> that may be grounded in epistemic sources like perception or inference, that does not mean that such practical knowledge is itself epistemically grounded in such sources. Using ‘cognition’ as a cover term for such knowledge or belief, I will say, using Tyler Burge’s (1992, 2011) familiar terminology (without committing to the detail of his account), that Anscombe takes such cognition to *enable* practical knowledge, even if it doesn’t epistemically *justify* it. For her purpose of responding to the interlocutor, Anscombe could be read as suggesting that such cognition is *sufficient* to enable practical knowledge of actions like painting a wall yellow. However, I think Anscombe also takes such cognition to serve as a *necessary* condition or a presupposition of *both* the agent’s acting intentionally *and* their having practical knowledge that they are doing so.<sup>16</sup> Or, if this is not her view, it is in any case the position I want to etch out. Thus, on this view, for an agent to be doing *Z* in doing *ABC* *and* to have practical knowledge of the same, they must have relevant cognition of “what is the case” and of “what can happen — say *Z* — if one does certain things, say *ABC*”. I’ll call the first condition “the Case Condition” and the second “the Belief Condition”.<sup>17</sup>

Do these two conditions necessarily hold of both intentional action and practical knowledge? In light of the loose description so far offered, it isn’t easy to say decisively, but there is surely a strong

---

<sup>15</sup> I assume that talk of ‘belief’ in place of Anscombe’s ‘opinion’ is a harmless switch. That may not be so, as she bothers to use both in *Intention*, though her usage is not entirely clear. (See also the following note.)

<sup>16</sup> See Anscombe (2011b, p.16) for a clear statement that she takes acting with an intention to require at least such an instrumental belief. (Though she there formulates the claim in a slightly stronger way.) Compare also her remarks in §22 of *Intention* that for an agent to act with the intention of trying to bring about some future state of affairs it must be that “we can understand the agent’s thinking it will or may be brought about by the action” (1963a, p.35).

<sup>17</sup> In the next chapter, I’ll go on to give this principle a few other possible names as part of a project of making its content more precise, as I understand it. Note that not everybody seems to think that Anscombe is offering two conditions here. Ford (2016, p.143) thinks there is a single condition, the one which I’ve dubbed the Belief Condition. To my mind, the passage is hard to parse if it only refers to one condition.

temptation to say that they do. Such conditions can, for instance, be read so that they are even consistent with the picture of the will that I just suggested we ought to discard. More generally, no one will deny that intentional action (let alone practical knowledge) will require some cognition on the agent's behalf about the relevant matters and some kind of background skill or ability or know how that they can bring to the task at hand. And that's another way of saying that some form of the Belief and Case Condition are surely uncontroversial. What philosophers may disagree about is the details.

Before we say more about that, let us consider how these two conditions, when read a certain way, might work to rule out carbon-copy cases. As a preliminary, take winning a lottery. This is a type of action which is very widely (if not universally) taken by philosophers to be something which we cannot do intentionally. Why not? Because the outcome is “too lucky”, or “out of the agent's control”, or “too accidental”. How is the notion of the accidental — to run with that term — to be understood in this context? Not in any one way, I don't think.<sup>18</sup> After all, at least in ordinary speech we often enough do say that someone pulled something off with a bit of luck or was lucky in doing what she did. Still, there are ways in which an intentional action can't be a mere happy accident, and we can invoke these to explain why we can't win a lottery intentionally. The most obvious thought is that winning is simply too accidental to count as intentional in the sense that there is an exceedingly low probability that an instance of an act of the first kind (*buying a ticket*) will lead to or be followed by an instance of an event of the second kind (*winning the lottery*).

---

<sup>18</sup> There is of course much philosophical debate about how to understand notions like *accident* and *luck*, both as they apply to intentional action and to knowledge. (And not everyone will want to run these terms together.) For relevant discussion in the case of intentional action, see Audi (1986), Mele & Moser (1994) and especially Hanser (2000). In the case of knowledge, for some recent discussion, see Yamada (2011), Schafer (2013), Pritchard (2014) and Hetherington (2014). Although the notion of non-accidentality will be important to what follows, I don't plan to offer any kind of analysis of it, but only highlight some important marks of the concept.

Now if we construe the Belief Condition so that it requires the agent to have beliefs with a certain likelihood of success, then we could obviously say that winning a lottery can't be intentional by dint of failing to satisfy that condition. And that would be a very natural way to construe the Belief Condition. As Anton Ford (2016, p.143) remarks, the typical characterization of such instrumental or means-end cognition in practical philosophy is one of "a run-of-the-mill belief whose propositional content is that, if one kind of thing were to happen, another kind of thing would happen".<sup>19</sup> That is just the kind of content we would imagine the agent to hold with various degrees of likelihood, depending on the agent's past experience and their knowledge of the circumstances.

Although I don't want to deny that such beliefs may be operative in agency, and even though this gives us a sense of how failing to satisfy the Belief Condition can render an action unintentional, this construal of that condition does not exhaust everything that we want to say about the content of such means-end beliefs (or means-end knowledge), and so does not exhaust the ways in which a doing may fail to count as intentional. For instance, a different issue concerns whether we should allow for only beliefs to play a role here, or also knowledge, as I'm suggesting (following Anscombe). Some philosophers argue by intuition about cases that belief isn't enough, and that we need instrumental knowledge for intentional action.<sup>20</sup> That seems to me to be a hard sell, not least because it makes it difficult to see how we could ever do anything intentionally for the first time, and where doing so would be a way to learn that the instrumental thought employed is indeed correct. A better approach, I think, is

---

<sup>19</sup> In fact, Ford sees this construal as pretty well universal, with the exception of Aristotle. Thus, he sees such a construal at work in Anscombe's text from §28 which I earlier cited. I'm not sure why though. It may be that he thinks that she means her action-descriptions '*ABC*' (or 'doing *ABC*') to be read as variables that range over *kinds* rather than particular actions or events. But Anscombe's notation seems to me to be intentionally loose, defying any sharp interpretation, and so not forcing that reading. We can't usually stick singular terms together as she does in '*ABC*', regardless of what they range over.

<sup>20</sup> See Gibbons (2001, pp.589-9; 2006, p.90; 2010, p.81), Sliwa (2017, p.129), and perhaps Schwenkler (2015, p.23) for accounts that require an agent to know the instrumental content operative in an action for it to be intentional.

the disjunctive one Anscombe proposes, which will allow us to invoke instrumental knowledge when it is needed, but resort to mere (true) belief when it is not.<sup>21</sup> And that would give us a distinct way for something to fail to be an intentional action: that the agent not only doesn't have a strong credence in the outcome, but doesn't know that her action is a way to bring it about.

Even once we've allowed for either means-end knowledge or belief, we might also want to say more about how the content of such cognition should be understood. Thus, I think we should say, with Ford (2016), that instrumental cognition may itself need to contain certain perceptual demonstrative elements, so that we have beliefs in play with shapes like "If I push *this* button in *this* way, then the bomb will explode" (*cf.* also Grünbaum 2013). And so, again, we would have a distinct way in which an agent might fail to satisfy the Belief Condition, for an agent may have an instrumental belief linking kinds of acts, but still lack the requisite perceptual cognition about the matters at hand. Moreover, I would argue that a construal of the relevant means-end cognition as a relation between kinds of acts will, perhaps surprisingly, routinely fail to properly capture the causal content of such cognition. For, 'Humeian' prejudices aside, it is a mistake to think that causality must consist in a relation between events, and so too is a mistake to think that cognition of causality must have that shape. If such an assumption did determine what could fall under the content of the Belief Condition, the condition would leave out the causal content itself — that is, the particular causal route that an agent adopts to effect their end.<sup>22</sup> It may be, for instance, that the kind of causal cognition involved in my painting a wall yellow — that I think that dunking the brush into the bucket is a way to get paint on its hairs, and that pushing the brush along the wall is a way to get the paint to stick on the wall, and so on — is best construed not as

---

<sup>21</sup> It would only help my argument in what follows if only instrumental belief were not enough for intentional action, so to make life harder for myself, I'll adopt Anscombe's weaker view.

<sup>22</sup> I take the term 'causal route' from Makin's (2000) discussion of Anscombe's notion of causality as 'derivativeness', which we'll discuss in detail in chapter 3.



involving relations between events, rather as involving a type of causality that may be ‘internal’ to the events themselves.<sup>23</sup>

Such remarks about causality will seem a bit dark, and they won’t be much clearer until we consider causality in more detail in chapter 3. For now, I make the point only to urge that the Belief Condition may involve more than what philosophers have often taken to be involved in means-end cognition, and there may thus be more ways to lack such cognition. Indeed, in the following chapter, I will argue that the Belief Condition is a necessary condition in a way that brings out yet another aspect of its content. But even abstracting from debate about the specific form the condition may take, we can use the Belief Condition to rule out certain carbon-copy cases as involving intentional action much in the same way as we can use it to rule out that winning a lottery could be intentional. Thus, recall Ludwig’s example of the agent who supposedly starts the car intentionally even though he thinks there is “zero possibility” of that happening, because he believes the car battery is dead. Ludwig’s case clearly fails the Belief Condition: he doesn’t have the right belief (or knowledge) to count as acting intentionally, at least under the description ‘starting this car’. He doesn’t believe that if he turns the key he can start the car, because he thinks the battery is dead. All he can do intentionally is *feign* starting the car. When the car does start, it’s a happy accident. We can offer a similar explanation for my refueling example. In that case, although I start out with the belief that pulling down the gas lever is a way to refuel the car, once I think the fuel dispenser is broken I revise my belief and I now think the dispenser is broken. (I revise my belief yet again when I learn that I was wrong.) Since I now think the fuel dispenser is broken, I could just as well have taken it out of the car before I went inside to get help, and so it is a happy accident that I fill up my car by leaving it in. Of course, in both this case and Ludwig’s, we can still say that we’re both *responsible* for what we do, but to say that we don’t need to say that we act intentionally. And so we can

---

<sup>23</sup> For some discussion of this line of thought, inspired by Anscombe, see Hornsby (2011).

accept that neither Ludwig nor I have practical knowledge of what we're doing without giving up *AP* because we don't act intentionally.

We seem then to have defused the threat posed to *AP* by at least some carbon-copy cases. But what about Davidson's example? This case seems different, precisely because it *doesn't* seem plausible to say that the agent lacks the kind of causal belief needed to satisfy the Belief Condition.<sup>24</sup> For Davidson's agent thinks something along these lines: pushing down the pen in such-and-such ways could be a way to get the impressions to go through onto all the pages. So this case isn't like Ludwig's, or like my own. (It would be like those cases, if, say, the agent thought that her pen was broken, but was just pressing down hard on the stack of paper for fun, all the while writing through all the copies.)

To properly deal with this case, we need to instead consider the role of the Case Condition. At a glance, it may not seem obvious that this could help — for what is Davidson's agent's supposed to know or believe about her situation which she doesn't, aside from the fact that he isn't writing through all the sheets? But here it matters how we construe the function of the condition. And as with the Belief Condition, I think philosophers tend to presupposes too narrow a construal of what is presupposed in involved in this condition, and I'll argue that with a better account of the condition we will be able to bring out its relevance to Davidson's example. But to do that, though, we need to first take up a different line of response.

#### 4. THE ROLE OF PROGRESSIVE ASPECT

Davidson in fact offers us very little context to his example. It may seem he doesn't need to. It is after all a very simple sort of scenario. However, Michael Thompson (2011) suggests that the case is in fact under-described, and that we need to imagine different ways of filling in further details before we can

---

<sup>24</sup> See McDowell (2013, p.393, *n.16*) for this point.

come to a clear verdict about it.<sup>25</sup> So here is one way that Thompson thinks Davidson's example could be spelled out further:

[Y]ou write on the top sheet, trying to make a good impression to get through all the carbon, then look to see if your impression made it through all of them. If it did, you stop. If it didn't, you remove the last properly impressed sheet and begin again. If necessary, you repeat. Even the man who has to go through five stages is all along, from the first feeble impression, making ten copies of the document, and he knows it, all along. (Thompson 2011, p.210)

This, Thompson tells us, is "the more ordinary case". And he is surely right about that. For we have to wonder, in reading Davidson's example: Why doesn't the agent just have a look to see how he is doing? Why should it matter if he can't *right at that second* see that he is managing to write through all the sheets? It is, after all, in the nature of the task of trying to write through a stack of such papers that one can't see everything that one is doing as one is doing it. Would it not require a special context to think the agent doesn't need to look to see how they're doing?

It is important, though, to notice that Thompson in fact bundles together two sub-cases into his description of the ordinary case. In the first sub-case, you succeed in one go; in the second, you don't. And what Thompson describes as part of the second sub-case — the possibility of various false starts or attempts made on the way to success — is surely a part of our ordinary predicament. But how does this re-description of the case not simply beg the question against Davidson? I mean: what stops Davidson from saying that in the first sub-case, the agent *doesn't* know that he is intentionally writing on all ten sheets *as* he is doing so, and in the second sub-case, he *does* know that he is intentionally writing, but *not* that he is intentionally writing through all ten sheets, but rather only knows that he has written through (say) the first five, if that is what he has done by that point?

---

<sup>25</sup> For similar responses, indebted to Thompson, see Wolfson (2012) and Small (2012).

A rejoinder we could offer on Thompson's behalf would begin with the claim that these questions are only pressing because we're ignoring the essential role of *progressive verbal aspect* in our ascriptions of intentional actions. When we say that someone is "writing on all ten sheets" we characterize them as engaged in an ongoing process, of which there will ordinarily be more to come. As Kevin Falvey (2000, p.22) points out, one of the central semantic properties of this form of predication, particularly in application to human and animal action, is its *broadness*. Our descriptions of what people are up to — say, 'doing *A*' — are characteristically semantically broad in that the description can be true of an agent even if they are not, at that particular point, doing anything that would clearly count as doing *A*, or even if they are not clearly doing something toward doing *A*. It can be true, to take an extreme sort of case, that I am writing a dissertation even though the moment at which the sentence 'He is writing a dissertation' is tokened is one where I happen to be (say) staring at a picture of Frege, thinking about his moustache. But the broadness of the progressive doesn't just allow for pauses in an ongoing bit of intentional activity; it also allows, to varying degrees, for slips, mistakes, false starts and retries of various sorts, without yet necessarily entailing the falsity of an action-description, though of course this semantic broadness will surely have a breaking point.<sup>26</sup>

Returning then to Thompson's description of the ordinary case, if we keep the role of progressive aspect firmly in view, we can see that there is no in principle problem with the idea that an agent can be doing something even if there are hiccups and restarts along the way. There is no problem, for example, in saying that I'm making tea even though when I started out doing so I accidentally reached for the coffee, only to realize my mistake and redirect my movement for the Russian Caravan — I'm still making tea the whole time. In the same way, we can say that Davidson *is* writing through all ten sheets even though right now he has only succeeded in getting through five so far ... and has now just started on the rest ... is now really struggling with the last few pages ... and so on. He can still be said to

---

<sup>26</sup> For a discussion that perhaps pushes that point too far, see Haddock (2011).

be writing through all ten the whole time. Getting clear about the role and centrality of progressive or imperfective aspect thus lets us see why we can say, as in Thompson's second sub-case, that in the ordinary case the agent is intentionally writing on all ten copies the whole time, and knows he is doing so.

But how does this help with the first half of the question? Why, to recall the complaint, can Davidson not reply that Thompson's first sub-case simply repeats his own description of the carbon-copier, and then conveniently elides the crucial detail that the agent doesn't believe he is succeeding as he is writing? Could Davidson not also in turn claim that Thompson's first sub-case is itself under-described? Could he not say that we need to fill it out by adding that the agent next goes on to take a look at this work and says "Oh, I did indeed write thorough all ten. Well, you learn something every day!"

How does the progressive help here? Is the idea that the progressive can also, as it were, swallow up the doubt or disbelief that would exist in the agent which Davidson thinks renders him ignorant of his success, thus making the re-description just suggested seem unnecessary? To some extent, I think the answer must be "Yes". For in so far as we accept that the progressive can allow for mistakes and restarts in the manner just sketched it must surely be able to carry with it whatever little bits of doubt or lack of confidence such events might bring with them in the agent's mind, so that he still knows what he is doing the whole time even in the light of a few dark moments. At some point, though, such doubt may be strong enough to render the case akin to the two we earlier set aside for failing the Belief Condition. If someone tells me that I didn't actually interleave enough sheets of carbon paper into the pile to begin with, then I'll have to drop the belief that pressing down hard with the pen in this way could be a way to write through all the sheets, at least as things stand. The question, though, is why Davidson can't inhabit the space between these extremes, where there is some doubt in one's success, but not enough to require revising one's belief about the possible efficacy of one's means.

It will help, I think, if we think of Thompson's invocation of the progressive as working in two stages in his response to Davidson. First, the progressive gives us room to say that an agent may in some cases be said to be intentionally 'doing *A*' even if at that moment they are not observably doing so. If we leave unspecified just *what* the progressive makes room for, then there doesn't seem to be anything about the invocation of the progressive in itself that Davidson, *qua* pusher of his carbon-copy case, should be worried about. His example would work just as well if, for example, he set it up so that the agent begins to write through all ten copies, puts the pen down for a little to reflect upon the state of the world, and then continues to write. In the latter scenario, Davidson will still insist (now helping himself to the progressive) that the agent is intentionally writing through all ten sheets but doesn't know that he is doing so. And we could say the same on Davidson's behalf for allowing the progressive to make room for slips, restarts, hesitation, and various other ways of intermittently achieving one's goal. Invoking the progressive to urge a generous sense of *what can count as doing intentionally* is then, in principle, a move that advocates of carbon-copy cases can happily help themselves to while still making trouble for *AP*.

The second and crucial step in Thompson's use of the progressive lies in the introduction of a *specific* item into what the progressive soaks up: namely, that the agent gets to *check* on her progress, as *part* of the ordinary course of her activity.<sup>27</sup> It is, I think, this specific use of the progressive to allow for observational knowledge to facilitate a more generous conception of *what can count as acting with practical knowledge* which I think must mark the central disagreement between Thompson and Davidson. For Thompson, his first sub-case does not beg the question against Davidson, because Thompson assumes

---

<sup>27</sup> In his *Life and Action*, Thompson (2008, p.91) is explicit that in rallying against the tendency of philosophers to view action as "discrete or atomic or pointlike or eye-blink-like units", he is trying to put in place an alternative picture which takes seriously the progressive character of action not only in involving activities that extend over time, but also specifically that includes, as their parts, the activity of *checking up on one's work*. As he says of someone *assassinating a political opponent*, on his approach, we will find an "attractive and likely resolution into *reaching for, raising, aiming, and firing* a gun, to saying nothing of *checking to see if the victim is done for and repeating as necessary*".

that it can be part of intentionally writing through all ten sheets that the agent may need to look to see how he is doing. If he hasn't succeeded, and needs to keep going, the first sub-case effectively turns into the second sub-case. The difference between Thompson and Davidson must then be that Davidson will want to insist that we can still have cases where an agent can act intentionally even if the agent *doesn't* get to check on his progress and doesn't know that he has succeeded. This must be how Davidson would view Thompson's first sub-case, if it is still to serve as a counterexample to *AP*.

It is not then surprising that Thompson goes on to remark that he reads Davidson as having had just such a scenario in mind:

The other one who doesn't know it, Davidson's man, must be under some strange mafia threat: he gets one chance, no checking, and he's dead if he doesn't manage it. Well, 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 2011, p.210)

Now I think Thompson's diagnosis here is correct, but his description of his own position, or one we could draw from it, unnecessarily concessive. We can, I will argue, use Thompson's line to provide ground for reasonably contesting the usage of 'intentional' in application to "Davidson's man". But this will require unpacking Thompson's remark a little further. For an obvious objection here is that Davidson's example *isn't* analogous to someone 'intentionally' winning a lottery at all. In one way, that is borne out by our earlier discussion wherein we explained the accidental and so non-intentional character of winning a lottery as a consequence of the agent lacking the requisite belief about how one wins a lottery. As we saw, it doesn't seem right to deny that the agent in Davidson's case lacks such a belief. The question, though, is whether there isn't some *other* salient sense in which we might elucidate the thought that the situation is akin to winning a lottery. I suggest there is.

## 5. TWO ASPECTS OF NON-ACCIDENTALITY

The basic idea is that in Thompson's last remark we can see the suggesting that Davidson's agent is like someone winning a lottery *because* he doesn't get to check. In saying this I don't mean, of course, that one normally gets to check how one is doing at the lottery before it is drawn. I mean that in so far as an agent is in a carbon-copier situation as Thompson reconstructs how things are for "Davidson's man", where he has no chance to check on his progress, then his success may be analogous to that of winning a lottery in being too accidental to count as intentional. In other words, there are different ways in which an action can be non-accidentally brought about. One dimension of (potential) accidentality is indexed to the Belief Condition, to the role played by an instrumental belief, or its absence, in the agent's activity. But another dimension is tied to the Case Condition, to the requirement that the agent possess the relevant speculative or observational input about how things are going with their work. In light of the account of the role of the progressive in characterizing our actions, my thought is that we can now also see how failing to meet the Case Condition can equally undermine the status of an action as intentional.

Now as in the case of the Belief Condition, there is, as flagged earlier, a standard way to understand the role of the Case Condition as a presupposition of intentional action and practical knowledge. As it is typically construed, the idea is that for an agent to act intentionally they must have various beliefs or knowledge about the materials upon which they act and the various objects and features of the environment in which they act. Significantly, though, such cognition tends to be restricted to matters that lie, as it were, in the 'background' to one's action.<sup>28</sup> To take the earlier case again, it will be said that for an agent to intentionally paint a wall yellow they would need to know things like "There's the brush", "There's the paint", and "There's the wall". And, as with the Belief Condition, my complaint here isn't that this is necessarily wrong, but that it is incomplete. If we are concerned with what is

---

<sup>28</sup> See, for instance, Falvey (2000, pp.34-5), Velleman (2007, p.20), and Schwenkler (2015, p.18).



required for having an intention to do such-and-such in such-and-such a way — as in the first part of the passage from Anscombe with which we started — then construing the Case Condition in this way may not leave anything out, especially if the intention is one for the future. But if are concerned with what is required to be doing something in doing something else — as I am — then we will surely want to say that the agent needs to have cognition not only of what lies in the background to her action, but also about what happens in the foreground, about what is happening, to use Anscombe’s phrase again, “the whole time” that she acts.

An example will help. Consider landing a plane by autopilot. Once the pilot has programmed the autopilot to ‘autoland’ (as, I’m told, is the jargon for this process), we can suppose that the pilot must still *monitor* or *track* the plane’s performance — imagine there are a range of warning lights or signals that may trigger if anything is not going according to plan, and that they often enough trigger so that the pilot has to manually adjust accordingly. We can then distinguish two ways in which the pilot can be said to be properly in control of the plane’s landing: he can set off the autoland process to land, and he can monitor its doing so. Accordingly, we can say there are two ways in which he can fail to be in control of its landing: he may not set the process off at all, or he may not monitor it properly. If he doesn’t start the autoland when he ought to, but some miraculous combination of air current and declining fuel supply causes the plane to safely land, the pilot will not have landed the plane intentionally. But if he does start the autoland, and then he puts up his feet and starts reading a book, all the while paying no attention to the plane’s controls at all, we will equally want to say that he hasn’t landed the plane intentionally, as the outcome may well be too accidental.<sup>29</sup> For the pilot to land the plane intentionally, he thus needs to know more about “what is the case” than what he might need to know at the point at which he sets the landing in motion. Similarly, to be intentionally painting the wall yellow, as well as knowing where the

---

<sup>29</sup> For an analysis of luck in terms of a lack of control which draws such a distinction between ‘effective’ and ‘tracking’ notions of control, see Broncano-Berrocal (2015).

brush is and so on, I'll need to know things like "The brush is in my hand", "The paint is getting on that part of the wall", and "I've covered the middle of the wall with enough paint".

The idea generalizes. Intentional actions take time and involve various steps or parts or stages — they are, as we've seen, characterized by progressive aspect, and unfold over time. Accordingly, what the Case Condition requires the agent to know or believe will in effect *update* as things progress. To do *Z* in doing *ABC* we need cognition about what is the case and how doing *ABC* is a way to do *Z*. In the thick of things, though, there is no single bit of cognition at work here, but rather something that takes on a variety of shapes as the circumstances demand. To illustrate with an example from Anscombe (2015, p.154): "If one knows that one cannot  $\phi$  by doing *A* unless *p* (mount or descend a step by treading a certain way unless there is a step there), then one cannot do *A* in order to  $\phi$  unless one *believes p*." But we don't just walk down a step; we walk down the stairs, and to do that we don't just need to believe that there is a step *there*, but also one *there*, and *there*, and so on; and we will need to be aided by cognition about whatever else comes up along the way, such as where the railings are, or that there is someone coming down the stairs directly in one's path, and so on. And this should bring into relief Thompson's remarks about the role of *checking* on one's progress. The possibility of checking to see what one has done or to see what is happening with the material upon which one is working is facilitated by the broadness of the progressive; but checking can be part and parcel of meeting the Case Condition, as checking on one's progress can also provide the sort of information the agent needs to continue to be acting intentionally. Considerations about the role of the progressive are not then a digression from our starting point, but a way of properly grasping the role of the Case Condition in intentional action and practical knowledge.

To be sure, there must be limits to what the Case Condition requires in any case. In particular, it can't be that the Case Condition requires the agent to have cognition of facts which would undermine the possibility that satisfying the condition (and the Belief Condition) can be a way to garner knowledge that serves a merely enabling rather than justificatory role in the agent's practical knowledge. It can't be,

then, that satisfying the Case Condition can require having speculative knowledge that one is doing  $A$  where one's practical knowledge has the same content. But to rule such cognition out is still to rule a lot in, a lot more than has sometimes been thought. How much though? What exactly is required to meet the Case Condition? And, while we're on the question, exactly what is needed to meet the Belief Condition in any particular case? (That question was left unaddressed earlier on.)

If we recall the passage by Anscombe cited earlier, she offered no sharp answer to these questions. She doesn't tell us how we know just what either condition would require in any particular case, nor even when they would call for knowledge rather than mere belief. But this doesn't seem like a shortcoming, for asking for sharp criteria here seems like a bad question, or at least a very intractable one. Indeed, if our conditions are to reasonably track our judgements of when an agent acts intentionally, we surely ought not to expect that very sharp criteria could be laid down. The story of how these conditions apply to particular cases may be a little messy, involving a range of contextual factors, leaving room for some disagreement about individual cases.

But that does not mean that satisfying our conditions (or failing to do so) will not mark out broad patterns or types of cases. For instance, we can think of Thompson's 'ordinary' cases as ones where we have belief or knowledge about how to achieve our, end and can help ourselves through observation or other means to whatever further knowledge or belief we may need as we go along, including information about what has happened or is happening to the material we're working on. In other words, the ordinary case is the one where we need speculative knowledge "the whole time", and get it. At another extreme, we would have the sort of case discussed earlier where the agent doesn't possess any relevant speculative knowledge about what they are up to as they act. Even if, as I suggested earlier, Anscombe's own examples of such cases were hard to make out, I wanted to allow for the possibility of such cases, and our condition can accommodate that possibility. Suppose, for instance, that the agent has knowledge that satisfies the Belief Condition, is in otherwise favorable epistemic circumstances that don't change over the course of her action, and is engaged in some minor activity —

the sort of thing that Anscombe (1963a, p.92) calls one of “the small activities of everyday life”. In that case, we may be happy to say that agent has practical knowledge of what they are doing even if there couldn’t, say, see what they were doing right at that moment.

That may even be the lot of someone making copies with carbon paper, if we offered a different kind of re-description of the situation. Suppose Gertrude is a seasoned carbon-copier. She practices writing through stacks of twenty sheets every morning before breakfast. She knows this is her familiar pile of papers and carbon paper in front of her, that this is her trusty pen. The mafia also have a gun to her head, and she only gets one go, with no chance to check — they need to move her on quickly to other stacks of paper to keep going. But she only has to write through ten sheets, and she does so, effortlessly. I am inclined to say: she does so intentionally and has practical knowledge of what she is doing, even if she doesn’t get to check. Thompson (2011, p.210) himself seems to give a more realistic case not entirely unlike like this one, when he says “there are some real estate closings that are like that” in reference to the carbon-copy case as Davidson’s construes it. But surely a realtor closes the deal intentionally and with knowledge of what he is doing. It’s the wider facts about the realtor’s epistemic situation and capacities that explain why we’re happy to say that.

Crucially, though, Davidson’s case, as Davidson describes it, is different. It doesn’t fall into the same patterns as either of the two sorts of cases just sketched. In this case, it matters that the agent doesn’t get to check on what he has done, because he doesn’t have a bit of instrumental knowledge that can take up the slack of not possessing such observational feedback — all he has is a mere belief about how to accomplish the task — and nor is he engaged in some small familiar matter, but a fairly *recherché* task. Moreover, he has some explicit ground to think that he isn’t succeeding but is, for whatever reason, unable to check to see that he in fact is succeeding. In this case, I say, the Case Condition is not properly met, and we can therefore say that the agent’s success is not intentional. And notice that we can say that without saying that in any case where the agent only has a mere belief about how to achieve their end they don’t act intentionally if they have some doubt about their success — as, for instance, when we do

something for the first time. That sort of case can exemplify a different type too, since unlike in Davidson's peculiar setup, nothing about such a case requires the agent not to be able to check on how things have gone, as we routinely do, wanting to find out if we pulled off the task, with the progressive again providing the requisite slack to not deny that the agent knows what she is doing just because it is the first time she has undertaken this particular task.

So where does all this leave us? I've argued that intentional action and practical knowledge jointly presuppose two forms of cognition, and that getting clear about the content of these conditions lets us see why we need not treat carbon-copy cases as *bona fide* cases of intentional action. I haven't offered precise criteria for how these conditions apply to any particular case — I doubt that would be possible — but I've argued that we can nevertheless see the operations of these conditions as roughly delineating certain classes of cases, one of which rules out the carbon-copy cases. But I haven't argued explicitly for the claim that the Belief Condition or the Case Condition are necessary conditions on intentional action. In the case of the Belief Condition, I think more can be said, and I'll return to the topic in the next chapter. But in the absence of a more direct argument for applying the Case Condition in the way I have, the advocates of the carbon-copy cases may be unimpressed. They may well be happy to say that the condition as I've interpreted it is necessary for having practical knowledge of what we are doing, but they will balk at the idea that it must hold as a condition on intentional action. The advocate of a carbon-copy case will simply want to say that "Davidson's man" doesn't need to have speculative knowledge of his progress to be intentionally succeeding. So it may not seem that my account has moved the debate has very far.

But we have made progress, in the following way. For we can now urge, in the first place, that the insistence that *AP* is false in the light of carbon-copy cases rests not on an unchallengeable intuition, but rather upon a substantive bit of theory. That theory is the one flagged at the outset, that intentional action is simply very easy to come by in the sense that our capacity to act intentionally does not depend upon cognition from outside the will. The humbler picture of the will that I'm trying to recommend

offers a different story, one upon which our doings can only properly count as non-accidental, and thus intentional, when we find ourselves in the more favorable epistemic circumstances than the carbon-copy cases suggest. And so, at the same time, setting aside carbon-copy cases as involving unintentional actions on the ground of violating either of our two conditions (as I've construed them) should not seem like the kind of thing that only someone defending a theory would do, but is rather motivated by the thought that doing so is one way to preserve the widely accepted idea that intentional actions cannot be accidentally brought about.

## II. ACTING IN THE LIGHT OF MY END

“[M]y idea of what I’m doing [is] as it were a miner’s lamp on [my] forehead which illuminates always just so far ahead as we go along.”

— J.L. Austin, ‘Three Ways of Spilling Ink’, p.284.

### 1. INTRODUCTION

If there is a necessary connection between intentional action and knowledge, how can we account for it? How do we explain the truth of *AP*? Suppose we start by recalling (again) Anscombe’s definition: an intentional action is “an action to which a certain sense of the question ‘Why?’ is given application”. At a first pass, that definition may seem to offer a very easy explanation of *AP*. Here it is: If an intentional action *is* one where the agent gives the question ‘Why are you doing *A*?’ application, then that explains why the agent must know that she is doing *A*, since her ability to answer the question will express her knowledge that she is doing *A*. On the other hand, if she refuses to give the question application, that will explain why she doesn’t know, because she can’t explain anything of which she lacks awareness. Now one reason why this is too easy, or not right anyway, is that Anscombe surely did not mean to imply that an action is intentional only when an agent *does* give the question application. We need to at least weaken the definition so that it means the agent is *able* to give it application. But even then, as Kieran Setiya points out, the matter is not straightforward:

What is clear is that intentional action is that to which the question “why?” *has* application, understood as a request for reasons. The possibility of explanation by reasons is distinctive of intentional action. It is a further question whether, and why, the question has application only when the agent “gives it application” because she is able to answer it. In particular, why must she have a true belief about what she is doing, in order to be doing it for reasons at all? Why isn’t it sufficient for reasons-explanation that she knows what she *intends* to do, or what she is *trying* to do? More strongly, why must she have any belief about what she is up to, at all? Why is it not enough for her action to be motivated by the right desire? (Setiya 2007, pp.26-7)

In this chapter, I want to show that this question can be answered, and not just in a way that shows the agent needs to believe that they’re doing such-and-such, but also that they must *know* that they are doing so. As in the previous chapter, my focus here will be on the particular case of doing something *in* doing something else, or of ‘doing *B* in doing *A*’, as I’ll often say. When philosophers consider the relationship between action and cognition they tend to take as their topic intention or intentional action. But my thought is that if focus on what is surely the central phenomenon of intentional action — doing one thing in doing another — then we will be able to offer a novel type of explanation of the truth of *AP*.

The basic idea will be to explain why we always know what we’re doing as a result of the interplay of two distinct cognitive factors in intentional action, so construed. First, I will argue in §2 that it is a necessary condition on doing one thing in doing another that the agent possess a certain type of instrumental cognition. In effect, we’ll again be looking at the Belief Condition of the previous chapter. I’ll offer a further clarification of its content so that we can see beliefs that accord with the condition as involving a particular kind of *de re* element. The immediate upshot will be that in the cases on which we’ll be focused, the possibility of the application of the question ‘Why?’ to the subject matter *does* imply that the agent must at least have this sort of belief to be acting intentionally. Second, I will argue that by focusing on cases of doing one thing in doing another, we will also see why the agent must have knowledge of her reasons. I argue for this in two steps: by arguing in §3 that in the cases at hand the agent can always offer a certain type of reasons explanation, and by arguing in §4 that this type of



explanation is only one that the agent can offer if she has knowledge of her reasons. As we'll see in §5, the combined effect of such belief and knowledge is the agent must also know what she is doing. And so *AP* is true.

## 2. THE BELIEF CONDITION

My first task is to show that intentional action necessarily involves a specific type of cognition. More precisely, I will argue that for an agent to intentionally do one thing in doing another it is necessary that they have a specific type of instrumental belief — a kind of *de re* belief about their own activity, which I'll call a belief *de actione*.<sup>30</sup> I will work up to spelling out the content of this belief in a few steps, and I'll argue for it on the ground that an intentional action must be non-accidentally brought about by its agent to count as such.

My starting point then is the basic assumption that an agent can act intentionally only if their action is brought about non-accidentally.<sup>31</sup> To illustrate, consider the following principle:<sup>32</sup>

*The Further Intention Principle:* If *X* is intentionally doing *A* with the intention of doing *B*, and *X* is successfully doing *B*, then *X* is intentionally doing *B*.

---

<sup>30</sup> Thanks to Tom Marré for suggesting the phrase 'belief *de actione*'.

<sup>31</sup> As in the previous chapter, I don't mean to deny that there are sense of 'luck' that can apply to intentional actions, but only that there are also salient ways in which intentional action must not be accidental, one of which I mean to bring out here.

<sup>32</sup> Compare the principles discussed in Hanser (2000), from whom I've borrowed the names of the two principles that follow.

This principle is false. An agent may aim at an outcome and succeed in bringing it off only by accident, thus failing to act intentionally. There are familiar examples. Suppose I'm crossing the street to get to the bus stop, but a freak blast of wind lifts me in the air and drops me at the bus stop. I've ended up where I want to be, but I haven't done so intentionally. Hence, we need to say that to be intentionally succeeding in one's further aim requires that one succeeds non-accidentally. Adding such a condition to the above principle, we get:

*The Non-Accidental Consequence Principle:* If  $X$  is intentionally doing  $A$  with the intention of doing  $B$ , and  $X$  is successfully and non-accidentally doing  $B$ ,  $X$  is intentionally doing  $B$ .

I think this principle true. But what does it mean to here invoke 'non-accidentality'? It surely requires more than that what happens accords with my intention, or with what I want to do, or even that it may accord (in some sense) with what I am trying to do, since all of these conditions can, for example, be satisfied in the example I've just described. The obvious suggestion is to say that what happens must also in some way be in accord with the content of an instrumental belief which is involved in my action. For my further intention to be fulfilled in a non-accidental way, it must be that what happens accords with the content of such a belief. My plan never included being carried along by the wind, so what happens deviates from my practical thought.<sup>33</sup>

---

<sup>33</sup> This role of the instrumental belief is often referred to as a "sensitivity condition" in the literature concerning the causal theory of action. Such a condition is not the only one that causalists have proffered in trying to deal with the problem of deviant causality (indeed, it is often taken to be insufficient for doing so), nor is this the only kind of accidentality that they may want to address — see Setiya (2007, p.32), for instance, for a discussion which invokes such a consideration to account for non-accidentality for 'non-basic' actions, but an alternative story for 'basic' ones. For an extended discussion of such issues from a causalist perspective, see Bishop (1989, Ch.4 & Ch.5). For the 'state of the art', see Aguilar & Buckareff (eds.) (2010), and also Hyman (2014) for an important discussion about the role of dispositions in causal theories, and Stout (2010b) for the claim that 'Aristotelian' approaches to causality can skirt the problem.

What kind of belief must it be? The content of such a belief is given, I will assume, by what would be revealed by the agent's answers to the question 'Why?'. (Alternatively: the content of the belief is given by the practical reasoning at work in the action.) But what kind of content would that reveal? As we saw in the previous chapter, a common answer is that such a belief is a "a run-of-the-mill belief whose propositional content is that, if one kind of thing were to happen, another kind of thing would happen" (Ford 2016, p.143). Is that all we need to say to cross the line from the *Further Intention Principle* to the *Non-Accidental Consequence Principle*? I'll suggest in a moment that it is not. Notice, though, that even if it is necessary for non-accidentally doing *B* in doing *A* to have a belief of this sort, then although that necessity would establish a connection between intentional action and cognition, it wouldn't seem to help us get us closer to establishing *AP*. I don't mean to point to the obvious fact that we would only have belief on the scene. The problem, to adapt a point Setiya (2007, p.33) makes, is that *any* kind of cognition of this general or hypothetical shape is not sufficient for having cognition that I'm doing *A* or that I'm doing *B*. For the content of the instrumental cognition, so construed, can be satisfied without requiring the existence of any particular actions, and so without any knowledge or belief about particular actions.

Must we in any case construe the content in this familiar way? There are undoubtedly reasons why the content of an instrumental belief has been so often described in this way, not least being that we often form an intention to do such-and-such, in which case it would be natural, if doing such-and-such required doing so-and-so, that we form such a belief in forming such an intention.<sup>34</sup> I suggested in the

---

<sup>34</sup> The element of generality might also be taken to capture the sense in which such a belief can codify a background ability or capacity, something that essentially involves generality. Or perhaps it is the thought that an instrumental belief can't properly track the precise causal path involved in executing an intention and therefore requires some greater generality (see Mele & Moser 1991, pp.240-1; also Stout 2010b, p.159). Or perhaps it is a symptom of the dictum accepted by some causal theorists of action that a cause must precede its effect, which would rule out mental states playing a causal role if they entailed the existence of particular occurrences, and so the instrumental belief could not be allowed to entail the existence of a particular belief. For an example of this last thought, see Stout's (2010b, p.160) discussion of why a causalist of a 'non-standard' type who wishes to employ intentions for the future to do causal work

previous chapter that doing so can blind us to various other features of instrumental cognition, and here I want to add one such further feature. To begin, we should note that the content of the belief is not always so described as to have a purely ‘general’ character. Thus, in ‘Actions, Reasons, and Causes’, Davidson suggested that when a belief-desire pair (or ‘primary reason’, in his jargon) rationalizes an action, the agent is characterized as having a ‘pro-attitude’ toward actions of a certain kind and “believing (or knowing, perceiving, noticing, remembering) that *his action* is of that kind” (1963, p.3-4; my emphasis).<sup>35</sup> And Michael Smith (2012, p.387) says that when an agent acts intentionally he must have “some belief ... that a basic action of his — specifically, his moving his body in the way under discussion — has a suitable chance of making things the way he desires them to be”. These remarks seem to describe not a ‘general’ belief with a content which doesn’t entail anything about particular actions but rather a ‘singular’ belief that can be taken to refer *de re* to a particular action. And so we seem to have two candidates for the shape of an instrumental belief entailed by an agent’s doing *B* in doing *A*: a general belief about how performing one kind of action could lead to a desired outcome, and a singular belief concerning how a *particular* action will bring about that outcome.

If an instrumental belief is necessary for a given action to be non-accidental, of what kind is it then? The argument I gave a few paragraphs back might have been taken to show that a general belief is necessary, but perhaps it shows that a singular one is instead, or as well. I think a singular belief is indeed necessary for doing *B* in doing *A* to satisfy the *Non-Accidental Consequence Principle*. That is, where ‘*X* is

---

cannot resort to intentions to do *A* where ‘do *A*’ is meant to pick out a particular action, since such states can have no object, as there is yet no such particular.

<sup>35</sup> Though, according to Hunter (2015), this specific belief seems to fall out of the picture in Davidson’s later papers, replaced by a general belief about performing an action of a certain type. Hunter also suggests that Davidson was confused in allowing for such a belief, on the ground just flagged, that Davidson also argued that a rationalizing cause (or at least the desire or ‘pro-attitude’ part of such a cause) must be general in the sense of not entailing the existence of the particular event that it brings about.

doing  $B$  in doing  $A$  describes a happening that satisfies the *Non-Accidental Consequence Principle*, I think a principle along the following lines must be true:

*The Belief De Re Principle*: If  $X$  is doing  $B$  in doing  $A$ , then  $X$  believes, of her doing  $A$ , that it is a way of doing  $B$ .

I say “along the following lines” because this particular principle is false. To see why, we need to consider another important feature about descriptions of events which involve progressive or imperfective aspect, also brought out by Thompson (2008, pp.134-7). If, for instance, I am crossing the road, it does not follow that I will cross the road. It may be that I will get hit by a bus along the way, or just change my mind half way across and turn around. Either way, it will still have been true that I was crossing the road, and while I was doing so I could have expressed that truth by saying “I am crossing the road”. But that means that for such a claim with progressive aspect to be true, reality need not contain a *particular* — an action or an event — of the sort that we could refer to using a singular term like the description ‘my crossing the street’.<sup>36</sup> In so far as our descriptions of an agent’s doing  $B$  in doing  $A$  involve progressive aspect, the *Belief De Re Principle* can be false, at least if we take the expression ‘of her doing  $A$ ’ to purport to have reference to a particular. For there need not be (or there need not occur) such a particular for the claim “I’m doing  $B$  in doing  $A$ ” to be true, even if I will in the end have done  $A$ . And so it can’t be necessary for anything that we have such a *de re* belief.<sup>37</sup>

---

<sup>36</sup> Though the truth of such a statement may well and usually will entail that certain particulars have come into existence which we could refer to, like the one we would pick out with ‘my moving my foot from the curb onto the street’. So there could be some kind of modified version of the *Belief De Re Principle* that would capture this; but I won’t pursue it. I mean ‘particular’ here to refer what the agent could (attempt to) refer to, in that context, with the nominal ‘my crossing the street’.

<sup>37</sup> See McDowell (2011b) for a related argument concerning the claim that intention in action requires certain *de re* attitudes towards actions as Wilson (1989) proposes. I’m indebted to McDowell’s paper in the next paragraph too.

But that doesn't refute the thought that we need *some* such belief to satisfy the *Non-Accidental Consequence Principle*. Rather, it calls for modification of the *Belief De Re Principle* so that it refers to the agent's *activity*, to what they are up to *in* doing *A*. That is, we need a principle that refers *de re* not to my crossing the street but to the *activity* in which I'm engaged *in* crossing the street. If I'm going to the bus stop in crossing the street, I still believe, of my activity of crossing, that *that* is a way to get to the bus stop. So we need a form of expression to replace 'of her doing *A*' in the above principle to instead pick out the activity that is involved, at the point in time at which we refer to it, in the agent's action of doing *A*. I suggest the following:

*The Belief De Actione Principle*: If *X* is doing *B* in doing *A*, then *X* believes, of her activity of doing *A*, that it is a way of doing *B*.

Does this principle follow from the *Non-Accidental Consequence Principle*? I suggest it does, and the argument has the same shape as earlier: without such a belief, the agent's doing of *B* may be merely accidental, and so not intentional. For even if having a general belief connecting two kinds of actions is necessary for acting non-accidentally, it doesn't itself seem sufficient. If I intentionally get to the bus stop in crossing the road, then it would not be enough to say that I have a general belief that crossing the road is a way to do that. If that belief isn't at work in my action in a more specific way, then my getting to the bus stop won't be intentional.

We might try to avoid saying that a singular belief is also necessary by recalling what other states might be at work in our agency when we do one thing in doing another. Suppose to the general belief that crossing the road is a way to get to the bus stop we add that I also believe that I am crossing the road. Is that enough to make getting to the bus stop intentional? No: we can still imagine a case where we are in those states but don't connect their content — I don't believe that *this* activity of crossing the

road is an instance of doing something that will leave me at the bus stop — and so where I don't intentionally arrive at the bus stop. The case as we originally described it would do. The same is true if we recall that I am acting with an intention, so that I also have the intention to get to the bus stop. We can make explicit that such an intention is at work and still repeat our scenario where the wind gets me there, because there remains no connection between the content of the general belief, the belief that I am crossing the road, and the intention to get to the bus stop.

We want to say, I think, that an instrumental belief must be operative in our intentionally doing *B* in doing *A*. And I think that requires the agent to think *of* her activity of doing *A* that *it* is a way to bring about her goal of doing *B* (or more precisely, because she *is* doing *A*, that her doing *A* *is* bringing about her doing *B*). Once I can say that I believe, of my activity of crossing the street, that *it* is a way to get to the bus stop, then it is possible that I've intentionally arrived at the bus stop. It may be that we need to say more — I'm not here trying to devise a set of necessary and sufficient conditions for non-accidental action — but I think we can see that the *Belief De Actione Principle* follows from the *Non-Accidental Consequence Principle*. If it didn't, then we shouldn't have any trouble with a train of thought like this: "I was intentionally crossing the street. I was crossing the street with the intention of getting to the bus stop. As it happens, I intentionally arrived at the bus stop. But even though I believed that crossing the street could be a way of getting to the bus stop, I had no view about whether my activity of crossing the street was how I got to the bus stop." And of course we do have trouble with that — it seems flatly incoherent. The *Belief De Actione Principle* explains why.

But if this is much is right, then the *Belief De Actione Principle* does answer the question broached at the outset about whether the applicability of the question 'Why?' should entail that the agent has some kind of cognition of what she is doing. So if we focus on intentionally doing one thing in doing another rather than intentional action *sans phrase*, we can see why a particular type of (true) belief that one is doing such-and-such *is* necessary to be intentionally doing such-and-such. The intention or desire to do *B* or even an attempt to do *B* (construed in a way that doesn't involve an instrumental thought about

how one is trying to do  $B$ ) may all be present in an agent who is bringing about  $B$  by doing  $A$ , as may a general belief that doing  $A$  would be a way to do  $B$ , but if the agent lacks a singular belief that fits the *Belief De Actione Principle*, then we won't be able to say that the agent is doing  $B$  intentionally. So we're a step closer to establishing  $AP$ .

### 3. THE AVAILABILITY OF NAÏVE EXPLANATION

But we're still some way off. What the *Belief De Actione Principle* does not explain is why the agent needs to *know* that she is doing  $B$  in doing  $A$  to be intentionally doing so. To get to that result, we will need to first consider what sorts of reasons explanations will be true when an agent is doing one thing in doing another, where by reasons explanations we mean explanations that purport to explain what the agent is doing by invoking the agent's reasons for acting — that is, in Jonathan Dancy's (2000, p.1) expression, “considerations in the light of which” the agent acts. I will argue that in the case of intentional actions in which we are interested, what Thompson (2008, Ch.5) calls a “naïve” explanation is always available. The significance of this will not be apparent immediately, but will emerge in §5.

In ordinary language, we explain actions in a wide variety of ways. Restricting our attention to reasons explanations (that is, setting aside, say, purely causal explanations), we can simplify the discussion in a familiar way by focusing on certain canonical forms or expressions of explanation. Following Setiya's (2009) typology, I will say that reasons explanations come in three main forms: (1) ‘ $X$  is doing  $A$  in order to do  $B$ ’, (2) ‘ $X$  is doing  $A$  because  $p$ ’, and (3) ‘ $X$  is doing  $A$  on the ground that  $p$ ’.<sup>38</sup> Of course, such

---

<sup>38</sup> Of course, it is a simplifying assumption that anything that could count as a reasons explanation in ordinary language could be bashed into one of these syntactic shapes. (I'll mentioned a few others, or perhaps variants on (3), a little later.) For further discussion of the forms and features of rational explanations, see Audi (1986), Dancy (2000), Davis (2005), Thompson (2008).



sentence forms can also be used to offer explanations which do not invoke reasons in the relevant sense, as when we offer purely causal explanations such as “Nina is coughing because she has a cold”, but we’re stipulating away such cases, or any others that don’t involve giving an agent’s reasons.

The kinds of explanations I want to focus on are a subcategory of (2). Following Thompson (2008, Ch.5), we can distinguish (roughly) between “naïve” and “sophisticated” forms of reasons explanation, where a naïve explanation explains what the agent is doing in terms of something else they are doing, while a sophisticated explanation instead works by invoking what the agent wants or intends or is trying to do. So: “I’m putting the pasta in the water because I’m cooking dinner” is naïve; but “I’m putting the pasta in the water because I want to cook dinner” is sophisticated. For Thompson, the interest in the notion of naïve explanations is whether such explanations are in some way prior to or more fundamental than their sophisticated counterparts.<sup>39</sup> For our purposes, we don’t need to take up the question. The issue for us is only whether it is true that when an agent is doing *B* in doing *A* there is a naïve explanation available of their action. Of course, we’re not asking here about any event that could be described as the agent’s “doing *B* in doing *A*” — this expression, as I have been using it, is *defined* by the *Non-Accidental Consequence Principle*, so that it applies to events where the agent is intentionally doing *A* with the intention of doing *B* and is non-accidentally succeeding in doing *B*. Our question, then, is whether the following principle is true:

*The Naïve Explanation Principle.* If *X* is doing *B* in doing *A*, then *X* is doing *A* because *X* is doing *B*.

It seems to me it is. The ‘because’ of reasons explanation is factive, and in the cases we’re imagining it is true that the agent is doing *A* and doing *B*. The agent is doing *A* with the intention of doing *B*, and

---

<sup>39</sup> See Thompson (2008, pp.117-8), Bishop (2011, p.216), and Wiland (2012, p.151). See also Dancy (1995, pp.426-8; 2000; 2008), though he is not concerned with naïve explanation.

they're succeeding non-accidentally. We also know by the *Belief De Actione Principle* that they correctly believe of their activity of doing *A* that it is a way to do *B*. What more could we need for naïve rationalization? The advocate of the carbon-copier will of course say we need a lot more; but we've set those cases aside in the previous chapter. However, because of the particular course which my argument is going to take, it is also worth noting that in typical carbon-copy cases the agent will still be able to offer a naïve rationalization of what they are doing. For example, "Davidson's man" may not believe that he is making an impression on every sheet, but he does believe that in holding the pen and moving it a certain way he is writing on the top sheet. This may not be true for every such case, but in so far as they carbon-copy cases involve doing one thing in doing another we should be able to find some *other* thing along the way that can naively rationalize what the agent is doing.<sup>40</sup>

#### 4. ACTING IN THE LIGHT OF A FACT

Having shown that there is a particular kind of reasons explanation available when one is doing one thing in doing another, I want to now turn to consider the *nature* of the reason involved. I'll start with some general questions about the nature of reasons for action. Just what are reasons? Earlier, I used Dancy's expression that reasons are considerations in the light of which we act. But what are such considerations? On one view, reasons are the mental states of the agent operative in their action. On another view, reasons are propositions, or whatever it is that we take to be designated by the that-clauses that we use to express the contents of our attitudes. And on yet another view, reasons are facts, or true propositions, or whatever facts are if they aren't true propositions. But I'll argue that reasons for actions

---

<sup>40</sup> Of course, the rationalization could involve a very inarticulate form of expression, and certainly wouldn't require action-descriptions that are common or familiar in ordinary language. It may just be something like "I'm moving my arm like this because I'm moving this pen like that".

can cut across these three ontological categories. My main target, though, will be to put on the table the more specific claim that when a fact is the consideration in the light of which we act, we must not only believe the fact but also *know* it.

Let us start then with a few remarks about how the different kinds of reasons show up in ordinary thought and talk. Can a mental state which I'm in be the reason for which I act? To focus on the case of belief, can the fact that I believe that *p*, or my believing that *p*, be the reason for which I act?<sup>41</sup> The answer seems to be clearly "Yes". The contexts involved seem special though. To give an example from John Hyman, via Dancy (2000, p.125, *n.3*), if I think there are pink rats in my boots, then the fact that I believe that (or my believing that) may serve as my reason for action if I (say) call the doctor or a psychiatrist in response to it. The example might not be the best, for if I did go on to seek help in that situation, that might be said to show that I don't really believe it after all. But consider a familiar sort of publisher's blurb: "If you believe that *p*, then this is the book for you!". We could imagine a strange circumstance where someone goes around purchasing books only if the blurbs promise to speak to the person's beliefs or interests in this way. This need not be irrational or unintelligible: imagine the person is being paid by the hour by some book hoarder to buy books of any kind with no time constraints, and they've decided to go with this criterion.<sup>42</sup>

So our mental states can at least *sometimes* serve as the reasons for which we act. But these cases, as I've said, seem special. The more typical sort of case is one where our reasons are given by the *contents* of our attitudes, not the attitudes themselves, or our having of the attitudes. If it is true that there are pink rats in my boots — maybe someone has painted a few rats as a practical joke — then if I set the cat onto them, or release the hounds, my reason is *that there are pink rats in my boots*, not that I believe that

---

<sup>41</sup> This is of course a simplifying assumption. Maybe advocates of the view that all reasons are psychological states — 'psychologism', as its sometimes called — will want to help themselves to more than beliefs in a full story about the motivation of action.

<sup>42</sup> For more cases of this sort, see Hyman (1999, p.444), Dancy (2000, p.124), and Hornsby (2007, p.288; 2008, p.247).

there are pink rats in my boots, or my believing as much. Or, to adapt my other example, the consideration under which we buy a book is usually *that the book is such-and-such* — interesting, cheap, or whatever, and not *that I believe that the book is such-and-such* — even though in both cases it may be true that I must still have that relevant belief to be able to act for the reason given by its content.

What must the content of such a belief consist in? The cases I've just described mention *true* beliefs, or facts, as I will say; but must the contents be true? When we act for the reason that *p*, it will often enough be that *p* is the case, so that we can say that our reason is *the fact that p*. And it seems to be true that when we explain our actions by giving our reasons, we take our reasons to be facts, or represent ourselves as taking them to be facts, at least in the case when we state our reasons in the present tense. That much is true even in the first set of cases above involving pink rats and buying books: it is the fact that I believe such-and-such (or the fact of my believing such-and-such) which serves as my reason.

However, it isn't universally accepted that reasons for action must be facts or true propositions. Dancy (2000; 2006; 2008; 2011) argues that there are forms of reasons explanation which are not 'factive' in that their truth does not entail that the subordinate clause which provides the explanans is true (all reasons explanations are factive in the place of their explanandum).<sup>43</sup> He gives as examples forms such as "X's reason for doing *A* was that *p*" or "X's reason for doing *A* is that *p*" or "The reason for which *X* is doing *A* is that *p*", and argues that their non-factivity is borne out by our being able to cancel the putative factual commitment without contradiction, as when we say "His reason for doing it was that it would increase his pension, but in fact he was quite wrong about that" (Dancy 2000, p.132). In effect, we might view these, and especially the last case, as close to our case (3) mentioned earlier — '*X* is doing *A* on the ground that *p*' — since this latter form also seem non-factive. If there are such non-factive forms,

---

<sup>43</sup> Obviously, "*X* is doing *A* in order to do *B*" is not factive in this sense, but then it is also not non-factive, since it's explanans isn't truth-apt.

then, perhaps a little surprisingly, reasons explanations in general need not commit us to the truth of their explanans, for the forms where that question can arise.<sup>44</sup>

Not everyone finds Dancy's examples to be convincing cases of non-factive explanation.<sup>45</sup> I will simply assume that Dancy is right that such forms exist, and that they can still do explanatory work by displaying to us the way in which the agent acts in the light of a false belief or a putative fact. What I want to question, though, is the extent of such non-factive explanations. I want to allow that other forms might be non-factive too. Thus, I will also follow Dancy (2000, p.133) in allowing for the possibility of what he calls an "appositional" form of the 'because' locution which is non-factive, as illustrated by a sentence like "He did this because, as he supposed, she had lied to him". (The thought here is that such a use of an expression like 'as he supposed' or 'as he believed' creates an intentional context, but without effectively preserving factivity, as would be the case if we appended instead 'he supposed that' or 'he believed that'.)<sup>46</sup> This isn't to challenge the factivity of 'because' in its more familiar appearance in reasons explanations — Dancy (2011, p.350) is clear that he doesn't mean to do that — but rather to introduce, or make explicit, the possibility that we can see a non-factive connective of the form 'because-

---

<sup>44</sup> As Dancy (2000, p.134) points out, the possibility of non-factive reasons explanation may go some way to defusing a familiar argument in favor of the idea that reasons must be psychological states. That view may seem compelling if we demand that explanation be factive, because in the cases where an agent can act on a false belief it seems like we can always still invoke her believing the falsehood as her reason. Dropping the demand for factivity can relieve that pressure.

<sup>45</sup> Although he isn't responding to Dancy, it is worth noting that Unger (1975, p.208) cites sentences of pretty well the same form as Dancy to illustrate what it would be *inconsistent* to say.

<sup>46</sup> See Turri (2009, p.502) for bafflement at how the thought that an agent can act on a false reason in the sense of how the reason appears to the agent, or how she understands the reason, could not but collapse back into the view that treats all reasons as psychological states.

as-*X*-supposed' at work in the logical structure of some of our claims.<sup>47</sup> But now consider this remark by John Bishop about uses of 'because' in naïve rationalizations:

I may explain my doing *A* by my doing *B* and believe that I am doing *B* when in fact I am not. But, if the explanans in this rationalization were indeed my action of doing *B*, then there would be no explanation of my doing *A* since the proposed explanans does not actually obtain. Yet we recognize that 'I am doing *A* because I am doing *B*' can be explanatory even if it is false that I am doing *B*. And that can only be because the explanatory force of 'I am doing *B*' does not rest on that claim being a true description of what is actually happening but, rather, on its expressing what I want or intend to happen. (Bishop 2011, p.216)

Bishop doesn't explicitly say that he thinks naïve rationalizations can be *true* even if the explanans is false, and if that is what he means, I'm not sure I can do more in response than give an incredulous stare. A non-factive line about 'because' is, I think, hard to swallow because in its standard use '*P* because *Q*' just means '*Q* explains why *P*', and *explaining why* is a factive relation, because explanation just is a relation between facts or truths.<sup>48</sup> So I will draw a line in the sand at this point and insist, against Bishop, that at least the ordinary 'because' is factive on both sides.<sup>49</sup> On the other hand, we could in fact meet Bishop halfway here if we help ourselves to Dancy's 'because-as-*X*-supposed' form, and take Bishop to be using 'because' in *that* sense.

So we've seen that there is ground to say that facts can serve as reasons for action, including facts about psychological states, and that an agent's take on a fact might serve as her reason for action. In

---

<sup>47</sup> See Hornsby (2007, p.292) and Hyman (2011, pp.358-60) for less concessive responses.

<sup>48</sup> See Hyman (2011, p.359) for a similar line.

<sup>49</sup> There is also the expression 'because of', which takes a singular term of the right-hand side, but I'll ignore this case here. For some discussion of this connective, see Steward (1997, Ch.5). It also comes up for mention in Anscombe (1983, p.91).

both cases, we've assumed, the agent believes her reason. The final possibility to consider is an explanatory context where we want to say that the agent must know her reason too. To see what such a context involves, consider the following description of an example from Jennifer Hornsby (2007, p.296; 2008, p.251):

As a result of being told so by his normally reliable friend, Edmond thinks that the ice in the middle of the pond is dangerously thin. Sensible fellow he is, he keeps to the edge. As it happens, the ice in the middle of pond *is* dangerously thin. So Edmond has a true belief. But his friend actually has no idea at all about the thickness of the ice, and he just wants to make sure that Edmond doesn't skate in the middle (for whatever reason). We can't then record that Edmond takes the fact that the ice is thin to be his reason for acting by saying "Edmond kept to the edge of the pond *because* the ice was thin in the middle".

The structure of this example will, of course, remind us of Gettier's (1963) famous counterexamples to the analysis of knowledge as justified true belief. In Hornsby's example, Edmond has a justified true belief, but there is a disconnect between what makes his belief true and the fact itself, and so it is sheer luck that he happens to be right. What is novel and important about this sort of example, is, in the first place, that it shows that having a true belief is not sufficient for an agent to act *in the light of a fact*. As Hornsby (2008, p.256) puts it, because Edmond only has a true belief that the ice is thin that isn't enough to let us see how that fact "registered" with him, or how his action can be seen as "responding to this fact" (2008, pp.251, 256; cf. Hyman 2010, p.411). Rather, for the fact to properly *weigh* with the agent, to play a certain role in his practical rationality, the agent needs to also have *knowledge* of it. Second, what the example shows is that we haven't told the full story about the content of the 'because' of reasons explanation if we only insist on its factivity: the fact it takes in as its explanans must be known by the agent as well. That is why we can't deploy it in this context.

Now, as McDowell (2013, pp.18-20) argues, if we accept that an agent must have knowledge of a fact to act in the light of it, that does not itself mean, as Dancy (2008, p.276) thinks, that we need to

affirm that other non-factive forms of explanation must be ruled out of hand. In Hornsby's example, Edmund can't act in the light of the fact that the ice is thin; but we can still truly say, for instance, that Edmund's reason for skating at the edge is that the ice is thin, as Dancy wants to say. As Hornsby points out, the point of her cases is not to try to "cover the whole ground [of reasons explanation] with a single sweep", but to urge that if we are to properly capture all the different *kinds* of intelligibility involved in acting in the light of considerations, then "we need the idea that knowledge is a spring of action" (2008, pp.248, 252). So I don't claim that cases like Edmund's establish the claim that if an agent is acting for the reason that *p*, they must know that *p*. What such cases establish is the following narrower principle:

*The Action-Knowledge Principle:* If *X* is doing *A* because *p*, then *X* knows that *p*.

Offering one mere example may not seem a very strong argument in support of such a substantive claim. But the examples can be readily multiplied, including to cases involving reasons for belief, demonstrating a strong and stable intuition that we can't help ourselves to the 'because' of reasons explanation if the agent lacks knowledge.<sup>50</sup> Assuming that we now have the *Action-Knowledge Principle* on the table, we can pull the pieces together and see why *AP* is true.

## 5. WHY WE ALWAYS KNOW WHAT WE'RE DOING

In the discussion of what is required for an agent to be acting in the light of a fact that we find in Hornsby, Hyman, and others, the fact that the agent is taken to act (or not act) in the light of is always some fact relevant to the matter at hand — that the ice is thin, that we're out of milk, that the shops are

---

<sup>50</sup> For similar cases concerning reasons for belief, see Unger (1975, pp.209-10), Hyman (1999, pp.447-8), Williamson (2000, pp.60-4), and Marcus (2012, pp.37-8).



open, etc. — but which does not concern what the agent is herself doing.<sup>51</sup> But, as the possibility of naïve explanations brings into sharp relief, we can and routinely do give as our reason for acting the further thing which we are doing in acting. If we focus on such cases, then it seems we will have to say that the agent must know what she is further doing if her doing that further thing *is* her reason for acting. Indeed, this falls directly out of the principles which we've so far adopted. Recall the *Naïve Explanation Principle* from §3:

*The Naïve Explanation Principle.* If  $X$  is doing  $B$  in doing  $A$ , then  $X$  is doing  $A$  because  $X$  is doing  $B$ .

If we now conjoin that principle to instances of the *Action-Knowledge Principle* where the variable in ‘ $X$  is doing  $A$  because  $p$ ’ is given by ‘ $X$  is doing  $B$ ’, then we can immediately conclude that whenever  $X$  is doing  $B$  in doing  $A$ ,  $X$  knows that she is doing  $B$ . And so now we're very close to *AP*: we've now connected action not just to belief, but to knowledge as well.

Should we say that the *Action-Knowledge Principle* encompasses naïve explanation? Or should we say that all that examples like Hornsby's show is that a principle along those lines is true but only where the variable or schema ‘ $p$ ’ ranges over facts which precisely *don't* concern what the agent is doing? I think that the *Action-Knowledge Principle* does include naïve rationalization. But I think that there are some issues that arise once we allow that the relevant facts include what the agent is doing, and it will be worth teasing them out a little. The first issue concerns the appropriateness of the talk of ‘acting in the light of a fact’. Although I didn't try to elucidate precisely how Hornsby and others have employed that expression, we saw that the point of it was to mark how knowledge of the relevant fact makes it possible to say that the fact rationally *impresses* itself on the agent, and that the agent in turn can rationally *respond*

---

<sup>51</sup> The only exception is Thompson (2013), to whom my discussion is much indebted, though I develop the argument for connecting naïve explanation and knowledge in a different way to Thompson.

to it.<sup>52</sup> When the relevant fact concerns what the agent is herself intentionally bringing about, such ‘response’ talk, at least at a first pass, can sound a bit off key, as though my own actions are somehow floating around out there waiting for me to rationally respond to them. But, as I say, I think this is largely a terminological point. Hyman (2010, p.410) says that when we adopt this metaphorical talk of ‘acting in the light of a fact’ or ‘being guided by a fact’ we don’t mean anything more unfamiliar than “that we took it into consideration, that it informed our reasoning, when we decided what to think or what to do”. Given a few minor changes to the tense and aspect of this sort of gloss, such a characterization seems to apply just as well to what we aim at in what we’re doing, so that the way in which the fact of what we’re doing ‘weighs’ with us can be as much a part of our practical rationality at work as is our knowledge of facts independent of our doings. But if it helps, we could instead talk of acting *in the light of an end*.

On the other hand, if we try to *argue* for the necessity of knowing one’s end to be able to act in the light of it in the same manner as Hornsby, then there a few differences that emerge, precisely because the relevant fact is not independent of the agent. Consider for instance the following case, designed to be analogous to Hornsby’s, though now with naïve explanation in play:

Errol is hanging a large painting on the wall so that it is level with the floor, but he knows he can’t trust his own sight because of parallax error. So he asks Adrian, who is on the other side of the room, to guide him. Adrian tells him to move the right side of the painting up a little higher, and as Errol does so Adrian yells out for him to stop because it is level. In fact, Adrian wasn’t paying any attention at all to what Errol was doing, but just by sheer luck he happens to be right, and the painting is level. Still, Errol could not have explained his action at the time by saying “I’m moving the right side up *because* I’m hanging the painting level”.

Now, in some ways, this is just like Hornsby’s case of Edmund: although it is true that hanging the painting level is something that Errol is doing, he doesn’t have knowledge of that fact, and so we can’t

---

<sup>52</sup> For such language, see, again, Hornsby (2008, pp.251, 256), Dancy (2011, p.346), Hyman (2010, p.411), and McDowell (2013, p.19).

say that he acts in the light of it, and so we can't invoke the naïve 'because' to explain what he is up to. But there are some differences too. In Hornsby's example, the 'good' case — the one where the agent can act in the light of a fact — is one where the agent knows the fact that  $p$ ; the 'bad' case — where they can't so act — is either one where the agent has a false belief that  $p$  or an accidentally true belief that  $p$ . In our example though, we can't break things up like this. The good case will be the same: the agent knows that she is doing  $B$ . But what the agent knows is that she is doing  $B$  *intentionally*. In the bad case we can include that she falsely believes that she is doing  $B$ ; but we can't also include the possibility that she has an accidentally true belief that she is doing  $B$  *intentionally* — or at least *I* can't, since that would be to give up on  $AP$ . So we have to describe the cases differently. But we can do this. We can say: in the bad cases, it is either false that the agent is doing  $B$  or that she is 'doing  $B$ ' in the sense of bringing it about, but where it is clear that this doesn't entail doing so intentionally.

There is a further difference even in how things stand in the good case. In Hornsby's example, Edmund gains knowledge of the fact that the ice is thin by testimony. In my case, although the end in the light of which Errol acts is that he is hanging the painting straight, we can't say that he gains knowledge of this fact by testimony — or, again, at least *I* can't if I want to preserve the distinctive character of practical knowledge as Anscombe presents it, because letting in such a possibility would make Errol's knowledge a form of *speculative* knowledge. That means we have to see the example as working a little differently. But we can do this too. If we recall from the previous chapter the role that speculative knowledge could play as an enabling condition on practical knowledge, we have to imagine the case here as one where Adrian offers Errol speculative knowledge about what is happening to the materials he is working on that can enable his knowing that he is hanging the painting straight, rather than knowledge of that fact itself. There are then differences between how we have to understand the cases working once we take facts about what the agent is doing into consideration, but I don't think they undermine the basic claim that in examples like that involving Errol and Adrian, we still can't say that the agent is doing  $A$  because he is doing  $B$  when the agent lacks knowledge of her reason.

Having brought out some of the nuances of how the *Action-Knowledge Principle* applies to the case of naïve explanation, let us consider what follows from it in this context. Consider the following four potential bits of cognition:

- (1). *X* believes, of her activity of doing *A*, that it is a way of doing *B*.
- (2). *X* knows that she is doing *B*.
- (3). *X* knows that she is doing *A*.
- (4). *X* knows that she is doing *B* in doing *A*.

And now suppose that *X* is doing *B* in doing *A*. In such a case, we know from the *Belief De Actione Principle* that (1) will be true. And in the same case the *Naïve Explanation Principle* follows, and so, as just mentioned, since the *Action-Knowledge Principle* is independently true, we can then also derive (2). But now if (1) and (2) are true, then, I think, (3) must follow. And if (1), (2), and (3) are true, then (4) must follow as well. And at that point, we've established *AP*, at least in the case where the agent is doing one thing in doing another.

This may seem too quick. The step from the first three claims to the last seems unquestionable to me, since if (1) must be true, and both (2) and (3) entail belief (which they do), and we're not supposing that for an agent to be doing one thing in doing another requires more than a true instrumental belief (of a certain sort) linking them, then (4) seems available. But what about the prior step from (1) and (2) to (3)? Might that inference not go through? Here is one example to try to bring out a possible problem. Suppose I know that I'm getting on my partner's nerves. I also believe, of my activity of tapping out a repetitive tune, that that is a way to get on her nerves. But one might worry that it doesn't follow from just those two claims that I must know that I'm tapping out a tune, because I might also think that I'm annoying her some other way — maybe it's because I'm humming the tune over and over. But what difference would this make? It is true I may think I'm getting on my partner's nerves in various ways. Since I know that I'm doing so, I believe that I am. So whatever relevant *de re*

beliefs I have will then have the shape that I believe, of my activity of doing such-and-such, that doing this *is* annoying my partner. And then it seems undeniable that I must know that I'm doing such-and-such, for each of the different ways in which I'm annoying her. So the fact that there may be more than one way of achieving one's end doesn't seem to undermine the argument.

Notice that the argument here can in a way repeat itself to bring out just what is known as the content of an intentional action. I have, throughout, spoken of doing and knowing one thing in doing another. But what is represented in the claim that "I'm doing *B* in doing *A*" may itself cover up further phases involved in my doing *A* and my doing *B*. Each may in turn unfold into ways of doing things — 'doing *C* in doing *A*' and 'doing *B* in doing *C*', and so on, may be the further phase or sub-actions involved in my doing *B* in doing *A*. But the content of the belief given by (1) will accordingly capture this, as the activity it records reaches out towards the agent's doing *B* in a way that is structured by the agent's practical reasoning — that is what such a belief will track. And so if we imagine that my doing *B* in doing *A* is some extended process, it is not just that we have shown that we have knowledge of the bit of activity at the end, and the bit at the start, but rather that we can see that we must also have knowledge of all the parts that come between too, at least in so far as they are phases that are marked out by the agent's answers to the question 'Why?'

In this way, I think we can see that *AP* is true. It may be said though, that even if this line of argument is to be accepted, it would only show that *AP* is true for intentional actions that involve successfully executing a further intention. Have we not then only offered a partial defense of *AP* proper? Are there not intentional actions where the agent is neither acting for some further reason nor through some particular means? If there are, the argument here won't apply to them, at least not directly. If there are such actions, some other story will be needed to cover them. Now, I don't want to deny that there are such actions. But must we say that they are intentional? I won't directly take up the question here, but I will return to it later, at the end of chapter 4. There, I'll suggest that Anscombe herself may have thought that such cases fall outside the scope of practical knowledge and are better described as a case of

what we can call “merely voluntary” action. I plan to follow her in that respect. But even if such a classification can’t be maintained, I can retreat to the view flagged at the outset of this chapter, that cases of doing one thing in doing another are surely the central phenomenon of intentional action, and if we have a story about why such actions always bring with them the self-knowledge of the agent, then I’ll be happy to rest content with having established that much.

### III. REASONS EXPLANATIONS AS CAUSAL EXPLANATIONS

“This is an original phenomenon of causality: one of its types ... No general theory about what causality is has to be introduced to justify acceptance of it. Nor does it have to be accommodated to any general theory, before it is accepted. It is just one of the things we mean by causality.”

— Anscombe, ‘Memory, “Experience”, and Causation’, p.127

“If this upsets our concept of causality then it is high time it was upset.”

— Wittgenstein, *Zettel*, §610, p.106e.

#### 1. INTRODUCTION

What is the relation between our actions and the reasons for which we act? If I start making dinner because I notice that it is 6 o'clock, or if I chop some mushrooms because I'm cooking pasta, or if I add some basil to the sauce because I know it will make it taste better, then I have, in each case, acted for a reason. At the least, the relationship between reason and action in such cases is *explanatory*: my reasons explain my action, or at least purport to. But what *kind* of explanatory relation is involved? In his 1963 classic ‘Actions, Reasons, and Causes’, Donald Davidson held that reasons explanations purport to *rationalize* an action: they show us “what the agent saw, or thought he saw, in his action” (Davidson 1963, p.3). To say only that much is, I think, to say something that everyone will accept. But as soon as we say more, disagreement will set in. Thus, a natural way to fill out Davidson's remark is that a rationalization tells us what *good* the agent saw in his action, and that is something which many will reject. In this

chapter, however, I want to revisit the further and main claim that Davidson argued for in his 1963 paper, which is that rationalizations are a form of *causal* explanation. Unlike the question of whether we act under ‘the guise of the good’, the question of whether reasons explanation are causal explanations may seem by now much more settled. For although Davidson was at the time running against a strong post-war current of ‘anti-causalism’ which took reasons explanations to be too bound up with certain teleological, normative, and empathetic considerations to possibly count as causal explanations, it has to be said that if we judge by the beliefs of contemporary philosophers of mind, action and ethical theory, then Davidson largely won the argument that we could view reasons explanations as a species of causal explanations without losing sight of their distinctive rational features.<sup>53</sup> Although there is still a vocal minority of anti-causalist holdouts, most philosophers follow Davidson in thinking reasons explanations are a species of causal explanation.<sup>54</sup>

And I will too. In §2, I will start by clarifying the terminology of ‘causalism’, ‘anti-causalism’ and ‘causal explanation’ as I understand it, and I’ll suggest that there is a weaker and stronger way to read Davidson’s argument in ‘Actions, Reasons, and Causes’, and that we only need the weaker (and valid) reading of the argument to see that reasons explanations are causal explanations. My aim, however, is not exactly to provide further support for the causalist side of the current debate. I will take as my point of departure the following remark about the state of the debate from one of its active participants:

Perhaps a source of part of this disagreement is a difference in how causation is conceived. In the literature defending or attacking causalism about action explanation, not much is said about what causation is. This is understandable: both sides may be thinking that, whatever the best account of causation is, their view about how actions are to be explained is the

---

<sup>53</sup> For helpful recent historical overviews of the background (and foreground) to Davidson’s argument, see Candlish & Damnjanovic (2013) and D’Oro & Sandis (2013).

<sup>54</sup> For varying views in this ballpark, Wilson (1989; 1997), Ginet (2002), Schon (2005), and Tanney (2009). See also the previous footnote and the references therein.



correct one to take, and they may want to avoid hitching their wagons to a specific theory of causation. Possibly differences in how causalists and anti-causalists ... understand causation help account for their different positions on whether causalism is true or false. Precisely because the body of literature does not have much to say about the nature of causation, the possibility just identified is difficult to explore. (Mele 2013, p.169)

I think Mele's hunch is right: how we understand the nature of causation is central to this debate. And yet, when we look to the philosophy of causation it can seem like it just doesn't matter which particular theory or analysis we opt for — debates between counterfactual theories or probabilistic theories or interventionist theories of causation can seem orthogonal to, or to cut across, the question of how we should construe reasons explanations. In effect, we have a kind of perfect storm where all we can do is hang on and watch the show continue, with no end in sight. That seems to be the implication of Mele's remark.

However, I suggest that there is an approach to the topic of causality that *does* make a difference to the discussion. On the approach I will explore and recommend, it will turn out that reasons explanations are causal explanations, but they are so in a manner that at least some anti-causalists should, in principle, be able to accept, at least if they accept the weak reading of Davidson's argument. On the other hand, although the picture of causality I adopt views reasons explanations as a species of causal explanation, it also has the effect of rendering some of the more familiar reductive and explanatory projects that most causalists find appealing entirely otiose. So Mele's hunch is right: how we construe causality matters. The account I offer, and the attendant picture of causal explanation, gives us a chance to follow Davidson, but by plotting a course that also sails between the Charybdis of the anti-causalist's obscurities and the Scylla of the causalist's overblown explanatory ambitions.

I don't devise the picture of causality *de novo*, but rather draw on the work of the other giant of contemporary philosophy of action: G.E.M. Anscombe. This may seem a strange source for this subject, because Anscombe has long been billed as an arch anti-causalist. I think that Anscombe's putatively anti-

causalist arguments have been misread, or at least overread, and I'd like to show why, but I'll relegate that discussion to an appendix. (See Appendix A.) Here, I want to focus on developing her positive story about causality. So, over the course of §§3–5, I will extract from Anscombe's famous 1971 lecture 'Causality and Determination' an account of causality that I call *causal pluralism*.<sup>55</sup> The view is not easily summarized in a way that lets us place it neatly alongside the extant options in the philosophy of causation. But the name is suggestive: for the causal pluralist, causality comes in a wide variety of flavors. As we'll see in §6, such a view of causality allows for a very capacious account of what causal explanation requires, one that we can exploit to show how reasons explanations are causal explanations — or, more precisely, we'll consider how the explanation of action by intention is a form of causal explanation. In §7, we'll consider the wider category of reasons explanations. But my story also has a kicker. As we'll see in §8, if the picture of causality here sketched is on the right track, then although we will also be in a position to see that the thought that we can employ a concept of causality to do serious explanatory work — of, say, a principle like *AP* — will seem seriously strained.

## 2. CAUSALISM, ANTI-CAUSALISM, AND CAUSAL EXPLANATION

To begin, then, just what does it mean to be a 'causalist'? And why should anyone care to be one anyway? I shall say, as I've been assuming already, that a causalist is someone who accepts that reasons explanations are causal explanations and an anti-causalist is someone who denies it. To try to help explicate that not very helpful definition we could survey the many different views that have gone by the name 'causalism'. But for our purposes, it will be best, I think, to start not by working through different

---

<sup>55</sup> The term 'causal pluralism' is also used by Psillos (2008) and Godfrey-Smith (2009) to describe Anscombe's view, though I'll offer a different interpretation of the view to theirs.

definitions of what causalism has been, but rather by attempting to elucidate it via what is often touted as the best argument for it. I have in mind the positive argument that Davidson gave in ‘Actions, Reasons, and Causes’. There are other arguments for causalism, but I focus on this argument both because of its status (it is sometimes referred to as ‘the master argument’),<sup>56</sup> and because working through the argument, and Anscombe’s response to it, will provide the essential backdrop to what follows.<sup>57</sup> Here is the argument:

Noting that nonteleological causal explanations do not display the element of justification provided by reasons [that is, of showing us what the agent thought ‘could be said’ for acting thus-and-so], some philosophers have concluded that the concept of cause that applies elsewhere cannot apply to the relation between reasons and actions, and that the pattern of justification provides, in the case of reasons, the required explanation. [...] Here it is necessary to decide what is being included under justification. It could be taken to cover only [...] that the agent have certain beliefs and attitudes in the light of which the action is reasonable. But then something essential has certainly been left out, for a person can have a reason for an action, and perform the action, and yet this reason not be the reason why he did it. Central to the relation between a reason and an action it explains is the idea that the agent performed the action *because* he had the reason. (Davidson 1963, p.9)

The distinction here is easy enough to illustrate. Suppose I want to get something to drink before I start work, and I know that getting a coffee from my local café would be a way of doing that; but suppose at the same time I also want to go sit in the sun for a little before I start work, and I also know that getting a coffee from my local café would do the trick (maybe they have great outdoor seats). If I now go and buy a coffee, either set of considerations could be offered in explanation of my action, in the sense of interpreting it, or rendering it intelligible. However — this is Davidson’s point — only one of them may

---

<sup>56</sup> Dancy (2000, p.161).

<sup>57</sup> For discussion of some of the other arguments in favor of causalism, see Child (1994, pp.98-9) and Davis (2005).

have in fact been (or included, if that is different) the *reason for which* I acted as I did. If Davidson is right, *this* is what we capture in saying that I went to the café *because* I wanted to get a coffee before work (assuming that that *is* why I went). The mere ascription of a reason — or, if the ascription holds true, the mere, as it were, *possession* of a reason — thus seems to fall short of capturing “the force of that ‘because’” (1963, p.9).

The pattern that *does* capture such an explanation, Davidson goes on to say, is the one given by “placing [the action] in the context of its cause; cause and effect form the sort of pattern that explains the effect, in a sense of ‘explain’ that we understand as well as any” (1963, p.10). And a little later in the paper he puts the idea (not, I think, equivalently: a point to which we return) by saying that what is needed to transform ‘X did *A* and X has reason *r* for doing *A*’ into ‘X did *A because* X has reason *r* for doing *A*’ is that X’s having reason *r* *caused* X’s doing *A*.<sup>58</sup> Since for Davidson a reasons explanation of this sort (typically) cites the attitudes or mental states which the speaker takes to rationalize her action, seeing how such explanations capture the content of this use of ‘because’ amounts to viewing such attitudes or states as the causes or causal conditions of the act in question.

So if a ‘causalist’ is anyone who accepts that reasons explanations are causal explanations, while an ‘anti-causalist’ denies that much, we can now add that the sense of ‘causal explanation’ involved is the one that Davidson’s challenge seeks to establish. This leaves open the possibility for a range of different causalist and anti-causalist positions, depending upon how exactly the notion of ‘causal explanation’ is understood. In particular, as a number of philosophers have urged, we can distinguish between two causalist positions that Davidson’s challenge might be thought to establish.<sup>59</sup> On the one hand, we might

---

<sup>58</sup> The syntactically fastidious must take these schemas (and the others that appear below) *cum grano salis*, since the ‘variables’ will typically require substitutions of more than what is explicitly marked here as up for substitution.

<sup>59</sup> See Hornsby (1993, 1999), Child (1994), and Steward (1997). I don’t claim these philosophers put the distinction in quite the way I do; the terminology of ‘minimal’ and ‘substantive’ accounts of causal explanation is my own.

say that Davidson's argument shows that reasons explanations are causal explanations in an intuitive or pre-theoretical sense.<sup>60</sup> I'll call this a 'minimal' picture of causal explanation. On the other hand, we might claim that Davidson's argument establishes that reasons explanations are causal explanations in a more specific sense, so that they are assumed to take on some or all of the features of what we have come to expect from a philosophical theory of causality, or of causal explanation, or of both. I'll call this a 'substantive' picture of causal explanation. I think we see the difference between the two kinds of position at work in the slide suggested above between whether Davidson's argument shows that to explain an action by a reason is a causal explanation and whether it shows that the agent's possession of that reason (say, the mental state she is in) causes her action. The difference is writ large in the bit of philosophical folklore that says that Davidson showed that "reasons are causes" — this (misleading) slogan goes far beyond anything that Davidson's argument necessarily shows, unless we simply assume that a causal explanation must work by citing the causes that brought the action about.<sup>61</sup> Such a view, which could obviously take very different forms, would be an example of what I'll call a more 'substantive' picture of causal explanation.

That there is a distinction to draw here can be brought out in different ways. In the first place, we can get the minimal picture in view merely by dint of some claims about the character of our ordinary explanatory talk. For there are clearly uses of 'because' which purport to be explanatory, but which it would be odd, from a naive perspective, to describe as causal explanations. If, for example, I say "I

---

<sup>60</sup> I don't mean to say that we necessarily have an intuitive grasp of the term of art 'causal explanation'. I mean we have an intuitive grasp, in Davidson's expression, of the force of the relevant 'because'.

<sup>61</sup> At least for Davidson, the terminology is potentially misleading because for him reasons — beliefs and desires — are states and not events (or entities of any sort), and only events can be causes. For Davidson, it is events 'associated with' such states that are, properly speaking, the causes of actions. Not all causalists agree. Some assume states can be causes, while others take the content of such states to do the relevant work. For some relevant discussion, see Dancy (2000), Davis (2005), and Setiya (2009).

broke the law because I parked on the yellow line” or “The bowler delivered a no-ball because she overstepped the crease”, I don’t seem to commit myself to claims about the causal efficacy of parking and overstepping, or anything along such lines.<sup>62</sup> Such claims seem rather to concern certain conventional or constitutive facts. On the other hand, there are uses of ‘because’ like “He calmed down because he took some medication” or “The car stopped because it ran out of fuel” which clearly *do* seem to be ‘causal’. And we can comfortably pepper or rephrase our talk about what happened to the man and the car using expressions like ‘brought about’, ‘led to’, ‘resulted in’, ‘effected’, ‘prevented’, or ‘interfered with’. For instance, it is perfectly fine to say things like, “As a result of running out of fuel, the car stopped”. This is a contrast with, to continue with our examples, our remarks about parking and overstepping creases. We can’t make sense of an utterance like “Your parking on the yellow line resulted in your breaking the law”.<sup>63</sup>

In the second place, our competent use of the causal ‘because’ and related causal idioms does not itself appear to carry any commitment to a particular metaphysics of causality of the more familiar stripe. We happily talk, for instance, of all sorts of things *as* being causes: people, objects, facts, properties, features, absences, omissions. All these items seem apt to do work in our everyday causal scheme. In the case of explanation, we are unperturbed about explaining an event in terms of a non-event (“The plant died because the gardener didn’t water it”), a non-event in terms of an event or state (“I didn’t read the book because it sounded stupid”), and even a non-event by a non-event (“The match

---

<sup>62</sup> These examples are from Child (1994, pp.91-2). I use these cases to help motivate the intuitive idea of causal explanation sketched below. Some philosophers deny that there are any non-causal explanations, but it isn’t obvious they would deny these claims about ordinary usage. The area is a bit of a terminological minefield: see Woodward (2003, pp.5-6).

<sup>63</sup> I don’t mean that ordinary language speaks with perfect clarity on such matters. I just mean that some such distinction is surely salient in our practice.

didn't light because it wasn't struck").<sup>64</sup> To demand that causal explanations be grounded in a unitary underlying relation between (say) efficacious particulars, or facts, or some other specific category or kind, or even just a serious sharpening up of this some of this messiness — a common preliminary to serious philosophical inquiries into the nature of causality — will be to take up commitments that go beyond what the minimal picture requires. Or so I claim.

Assuming that we can in this way distinguish between 'minimal' and 'substantive' pictures of causal explanation, we can distinguish two ways to read Davidson's challenge. In the first place, if we read Davidson's argument as an attempt to establish that reasons explanations are causal explanations in the sense of the minimal picture, then we take it to have the force of something akin to a Wittgensteinian 'reminder', one that makes us recall the fact, if philosophy has taught us to forget it, that our practice includes this causal 'because' and that it encompasses the rationalization of action. Read this way, the issue is squarely joined with the anti-causalist, because the anti-causalist insists the usage of 'because' and the various attendant causal idioms stressed above cannot *au fond* be understood in the manner Davidson suggests. In some sense, the anti-causalist may be happy to call such talk 'causal' talk, but she will insist that this 'because' and the attendant idioms do not mark or express any kind of causal understanding of the coming about of a state or an event.<sup>65</sup> As levelled against the anti-causalist position, I think the argument works: it shows that reasons explanations are causal explanations.

However, we can also see Davidson's challenge as an attempt to establish that reasons explanations are causal explanations in the sense of the substantive picture. Of course, there are many

---

<sup>64</sup> For discussion of such cases in the current context of causal explanation, see Steward (1997) and Hornsby (1999).

<sup>65</sup> My terminology may be a little idiosyncratic. Although this captures the 'anti-causalists' of Davidson's day well enough, more recent anti-causalists — like Wilson (1989; 1997), Ginet (2002), Schon (2005), and Tanney (2009) — sometimes only define their view negatively, telling us that they are rejecting a particular species of causal explanation as capturing reasons explanation, but they don't always tell us what they mean by the genus, so it isn't clear if they count as anti-causalist on my terminology.

such pictures. But whatever particular picture we choose, the argument is invalid, since the bare point that rationalization involves a causal ‘because’ or a kind of causal understanding does not itself force us to adopt some specific metaphysics or epistemology that goes beyond what we find in the minimal picture.

### 3. KNOWLEDGE OF CAUSALITY

What I want to do, then, is develop an account of what causal explanation involves which doesn’t go past the minimal story, which itself has as minimal commitments as possible. As David Lewis (1986) points out, an account of causal explanation is not the same thing as an account of causation, but it is hard to deny that the story we tell about one will to some extent be influenced by the story we tell about the other. What I want to do here is put in place the picture of causality that I think will better facilitate a minimal story about causal explanation than the available options.

To do that, I will mine Anscombe’s 1971 lecture ‘Causality and Determination’. Although the paper is duly famous for its negative claims, the strikingly distinct positive picture of causality therein sketched is not well appreciated, and almost completely undiscussed in the philosophy of action.<sup>66</sup> A good deal of contemporary philosophy of causation views a causal claim in the way that Carl Hempel (1965, p.349) memorably described as “like a note saying there is a treasure hidden somewhere” — with, of course, the philosopher’s theory providing the map to the booty. In that sense, Hume casts a long shadow. But early on in ‘Causality and Determination’, Anscombe puts in place a very different view.

---

<sup>66</sup> Hornsby (2011) is an exception, to whom I’m much indebted, though I tell a very different story to her, and she takes Anscombe to deny that reasons explanations are causal explanations. More generally, in helping me with Anscombe’s difficult paper I am indebted to Makin (2000), Bogen (2008), Godfrey-Smith (2009), and Psillos (2009), though none develop Anscombe’s views in the context of her philosophy of action, and I don’t follow them in the details, as we’ll see.



After a brief survey of what she sees as the ubiquitous role in the history of philosophy of ‘necessitation’ theories of causality — that is, theories that require (singular) causal relations to instantiate laws or universal generalizations linking events of the sort mentioned as cause and effect — Anscombe claims that it is at least *prima facie* obvious that we often have knowledge of a cause in ignorance of such laws or generalizations. In support of the assertion, she gives the example of tracing the contraction of a disease to the sole contact the infected person had with a carrier. In reply to the objection that such knowledge could be at best partial, Anscombe raises this possibility:

[M]ight it not be like this: knowledge of causes is possible without any satisfactory grasp of what is involved in causation? Compare the possibility of wanting clarification of ‘valency’ and ‘long-run frequency’, which yet have been handled by chemists and statisticians without such clarification ... Thus one of the familiar philosophic analyses of causality, or a new one in the same line, may be correct, though knowledge of it is not necessary for knowledge of causes. (Anscombe 1971, p.136)

There are two things I want to note about this passage. First, although ‘Causality and Determination’ is focused primarily on the necessitation view of causality, it is clear from this paragraph that Anscombe means to challenge a wider range of theories — she has in view both the (then) alternative analyses of causality on offer *and* any analysis or theory ‘in the same line’.<sup>67</sup> Second, as I read ‘Causality and Determination’, Anscombe is here sketching the overall shape of the argument of the paper.<sup>68</sup> The

---

<sup>67</sup> Perhaps “clear” is too strong. Maybe Anscombe only means analyses ‘in the same line’ as necessitation theories, i.e. ones that are bound up with nomic concepts. That is a possible reading of the passage, but it strikes me as a very unlikely one, given various other remarks Anscombe makes in ‘Causality and Determination’ and in other places, such as the remark of hers that serves as my epigraph.

<sup>68</sup> The paper, or the sections of the first part with which I am concerned, is usually read as primarily about the narrower topic of our *perception* of causality rather than our *knowledge* of it, with the overall shape of the argument thought to run something like “We can perceive efficacy, so necessitation theories are wrong”, or “We can perceive efficacy, so no analysis of causality is possible” (cf. Beebe 2003; Sosa & Tooley 1993). But her well-known rebuke to Hume’s claim

argument is that no such theories of causality are necessary or sufficient for capturing the content of our ordinary causal knowledge. To better grasp the argument — this is our goal over the next few sections — we need to know more about how Anscombe construes our ordinary causal knowledge; how she envisions one of the familiar theories of causality being ‘correct’ while nevertheless failing to capture the content of such knowledge; and what she means by a theory being ‘in the same line’.

Anscombe offers material for beginning on answering these questions in the following, oft-cited passage:

There is something to observe here, that lies under our noses. It is little attended to, and yet still so obvious as to seem trite. It is this: causality consists in the derivativeness of an effect from its causes. This is the core, the common feature, of causality in its various kinds. Effects derive from, arise out of, come of, their causes. For example, everyone will grant that physical parenthood is a causal relation. Here the derivation is material, by fission. Now analysis in terms of necessity or universality does not tell us of this derivedness of the effect; rather it forgets about that. For the necessity will be that of laws of nature; through it *we* shall be able to derive knowledge of the effect from knowledge of the cause, or vice versa, but that does not show us the cause as source of the effect. (Anscombe 1971, p.136)

And she goes on to claim that the lack of entailment between the two types of knowledge works in the other direction too. Commingling the ‘derive from’ and ‘come of’ talk, she says:

---

that we cannot perceive efficacy is made as part of a response to an objection that comes *after* the argument of the paper (on my reading) has already been put in place. She herself says, “my present topic is not the philosophy of perception” (Anscombe 1971, p.138). I should also remark here upon the tendency of writers to lump Anscombe together with other so-called ‘singularists’ about causation (like David Armstrong) who not only argue that causality is observable, but take a paradigm case of it to be the observation of the activity of the will, or the inner perception of bodily movements and sensations. This adds another layer of confusion, since Anscombe’s views on that score are quite different. For discussion of the latter, see McDowell (2011a).

If *A* comes from *B*, this does not imply that every *A*-like thing comes from some *B*-like thing or set-up or that every *B*-like thing or set-up has an *A*-like thing coming from it; or that given *B*, *A* had to come from it, or that given *A*, there had to be *B* for it to come from. Any of these may be true, but if any is that will be an additional fact, not comprised in *A*'s coming from *B*. If we take “coming from” in the sense of travel, this is perfectly evident. (Anscombe 1971, p.136)

It is implicit in the first of these two quotes that the “knowledge of causes” from our previous citation at the opening of this section is what she here refers to as knowledge of “the cause as source of the effect”. This latter knowledge is distinguished from mere “knowledge of the effect [given] knowledge of the cause, or vice versa”. Now, if we gloss the first sort of knowledge about derivation as involving a capacity that provides information about *how* the effect has been brought about, and we also gloss the second sort of knowledge as involving (at most) a capacity for accounting for the existence of the effect given knowledge of the cause (or vice versa) — say (most obviously) by predicting the occurrence of the effect given knowledge of the cause, or retrodicting the existence of the cause from knowledge of the effect — then this would give us some purchase on what Anscombe may have in mind by a theory being merely ‘correct’ without it yet providing knowledge of ‘the cause as source of the effect’, since merely being able to predict an effect will not yet tell us anything about how a cause works.

How is such talk of the ‘work’ of a cause to be understood? It is, I think, easy to misread Anscombe here. Jaegwon Kim (2007, p.235), for instance, suggests that Anscombe’s notion of ‘derivativeness’ should be taken to refer to a particular spatiotemporally contiguous physical process between cause and effect, or between a series of such causes and effects. The examples of cell fission and travel may be taken to suggest as much, and unlike Anscombe’s seemingly trite characterization, a gloss along these lines would hold out a promise for a sharper account. But although the issue is not clear in ‘Causality and Determination’, Anscombe in fact has a much broader sense of ‘derivativeness’ in mind. Thus, in ‘On Being in Good Faith’, she says:

There are [...] two different ways in which one thing is said to come of another. First, directly, when the other acts, so as to produce it — e.g. something gets hot by being heated by some other hot thing. Second, indirectly, from the mere fact that the other thing does *not* act. For example, we say a ship sank because of the pilot, if he stopped navigating it. (Anscombe 2008a, p.109; dating originally from the late 1950s or early 1960s)

Examples like that of the pilot and the ship rule out that causality must consist in a spatiotemporally continuous physical process.<sup>69</sup> Moreover, in introducing this second ‘indirect’ type of causality Anscombe means to include not only cases where the pilot makes a deliberate decision not to act so as to allow the ship to sink, but also cases that include what she calls an ‘omission’ — that is, cases where the agent need make no such decision in failing to act but is described as the cause of an event on grounds that involve ‘normative’ considerations concerning their station and status. As she puts it in another early paper, the pilot can be the cause of the ship’s sinking if his failure is due to ignorance for which he is culpable, where deciding the latter matter will involve us in questions as to what it is ‘necessary’ for the agent to know, in a sense of things being necessary “when without them some good can’t be got or some evil avoided” (Anscombe 1963c, p.9).

Evidently, Anscombe’s understanding of causality is generous, far more so than many philosophers would like.<sup>70</sup> But we still lack a clear sense of it. For some philosophers, Anscombe’s discussion simply falls short here. Thus, Michael Tooley (1987, p.189) complains that in Anscombe’s paper “everything turns upon the notion of an effect’s deriving from, arising out of, its cause, and no

---

<sup>69</sup> I don’t mean to imply that she views the first ‘direct’ way in which something comes of another as necessarily explicable in terms of such a process.

<sup>70</sup> According to Peter Menzies’s (2009) survey, one of the few things that philosophers of causation have widely agreed upon is precisely that causation is a ‘natural’ relation that cannot entangle us in ‘normative’ considerations, but that now seems too strong. For a small sample of recent discussion of ‘omissions’ in Anscombe’s narrow sense, see Alvarez (2001), Menzies (2009), Hitchcock & Knobe (2009), and Danks *et al.* (2013). For debate about the wider notion of ‘negative’ causation or causation by ‘absence’ see Dowe (2004) and Schaffer (2004).

analysis of this crucial notion is offered”. But while it is true that Anscombe offers no further ‘analysis’, she does have more to say. Coming at the question in a manner not uncharacteristic of someone who follows in Wittgenstein’s tracks, Anscombe goes on to offer the following elucidatory remark concerning our acquisition and expression of causal knowledge:

The truthful — though unhelpful — answer to the question: “How did we come by our primary knowledge of causality?” is that in learning to speak we learned the linguistic representation and application of a host of causal concepts. Very many of them were represented by transitive and other verbs of action used in reporting what is observed. Others — a good example is “infect” — form, not observation statements, but rather expressions of causal hypotheses. (Anscombe 1971, p.137)

The answer to the question is unhelpful, one presumes, because the acquired concepts are those that are apt to express causal knowledge. Conspicuously absent here is the causal concept we might have thought most obviously suitable to express such knowledge: the one expressed by ‘cause’. But ‘cause’ does come up for mention in a remark that follows the previous one, and is brought into connection with the above causal concepts in an intriguing way:

The word “cause” itself is highly general. How does someone show that he has the concept *cause*? We may wish to say: only by having such a word in his vocabulary. If so, then the manifest possession of the concept presupposes the mastery of much else in language. I mean: the word “cause” can be *added* to a language in which are already represented many causal concepts. A small selection: *scrape, push, wet, carry, eat, burn, knock over, keep off, squash, make (e.g. noises, paper boats), hurt*. But if we care to imagine languages in which no special causal concepts are represented, then no description of the use of a word in such languages will be able to present it as meaning *cause*. (Anscombe 1971, p.137)

The point of the passage is surely to mark some kind of dependency between what Anscombe calls ‘special’ causal concepts and the ‘general’ concept expressed by ‘cause’. Her exact meaning does not

exactly jump off the page though, let alone her grounds. The most pressing question, I think, is what is meant by the idea of being able to ‘add’ a word or concept to a language that already contains others.

If we look to the literature for enlightenment, we do find something like a standard reading of the passage available, or at least a standard way of trying to explicate what Anscombe might be up to.<sup>71</sup> The basic idea behind this reading is that Anscombe’s point in the passage is to get us to see that the expression ‘cause’ is in principle eliminable or expendable from our language. On this view, she thinks the content of any claim made employing ‘cause’ could be paraphrased away without remainder using sentences with special causal concepts substituted for occurrences of ‘cause’, modulo whatever syntactic adjustments would be required. Thus, Peter Godfrey-Smith condenses Anscombe’s claim into the following conjunction:

1. ‘*C* was the cause of *E*’ is true iff the relation between *C* and *E* can also be described using some member of set *S*, or can be described as a chain of relations each of which can be described using some member of *S*.
2. *S* is a set of causal verbs and other linguistic formulas which represent ‘special causal concepts’ in Anscombe’s sense. (Godfrey-Smith 2009, p.333; cf. Psillos 2008)

The thought here is that this is a kind of ‘minimalism’ about causality akin to that which we find in discussions of truth, so that the first claim here captures the thought that we use the word ‘cause’ to make certain convenient abstractions and generalizations, much as the minimalist about truth thinks the word ‘true’ has many such convenient functions, even if it doesn’t denote a substantive property. Now this interpretation certainly has the virtue of clarity. But if that were what Anscombe meant, then she would face the following objections: (1) it isn’t clear how to determine the extension of *S* (Snowdon 2011, p.123); (2) it won’t be clear how we perform the operation of abstraction in the first condition,

---

<sup>71</sup> For versions of this reading, more and less explicit, and with difference in detail, see Psillos (2008), Godfrey-Smith (2009), and Snowdon (2011).

since Anscombe's examples of special causal concepts seem to describe causal transactions between objects — they describe *substance* causation we might say — rather than relations between events or facts (Godfrey-Smith 2009, p.333); and (3) she won't be able to explain our use of the word 'cause' when we're *searching* for causes, as we routinely do (Snowdon 2011, p.122, Godfrey-Smith 2008, p.334).<sup>72</sup>

But this can't be what Anscombe has in mind. There are small replies to make on her behalf: (1) we can surely grasp the rough extension of *S* through examples;<sup>73</sup> (2) it is obviously no part of Anscombe's view that we must be able to reformulate causal claims into relational claims between events or states; and (3) her subject is the expression of causal *knowledge*, not causal *ignorance*. The larger issue, though, is that this reading of the passage transforms Anscombe's claim about the conditions under which it would be possible to *add* the word 'cause' to a hypothetical language into a claim about the possibility of *removing* it from our actual language. Why should we read her like that? I think there are two assumptions at work here. The first is the idea that she means to reduce our understanding of a causal concept to something that could be captured or expressed by a sentence. It is understandable that the

---

<sup>72</sup> As Snowdon (2011, p.122) puts it, "We are completely happy to ask whether smoking is a cause of cancer, or low interest rates a cause of inflation, and whether banker's bonuses were a cause of the current economic crash, and so on" — all cases where we shouldn't be able to intelligibly ask the question if we needed to grasp a special causal concept to do so.

<sup>73</sup> We can surely offer a few exemplars of non-causal concepts alongside her exemplars: mathematical and logical concepts, for instance, must count as non-causal, if anything does. It is tempting to look to the well-known entailment that holds between many of the transitive verbs that Anscombe lists and the general concept of cause for a possible criterion — for example: if *X* moves *Y* then *X* causes *Y* to move, or if *X* butters *Y*, then *X* causes *Y* to be buttered (for some relevant discussion of this schema, see Thomson 1987). But while this schema surely captures something central to our grasp of special causal concepts, I would not want to say on Anscombe's behalf that adherence to it can be used to define the class in question. Anscombe offers no criterion for what counts as a 'special causal concept', and I don't think we should try to proffer one on her behalf, for a reason that will be clear in a moment.

remark in the above passage that having a concept can be marked by having a word could encourage that. But that isn't Anscombe's meaning. Rather she thinks this:

The competent use of language is *a* criterion for the possession of the concepts symbolized in it, and so we are at liberty to say: to have such-and-such linguistic practices is to have such-and-such concepts. "Linguistic practice" here does not mean merely the production of words properly arranged into sentences on occasions which we vaguely call "suitable". It is important that it includes activities *other* than the production of language, into which a use of language is interwoven. For example, activities of measuring, of weighing, of giving and receiving and putting into special places, of moving about in a huge variety of ways, of consulting tables and calendars and signs and acting in a way which is connected with that consultation. (Anscombe 1976, p.117)

This is why, immediately after giving her list of special causal verbs, Anscombe goes on to say that learning such verbs cannot be pried apart from learning "words for natural kinds of stuff, nor yet words equivalent to 'body', 'wind', or 'fire'" (Anscombe 1971, p.137). So it seems a mistake to think that Anscombe is suggesting that we can somehow distill our causal understanding into little bits of knowledge given by sentences of the form 'Substantive + Special Causal Verb + Substantive'.

The deeper assumption, I think, is the idea that Anscombe must in some way view the abstract concept of *cause* as a sort of second-class citizen, as something not on a par with 'ordinary' or 'commonsense' special causal concepts. Thus, Snowdon (2011, p.122-123) remarks that in describing 'cause' as a "highly general" concept, Anscombe must view it "as akin to highly general categories that philosophers have introduced and use, categories such as a material continuant, or abstract entity, which themselves do not figure in ordinary thought". But what does it mean for something to 'figure in' ordinary thought? As we've seen, Snowdon takes it to mean that the fact that people routinely use the word 'cause' and other cognate 'abstract' causal terminology is a problem for Anscombe. Similarly, Godfrey-Smith (2009, p.334) says that our everyday talk about 'causes' or 'making a difference' presents "so abstract a mode of assessment that it is hard to make sense of within an Anscombe-inspired



approach”. But I think we only need to state this worry to see that it must distort something in Anscombe’s position.

#### 4. AN ANALOGY WITH *INTENTION*

We can see what is wrong with the standard reading by considering the methodology of Anscombe’s treatment of ‘intention’, at least on the account of it due to Michael Thompson (1992). To understand what Anscombe is up to in *Intention*, Thompson asks us to first suppose we are confronted with a conception of intention that we often come across in philosophy, one wherein the concept is viewed through “a purely technical conception, [...] as if its relation to reality were to be secured through a philosopher’s or jurist’s definition” (Thompson 1992, p.17). Such a conception, Thompson points out, will force a putatively basic ethical principle that presupposes the concept — such as the doctrine of double effect — to take on a peculiar epistemic status: it will render it a hypothesis, one that can only earn its keep through its ability to explain our ‘moral intuitions’. Faced with this sort of technical approach, “the correct response”, says Thompson, “was to write *Intention*”:

For, in the light of Anscombe's teaching, if it is true, we can at last see how the vocabulary of intention need not be viewed as a foreign importation when it turns up in a discussion of the foundations of ethics or in a theory of rationality, and how a person who formulates basic principles in terms of it need not be guilty of dispiriting innovation. (Thompson 1992, pp.18-19)

According to Thompson, *Intention* does this both by revealing the agent’s answers to the series of questions ‘Why?’ as the ‘primitive’ form of expression of intention and by showing us how this primitive form of expression provides the core content of a more reflective form of expression, one involving the first-personal, present-tense use of the verb ‘intend’. With the more sophisticated form of expression in

hand, we can then talk about the same reality given by the more primitive form, and we can do so in a philosophical discussion of an ethical principle like that of ‘double effect’. As Thompson puts it:

What Anscombe does is to bring out the — to us — hidden character of the vocabulary of intention as a merely mediated and articulate mode of expression of *the point of view that human beings — agents — themselves characteristically take in thinking practically*. This is the essential difference between such a concept as *intention*, and that of a sacred scripture, of a brain state or of natural selection. A human being can think practically without having a proper concept of intention (which presupposes the wider range of powers), and without understanding any verb for intention. But this vocabulary is unlike the neurological in that the description in its terms of a person who lacks the vocabulary does not put us onto any facts, any reality, that any deliberating person is not in a position to grasp. Thus the proponent of the doctrine of double effect need not found his principle [in Rousseau’s famous words] on “the belated lessons of wisdom”, nor is he “forced to make a man a philosopher before making him a man”. (Thompson 1992, p.20)

Anscombe’s discussion of causality in ‘Causality and Determination’ is far sketchier than the discussion of intention in *Intention*. But I want to suggest that Anscombe’s talk of ‘adding’ the concept of *cause* to a language should be read in just the same way Thompson suggests we read *Intention*.<sup>74</sup> We might illuminate the analogy by reformulating Thompson’s remarks in the following way: what Anscombe purports to show us in *Intention* is how a word like ‘intend’ could be added to a language which already contains the question ‘Why?’ and its canonical ‘primitive’ answers. Or like this: if we care to imagine languages in which the ‘Why?’ question and its range of answers is not represented, then no description of the use of a word in such languages will be able to present it as meaning ‘intend’.

---

<sup>74</sup> Compare also in this connection Anscombe’s (1976, p.121) remarks about ‘a right’: “the *concept* ‘a right’ may indeed not exist [in a human society], but the restriction on doing what it is another’s to do (for example) or on preventing him, which are accomplished by the art of linguistic practice – these seem to be found everywhere. The concept ‘a right’, like that of ‘a relation’, is the product of sophisticated reflection on the data”.

Working the analogy back in the other direction, we can say that Anscombe's aim in 'Causality and Determination' is to bring out that the general concept of *cause* draws its content from the more 'primitive' forms of expression given by our stock of special causal concepts. As Judith Jarvis Thomson (1987, p.110) nicely puts it in a related connection, 'cause' is "a latecomer in a language, and is learned late". If this interpretation is right, then Anscombe's point is not to show that we can or should do away with the general concept of *cause*, or that the word 'cause' has a 'technical' status with no currency in ordinary thought and talk. The aim is rather to show how a philosophical inquiry into causality, or into doctrines that employ the concept, even if it ascends to a level of abstraction which does move beyond our ordinary causal vocabulary, nevertheless need not put us onto a distinct reality from that characteristically represented by a human agent's ordinary thought and talk and about how, among much else, things in the world scrape, push, cut, wet, squash, and hurt one another.

## 5. CAUSAL PLURALISM

Let us take stock. We wanted to know how Anscombe construes 'knowledge of causality'; what she means by describing a theory of causality as falling into a certain class (being 'in the same line' as other theories); and why she thinks such a theory could be 'correct' about causality while failing to capture the content of our causal knowledge. I suggested the answer to the last question was that such a theory may let us infer causes from effects (or vice-versa) while leaving out *how* a cause (or the causality) works. We also saw that Anscombe does not always take the work of a cause to involve some type of physical mechanism; a more general account was wanted. And I think we can now say that for Anscombe to have causal knowledge — to grasp how a cause works — is just to have the kind of knowledge that we canonically express using special causal concepts. "What caused the tree to fall over? The wind blew it over". To know the wind as the causal source of the effect is, in this case, to know that the wind blew

the tree over. On her view, there may be no more and no less to distinguishing a causal relationship from a merely correlational one than grasping such a brute fact. This, I think, is the radical heart of Anscombe's approach to causality.

We now also have an idea of what Anscombe may mean by a theory of causation being 'in the same line' as the necessitation account or the various other extant analyses. She means this: such a theory will be one that does not make essential (ineliminable) usage of our grasp of special causal concepts in its formulation; rather, it will attempt to understand causality from a perspective that is, to borrow John McDowell's (1994) phrase, 'sideways on', from outside this sphere of concepts. This is the point of the closing claim of the passage, that "if we care to imagine languages in which no special causal concepts are represented, then no description of the use of a word in such languages will be able to present it as meaning *cause*".<sup>75</sup> It is not just that a language must contain special causal concepts for an account of 'cause' (or one of its types) to be expressible within it; the account itself must be restricted by the content of such concepts; it must be given either in terms of such concepts, or by material which cannot be grasped without them.<sup>76</sup>

If Anscombe is right, we can't grasp the concept *causality* independently of grasping special causal concepts and the concepts and abilities that come with them; the only *general* characterization of the concept across its different types that we have is given by the idea 'derivativeness', something itself

---

<sup>75</sup> To be clear, Anscombe's point is, in the Wittgensteinian jargon, about the description of the 'use' of a word and not a 'theory' or 'analysis' of a word or concept. I am assuming that her claim will carry over to explanatory projects of the latter sort just as well, even if she herself shows no interest in them, in this context.

<sup>76</sup> Of course, some philosophers also eschew the goal of providing an analysis of causation, and some believe that there are different senses of causality that need to be kept apart; but none, so far as I can see, go for quite so radical a view as Anscombe. For some discussion of the matter, and comparison with views that lean in Anscombe's direction, see Psillos (2008).

grasped *through* special causal concepts.<sup>77</sup> If I understand her, Anscombe would say that *causality* is a ‘categorical’ concept; or, in the idiom she often prefers, a ‘grammatical category’. To have given up on the various familiar reductive and sideways-on accounts of causation is not to say that we have ruled out saying anything about how the genus may relate to its species.<sup>78</sup> In particular, it can still be the case, as Anscombe (1981d, p.201) says, that such categorial distinctions “turn out to be somewhat generic: there are ‘categorical’ differences within each kind”. But when Anscombe, rather remarkably to most modern ears, chides Aristotle for introducing *too few* causes — “four is not enough”! (Anscombe 1983, p.91) — she is not urging us to begin writing treatises with titles like *On Scraping*. We can call those bottom level special causal concepts *kinds* of causalities if we like, but when, say, Anscombe talks of what she calls ‘mental causes’ as “a sort of causality or a sense of ‘causality’”, she is locating a kind higher up the categorial ladder, one which we grasp through the rungs below.<sup>79</sup> It is these types of causalities which are suitable objects of theoretical inquiry and philosophical reflection. And so what I think is really

---

<sup>77</sup> To be clear, I don’t think this means that Anscombe would somehow rule out *any* kind of further philosophical reflection on how the concept relates to other concepts. To my mind, is hard to imagine, for example, how we could not think that the notion of causality is tied to the concept of *change*, and I can’t see Anscombe saying anything to rule out pursuing the question. For a relevant recent discussion of the latter, see Ford (2014).

<sup>78</sup> For instance, Anton Ford (2014, p.29) suggests that Anscombe’s position can be profitably understood as treating causality as a *determinable* genus with various *determinate* species: much as *red* and *blue* are determinates of *color*, so *kicking* and *pushing* would be determinates of *causality*. I think it is hard to say whether that is a profitable way to proceed, as the determinate/determinable framework remains contested, and precisely in ways that matter to its possible application here. Thus, some philosophers hold that the framework can only be legitimately invoked if the putative determinations are comparable along certain dimensions of variations, as, for instance, the determinates of *color* vary in degree of hue, saturation and brightness (see, e.g. Setiya 2009, pp.141-2). That doesn’t seem to bode well for applying it to causality — what would the dimensions be? On the other hand, Ford (2011, pp.84-5, *n.*16, 17) suggests there is an older, pre-Johnsonian construal of the relation that runs from Aristotle to Brentano which doesn’t impose this stricture on its application. And there are other points of disagreement about how to understand this framework, including whether typical cases really are typical. For the state of the art, including discussion of many of these issues, see Wilson (2017).

<sup>79</sup> See Appendix A for detailed discussion what Anscombe means by a ‘mental cause’.

distinctive of Anscombe's approach to causality — the approach I'm calling her 'causal pluralism' — is not just that there are different kinds of causalities, but that our grasp of the different kinds is bound up with certain clusters of concepts.

For Anscombe, a causal inquiry, whether of something as quotidian as the wind blowing things over or as rarified as chromosomes dividing by mitosis, will involve a piecemeal investigation into various local matters of fact, in such a way that the type of causality under investigation may be inextricably tied up with the particular facts at hand, which will in turn be given to us through our grasp of the special causal concepts that refer to the activities or processes involved and the other non-linguistic activities involved in our interaction with the relevant subject matters.<sup>80</sup>

Of course, to the hard-headed philosopher of causation, this may sound like a dissolution of the topic. But the proof is in the pudding. The question is whether there is any substance to the thought that we can pursue causal inquiries in this vein, whether doing so provides insight and understanding. Let us try to bring out, in a cursory way, how this might work in the case of our subject matter, intentional action. At the end of her paper 'The Causation of Action', Anscombe asks us to consider "the causalities especially involved in a history of people's dealings with one another". I cite at length:

The first thing to note is: these causalities are mostly to be understood derivatively. The derivation is from the understanding of action as intentional, calculated, voluntary, impulsive, involuntary, reluctant, concessive, passionate, etc. The first thing we know, upon the whole, is what proceedings are parleys, agreements, quarrels, struggles, embassies, wars, pressures, pursuits of given ends, routines, institutional practices of all sorts. That is to say: in our descriptions of their histories, we apply such conceptions of what people are engaged in. In the context of such application, then, the causalities to which we ascribe such events

---

<sup>80</sup> For a very helpful discussion of examples of how this works out in particular scientific inquiries, and for the helpful description of the inquiries as 'piecemeal', see Bogen (2008). For related further discussion which is not as closely tied to Anscombe's position, but which is still inspired by it, see Cartwright (2004).

can, so to speak, get a foothold. Given the idea of an engagement to marry, say, you can look for its causal antecedents. Or again, this man was travelling from Aix to Ghent. What for? He was a messenger taking news. So in the situation in which the news was generated, and in which there was a requirement that he should take it, together with the instructions of whoever sent him off, and the exigencies of route or difficulties posed by his means of carrying out the purpose, together with accidental encounters and concatenation of events with aspects of temperament and facts of people's excitements — all these will contribute causalities of various kinds to the event of his arrival or non-arrival at his destination. The causalities will, for example, include negations. Because this man did not know this language, he went this way rather than that: a very different sort of causality from that of the issuing of a certain order to him at a certain moment. Or again: *because* he was quick tempered, he got into a rage because of a supposed insult, and *because* of that all unawares escaped certain dangers or involuntarily fell into other ones. (Anscombe 1983, p.107-8)

The parallel with the doctrine in 'Causality and Determination' is striking. In the domain of human action,<sup>81</sup> our 'primary knowledge' is of which proceedings are *parleys, agreements, struggles, marrying, travelling* etc. It is *through* that primary knowledge that we grasp the different causalities involved here. To grasp the causalities involved in this region appears to involve a very disparate set of special causal concepts — we seem to have moved some way up the ladder from *pushing* and *making, kicking* and *scraping* — and it may be hard to grasp their common core, but if we again slide back down the ladder a little and draw some distinctions among the different kinds of causalities being lumped together, we may also find clusters of concepts that express our 'primary knowledge' of the specific causalities involved. This is just Anscombe's view about intentional action, as I understand it. That is, we grasp the "the causality whose concept we use when we explain an action by its intention" (Anscombe 1979, p.80) through the 'primary knowledge' given by our grasp of certain special causal concepts, what Anscombe (1963a, p.86) calls

---

<sup>81</sup> In the final chapter, we'll see that Anscombe construes the category of 'human action' as having a wider extension than that of 'intentional human action'.

“intention-dependent concepts”. And to grasp the structure of the causality involved here, we don’t have to write the treatise, since Anscombe wrote it already.

## 6. CAUSAL EXPLANATION BY INTENTION

To return to the main thread, how does causal pluralism help us see whether reasons explanations are causal explanations? It does so in two ways. First, by allowing us to sketch, in a very rough way, a general picture of what a causal explanation is. And second, by allowing us to narrow in on the particular type of causality that is characteristic of the explanation of intentional action.

Now I don’t mean to say that there is logical entailment from Anscombe’s picture of causality to a particular account of causal explanation. But the causal pluralist story surely suggests a picture of what causal information or causal understanding would consist in. Without offering a detailed picture of the syntax, semantics, and pragmatics of causal explanation, the picture I think it suggest is this: a causal explanation explains an event by *providing information about the derivation of an effect from a cause*.<sup>82</sup> Like Anscombe’s account of causality, this may well appear trite. Indeed, it doesn’t add much to Anscombe’s account of causality at all; it merely reformulates it in a slightly different key. Obviously, it is not an attempt at a reductive analysis of ‘causal explanation’. But, like Anscombe’s account, it is not trivial: it places, in different ways, both stricter and slacker demands on causal explanations than those that philosophers normally recognize.<sup>83</sup>

---

<sup>82</sup> I use ‘event’ here, but I want to allow for whatever kinds of results the causal pluralist picture will allow for. It will include omissions, for instance (or ‘negations’, as Anscombe says).

<sup>83</sup> Though compare David Lewis’s (1986, pp.219-220) account of causal explanation as providing information about the causal history of an event. Lewis’s view resembles the current account in also allowing for different kinds of causal information, though I think the current account is account is less wedded to a particular metaphysical story than Lewis’s.



It is a stricter account in that the requirement to be told of how a cause (or a bit of causality) *works* means that if we want to know why (say) the door closed, we haven't yet been told if all we know is (say) that there is a lawful regularity between events of this sort and those of a kind that preceded it, or that a physical quantity has been preserved between it and the door's closing, and so on (though any of these things may in fact be true). But, as just noted, it is more relaxed in that it carries over from the account of causality the possibility of a wide range of different kinds of causal information. One sort of information would identify a state or event as an antecedent or concomitant cause of the event mentioned in the explanandum. Another sort would mention the absence of such a state or event and how that affects what happens. Yet another sort (the one we'll focus on momentarily) would locate the effect as a part of an ongoing process of one particular type. And there will be other forms too —really, anything we find in our ground-level explanations employing special causal concepts could be available.

So what about reasons explanations? First, let us recall some of the different guises they can take. Following Setiya (2009) again, we can take canonical reasons explanations to have one of these three shapes: (1) 'X is doing *A* in order to do *B*', (2) 'X is doing *A* because *p*', and (3) 'X is doing *A* on the ground that *p*'. In the second and third cases, '*p*' stands for a variety of items, but most typically: ascriptions of mental states to the agent, facts which do not mention the agent, or facts about what the agent is doing. Now although I've spoken so far of 'reasons explanations', I here in fact want to make a simplifying assumption and restrict our attention only to instances of (1), (2), and (3) where the agent could just as well explain their action by offering an instance of (1), or a instance of (2) in the shape of a naïve rationalization, or just by mentioning the intention with which they are acting. Often enough, the agent will of course be able to do this. If he can, let us say he has given an "explanation by intention".

Why the restriction? The reason is that my aim here is to try to see how to specifically develop "the causality whose concept we use when we explain an action by its intention", as I cited Anscombe

putting it in the previous section. This, I think, involves a narrower category of causality than that which encompasses everything that we might want to invoke to show how reasons explanations in general are causal explanations. (In the next section, I'll consider what we should say about cases that fall outside the restrictions — instances of (2) and (3) where we can't take the agent to be acting with or through an intention.) So, the question then is how an explanation by intention of the form 'X is doing *A* in order to do *B*' or 'X is doing *A* because X is doing *B*' could provide causal information about the derivation of the effect mentioned in its explanandum, *that X is doing A*. As a starting point, we might consider our earlier case of the explanation "The tent fell over because the wind blew it over". There, the explanans names a cause and provides explicit information about how the effect described by the explanandum is brought about. On the other hand, "I went to the café in order to get some coffee" does not name a cause, and may not seem to tell us how something is brought about, given standard conceptions of causality. But that our sentence doesn't name a cause is not an inherent problem, since a causal explanation need not name a cause in its explanans.<sup>84</sup> The harder question is to see just what kind of causal information such a claim provides.

In light of the previous chapters, my answer should not be surprising. As Anscombe says in *Intention*, when an agent says that they are doing *A* with the intention of doing *B* or in order to do *B*, "the future state of affairs [given by the intention with which the action is done] must be such that we can understand the agent's thinking it will or may be brought about by the action about which he is being questioned" (1963a, p.34). A little more formally: "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 stage in proceedings of which P is an earlier stage" (1963a, p.35). This is again a formulation of what I've been calling the Belief Condition. Anscombe's remark here concerns acting with a future intention, but the requirement that we be able to see how the intentions with which an agent acts can intelligibly mark 'stages' in some

---

<sup>84</sup> See Davis (2005), for example.

‘proceedings’ applies just as well to the further intentions operative in what one is doing now. The thought here is that what such an instrumental or calculative thought conveys is information about a causal process. For if the information conveyed tells us that the agent’s doing Q is a stage of a proceeding with P as an earlier stage; but this is itself causal information about P, and about the causal proceeding in which P is involved.

It is important here to keep in mind how generous a conception of causal information is intended to be allowed for as to what could count as a ‘stage’ in a ‘proceeding’. Anscombe (1963a, p.79) says that practical reasoning involves calculating how to obtain something which is at some distance from one, but that things may be at a distance in different ways: in time, in space, or by being given by a ‘wider’ description of what one is doing. What most readily comes to mind if we think of causal information that would be given by an instrumental belief is of the sort that describes a causal relation between events: “I’m pressing the button in order to set off the bomb”. But if I say, “I’m walking down the street because I’m walking to school”, we don’t convey the causal information that the event of my walking down the road causes the event of my walking to school. Similarly, if I’m asked why I’m lying down and I say “I’m resting” — Anscombe’s example of a giving a ‘wider’ description — then I don’t convey causal information that fits an event-causal model. But in all these cases I take it that we have explanations by intention that express causal understanding, as they involve an idea or belief about a way of doing things. On this view, what makes an idea or belief of a ‘way’ of doing things count as consisting in *bona fide* causal information is determined by whether it would count as such from the point of view of the minimal causalism which we started out from. If it is intelligible that the information involved tells us how the agent’s activity ‘brought about’, ‘led to’, ‘resulted in’, or ‘effected’ certain matters, or if it makes sense to consider what could have ‘prevented’ or ‘interfered with’ it, then we are dealing with causal information; if not, not. And the same goes for how we should understand the slightly more reflective idiom that an explanation by intention tells us something about the ‘stages’ in a ‘proceeding’ or the ‘parts’

and ‘phrases’ of a ‘process’ — our grasp of all of these comes from within the minimal causalist picture, or is at least constrained by it.

That isn’t to say that there may not be cases where it could be hard to see how an explanation by intention provides causal information. Consider, say, someone saying “I’m breaking eggs because I’m making an omelette”. Ordinarily, in offering this explanation we would take the agent to think that breaking eggs is a way to make an omelette, a part or a phase of that particular causal process. But what if someone says “I’m breaking eggs because I’m making an omelette” even when they’re making an egg-free omelette but are cracking some eggs in sacrifice to the omelette god to make sure all goes well?<sup>85</sup> In this kind of case, the explanation needs “filling in”, as Anscombe (1963a, p.35) puts it. If we have no idea about this person’s weird omelette god views, then we’ll understand the explanation in the normal way, but if we’re not sure what he means — look, he’s breaking the eggs into the garbage — we have to ask for some filling in, and so we’ll get more information like “I’m breaking the eggs because I’m trying to appease the omelette god”, and “I’m trying to appease the omelette god because I need to do that to make my egg-free omelette”, and again we can see how such explanations could at least purport to provide causal information. (Though, of course, since we don’t believe in omelette gods, we can’t take either of these explanations as conveying information about *echt* causal processes.)

But why must we say that the explanation of an action by intention involves causal information of *that* particular kind? Why not say that information about the intention’s role as an antecedent cause is what counts or perhaps of some other sort? In fact, we can even wonder whether the sort of information suggested really does count as a form of causal explanation at all. For, as was suggested earlier, the question of whether reasons explanations are causal explanations is often glossed as the question of whether such explanations answer questions like ‘Why did *X*’s *A*-ing occur?’ That is the sort of question that the anti-causalist thinks a reasons explanation doesn’t answer. But the causal information that our

---

<sup>85</sup> Thanks to Kieran Setiya for this example.

account delivers doesn't seem to answer that question either. If the causal information imparted tells us something to the effect that doing *A* is *X*'s way of doing *B*, the question we seem to answer is not "Why is *X* doing *A*?" but rather "How is *X* doing *B*?". The account seems to speak to a different explanandum.

The answer to this objection is part of the justification for why the sort of causal information given in a calculative belief is central. For although an agent can surely offer different answers in response to the questions "Why is *X* doing *A*?" and "How did *X* do *B*?", even when doing *A* is *X*'s means to doing *B*, what Anscombe's account of the question 'Why?' brings out is that the question 'Why?' is in fact the dual or a mirror image of the question 'How?', in these cases. Let us recall Anscombe's case of the man pumping poisoned water into the house of party chiefs. In this example, the question 'Why?' can be used to extract a particular series of descriptions that are true of the man by the repeated application of the question to the answers elicited from him:

Why are you (*A*) moving your arm up and down?

— I'm (*B*) operating the pump.

Why are you (*B*) operating the pump?

— I'm (*C*) replenishing the water supply.

Why are you (*C*) replenishing the water supply?

— I'm (*D*) poisoning the inhabitants.

In this series — not necessarily the only one of its sort to characterize what the man is doing — each of the first three descriptions is related to the one that follows it as means to end. But when descriptions stand in such a series, it is also the case that they "can make an appearance in answer to the question 'How?'" (1963a, p.46). Thus, the last three descriptions in the series equally stand as ends to means of their preceding description. So we can illustrate the series just as well like this:

How are you (*D*) poisoning the inhabitants?

— I'm (*C*) replenishing the water supply.

How are you (*C*) replenishing the water supply?

— I'm (*B*) operating the pump.

How are you (*B*) operating the pump?

— I'm (*A*) moving my arm up and down.

The point is not to illustrate something that must occur in speech or thought, but rather to bring out “an order which is there whenever actions are done with intentions” (1963a, p.80). The different why and how-questions, with their different attendant answers, are both ways of getting at that order. But what they get at is a causal order: what is here illustrated is the form of the causal structure that exists *within* an agent’s intentional activity. This is the causal information that gets conveyed by an explanation by intention.<sup>86</sup>

## 7. OTHER FORMS OF CAUSAL EXPLANATION

But what about reasons explanations which don’t involve the agent acting with a further intention of the kind that we could express by saying “I’m doing *A* in order to do *B*” or “Because I’m doing *B*”? Consider, for instance, what Rosalind Hursthouse (1991) called “arational” actions, actions that cannot be explained by the agent’s having a belief about some further end, but rather by citing the agent’s experiences or emotions, such as when I run away because I was frightened, or when I shout at my

---

<sup>86</sup> For discussion and different perspectives about how we might see causality operative ‘in’ an action, or in a bit of intentional activity, see, Hornsby (2011; 2012; 2015), Lavin (2016), Thompson (2008, Ch.7) and Stout (2007; 2010a).

television set when the news makes me mad. Or consider Anscombe's own discussion of what she calls "motive-in-general" and "backward-looking motive", as opposed to "forward-looking motive", or explanation by intention. The first sort of motive may be expressed by explanations of action like "I did it out of love" or "I did it in admiration of her". According to Anscombe, such answers may *appear* to felicitously describe the agent as acting "in order to" bring about some state of affairs, but in fact only express the *spirit* in which the action is undertaken. As she puts it, such explanations only "interpret" our actions, or "put them in a certain light" (1963a, p.21). We find the second sort of motive at work in explanations when we invoke (say) revenge: 'I killed him because he killed my brother' (1963a, p.20). Here too, Anscombe thinks, there is no calculative thought involved (1963a, p.65). Can a causal pluralist offer a story about such cases, or offer a story about reasons explanations in general which encompasses all these different kinds of explanations?

Of course, such phenomena (and the discussion of them) deserve much more careful discussion. I won't undertake that here.<sup>87</sup> Instead, I'll offer only a schematic response for how a causal pluralist might approach such cases. First, I'm not inclined to say on the causal pluralist's behalf that we must try to force such cases into the mold of the previous section, to (say) bring out how arational actions or interpretative or backward-looking motives *really* do involve a calculative thought.<sup>88</sup> We can grant that such explanations don't involve "the causality whose concept we use when we explain an action by its intention". For not all reasons explanations have to be causal explanations in *that* sense. Does it cede ground to a competing causalist theory — one that (say) seeks to explain intentional actions by invoking antecedent causes — to say as much? I don't think so. On the account of causal explanation I've proposed, there is no in principle problem with, say, construing explanations by emotion or by motive as

---

<sup>87</sup> Hursthouse's paper is much discussed; for the more neglected case of Anscombe's notion of non-forward-looking motive, see Müller (2011) and, at a greater remove, Teichmann (2013).

<sup>88</sup> For some cautionary remarks on this score, see Anscombe (1981c, p.viii).

forms of causal explanations. Perhaps they are worth singling out as distinctive forms of causal explanation — perhaps they even involve kinds of causes that should be distinguished<sup>89</sup> — and perhaps they can be grouped together under a common genus of causal explanation which is narrower than the account from which we started, but wider than our account of explanation by intention. In any case, for the causal pluralist, such possibilities only present an invitation for further inquiry, and it hard to see why we could not devise positive answers to these questions, given the causal pluralist’s capacious picture of causal explanation. These are the sort of questions for which the causal pluralist story is made for, as it provides a framework within which we can develop alternative causalist stories about subject matters whose distinctive features cry out for distinctive treatment.

## 8. THE EXPLANATORY LIMITS OF CAUSALISM

Let us go back to the start. Recall that the challenge that Davidson set was to show how our account of reasons explanations secured that the agent’s reasons were causally at work in her action. What I’ve tried to do here is show how we can extract from Anscombe’s work an account of causality and causal explanation that comports with that minimal causalist claim, at least when we focus on the central phenomenon of explanation by intention. For the picture of causal explanation here sketched meets Davidson’s challenge, because the account does not allow for the possibility that a reasons explanation might work solely by dint of the agent possessing the reasons in question rather than her reasons being ‘operative’ in her action, albeit not in the fashion envisioned by Davidson. At the same time, if we can

---

<sup>89</sup> Compare Anscombe (1978, p.106) on mental causes: “It would be of interest to discuss the causality – i.e. how many different types there are here”. And notice too that although Anscombe (1963a, p.20) says that motive-in-general does not involve causal explanation of one sort (the sort most in the air at the time she wrote *Intention*), she also says that such motives “influence” the agent’s actions — more minimalist causalist idiom — where this contrasts with having a “ground” (a reason) for action. Though for a different interpretation, see Hyman (2013, pp.97-8).



buy into Anscombe's approach to causality and causal explanation, then it may not be quite right to say, as Davidson (1963, p.10) does, that it must be the case that reasons explanations "explain an event [...] by placing it in the context of its cause". But we can still say that reasons explanations do their work by placing an action in the context of its causality.

At the same time, the particular causalist view I've developed is quite remote from the more familiar forms of causalism in the literature. Of course, there are many different forms of causalism. Some causalists want to give an account of intentional or voluntary actions; others want to explain what it means to act for reasons; or what it means to act in a 'full-blooded' way. And they don't all agree on what it means to give an account: some want *a priori* analyses; others will settle for metaphysical definitions; others will be happy with necessary conditions. And they don't agree on what materials go into the account: some settle for beliefs and desires; others think intentions are important; others want more psychological items yet.<sup>90</sup> What binds the different forms of causalist together is the idea that the concept of causality is, in some way, an important explanatory concept. Indeed, a causalist may want to argue that principles like *AP* or the *Belief De Actione Principle* or other that seem to offer a necessary connection between intentional action (or acting for a reason) and the presence of various psychological states in the agent should themselves be explained by the efficient causal role that such states play in bringing about the action.

From the outlook of the causal story offered here, any such project must look very suspicious. For if the story I've told is plausible, then we don't really have any independent grasp on the concept of causality that is relevant to a domain from outside the set of concepts that characterize the domain itself. The point isn't that what causalists say will be necessarily false; it is rather that it could not be illuminatingly true. There can't, for instance, be much point in saying that *X* is an intentional action if it

---

<sup>90</sup> See Davidson (1963), Bishop (1989), Davis (2005), Velleman (2007), Bratman (1999, 2009), Setiya (2009), Smith (2010; 2012), Mele (2013), Hyman (2014).

caused in a certain way if the concept of causality involved, for it to be relevant to the subject matter, could only be one that can be grasped through concepts like *marrying* and *carrying a book*. But this, according to Anscombe, is how things are with intentional action. The causality involved in intentional action is that displayed by the *A–D* order brought out by the series of ‘Why?’ or ‘How?’ questions, and our primary knowledge of that type of causality is given to us through our grasp of intention-dependent concepts like *walking*, *marrying*, *holding a cup*, *playing guitar*, *turning on a light*, and so on. Any attempt to illuminate what it is for something to be an action of one of these types or to perform one of them for reasons by deploying the concept of *cause* would then seem to either invoke a causal concept that isn’t relevant to the subject matter, or it would invoke one which is, but whose grasp presupposes an understanding of the target concept.

This is not yet to *refute* these more familiar causalist projects. In the first place, I haven’t tried to defend causal pluralism against objections, but only tried to tease out some central elements of what this neglected position may involve. (I do think it can be defended, but that would require a dissertation of its own.<sup>91</sup>) Nor have I tried to show that the generous picture of causal explanation offered here is compulsory. Again, I’ve only sketched the view without defending it against all and sundry. Nor, for that matter, is it clear just when circularity in definition is vicious, so a causalist may be able to accept my account and lean on the possibility of benign circularity in definitions or theories to get some kind of explanation going.<sup>92</sup> But my goal here hasn’t been refutation. The aim here has only been to develop an

---

<sup>91</sup> Peter Godfrey-Smith (2008, p.334) says he is unaware of any explicit defense of Anscombe’s position (as he reads it), which is not quite true — I cited a few at outset (see especially Bogen 2008 and Psillos 2009 for very helpful defenses). A defense of the view could start by mining a good deal of the literature on ‘singular’ causation and related positions (much of which Anscombe inspired) to shore up parts of the causal pluralist position: e.g. Cartwright (2004), Stein (2016), Carroll (2009), Moore (2009) and Rota (2009) are all helpful here. Along with positive arguments in favor of the view, we could supplement the case for it by rehearsing the familiar barrage of objections to the extant alternatives.

<sup>92</sup> For one important account in this respect see Humberstone (1997) and Burgess (2008); compare Ford (2011) for a different view.

account that allows us to see that the explanatory aims of the causalists need not be ones that we have to buy into, since we can see how reasons explanations can be causal explanations without going so far, and we have done so with a framework that can explain precisely why those projects haven't and can't succeed.

#### IV. ANSCOMBE ON THE NATURE OF PRACTICAL KNOWLEDGE

“Action does require that what the agent does is intentional under some description, and this in turn requires, I think, that what the agent does is known to him under some description.”

— Davidson, ‘Agency’, p.50.

“[A]n act does not merely have many descriptions, under some of which it is indeed not intentional: it has several under which it is intentional.”

— Anscombe, ‘Action, Intention, and “Double Effect”’, p.223.

##### 1. INTRODUCTION

Having defended the claim that there is a necessary connection between intentional action and knowledge against potential counterexamples, offered my own account of why the claim is true, and presented an account of the causality involved in reasons explanations which undercuts some more familiar causalist approaches to this claim, I want to end by considering the character of the knowledge in question, or of “practical knowledge”. I’ll do this by wading into a recent interpretative issue about how Anscombe understands this notion. Although I will mostly be concerned with getting Anscombe’s account properly in view, I think her story about the nature of practical knowledge is one we should adopt, and we will see that the account I’ve offered in the previous chapters dovetails with it, if, that is, my interpretation of Anscombe’s position is correct. So, let us then start by recalling her provocative introduction of the term in *Intention*:

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 the facts. The facts, reality, are prior, and dictate what is to be said, if it is knowledge. (Anscombe 1963a, p.57)

Toward the end of the book, she completes the thought:

Practical knowledge [in Aquinas's phrase] is the 'cause of what it understands', unlike 'speculative' knowledge which is 'derived from the objects known'. (Anscombe 1963a, pp. 87)

This all sounds very exciting, if also a bit dark. But one point upon which all of Anscombe's readers seem to have agreed is that in this last quotation from §48 Anscombe is *defining* practical knowledge. Given this common assumption, the natural interpretation is that Anscombe's view is that it is the *causality* of practical knowledge — i.e. that practical knowledge is 'the cause of what it understands' — that captures what is essential to it.<sup>93</sup> Taking for granted that this is Anscombe's definition of practical knowledge, the discussion has largely focused on trying to understand what Anscombe's Thomistic thesis means, and whether there are any good reasons to accept it.<sup>94</sup>

---

<sup>93</sup> For this standard way of reading Anscombe, see, among many others, Moran (2004, p.47), Newstead (2009, p.158), Paul (2009, p.1), McDowell (2011a, p.142), Rödl (2011, p. 211), Stoutland (2011, p.32), and Schwenkler (2015, p.13).

<sup>94</sup> Not everyone thinks that Anscombe's story about practical knowledge is a coherent one. Much of the early reception of *Intention* seems to have been skeptical on this score: see, among others, Jones (1961), Martin (1971), and Danto (1976). Over the course of the last few years, a growing literature has emerged which gives Anscombe a much more patient and sympathetic reception: see, for instance, most of the references in the previous footnote, and also the other essays collected in Ford *et al.* (2011).

However, in his recent paper ‘Anscombe on Practical Knowledge’, Kieran Setiya (2017) argues that the discussion rests on a mistake. He argues that if we look more closely at the assumptions which Anscombe makes in the lead up to the above remark from §48, we see that Anscombe *restricts* the identity of practical knowledge with its causality to actions falling under certain kinds of descriptions. On this unconventional interpretation, Anscombe believes there are two *kinds* of practical knowledge: one kind is the cause of what it understands, the other is not. If Setiya is right, then Anscombe’s story about practical knowledge is more complicated than it first appears.

I want to consider Setiya’s account in detail, offer some criticism of it, and present an alternative story about what Anscombe is up to in the relevant passage. As in Setiya’s paper, the discussion will be largely exegetical: I’ll be concerned with what Anscombe thinks and why she thinks it, and not so much with whether she ought to think it. I will start in §2 with some necessary stage-setting concerning Anscombe’s talk about the ‘form’ of descriptions and action. In §3, I’ll sketch Setiya’s interpretation of the relevant passage from §48 of *Intention* and raise some problems for Setiya in §4 and §6. I will suggest that it will be important for understanding Anscombe’s argument in §48 if we take up the underdiscussed topic of her account of *voluntary* actions. Having Anscombe’s account of voluntary actions properly in view will be important both for seeing what is wrong with Setiya’s interpretation and for my own alternative interpretation, which I’ll weave together over §§5–7. The upshot, in §8, will be a novel account of what Anscombe means when she says that practical knowledge is the cause of what it understands, such that this involves a particular kind of dependence between practical thought and reality. An important feature of the interpretation I offer, developed further in §9, will be to take seriously the possibility that for Anscombe acting with intentions may not just be the central phenomenon in action but may *constitute* the subject matter.

## 2. THE FORM OF DESCRIPTIONS OF INTENTIONAL ACTIONS

It will be helpful to begin with a few very general remarks about Anscombe's project in *Intention*. A key claim of the book is that 'intentional' does not denote a property or feature of events or happenings which would otherwise be the same in its absence. In §19, Anscombe argues that in describing something as 'intentional' we "do not mention *any* extra feature attaching to an action at the time it is done" (1963a, p.28). And in §47 she argues that it is a mistake to think that "events which are characterizable as intentional or unintentional are a certain natural class, 'intentional' being an extra property which a philosopher must try to describe" (1963a, p.84). As mentioned in the previous chapter, one such extra property would be *being caused by such-and-such a mental state*, a property that could be applied to an event such as a bodily movement that could be identified independently of whether the movement is an intentional action or not. But this is just one such property: Anscombe's position appears to be that there is *no* such feature or property of *any* kind.

Her arguments in support of these negative theses are notoriously difficult, and I won't take them up here.<sup>95</sup> More important for us is her positive alternative, that 'intentional' refers to "a *form* of description of events" (1963a, p.84). What does this mean? I assume that Anscombe's point in introducing this Aristotelian talk of 'the form' of something is to mark a basic logical distinction, one not given to us through empirical inquiry but rather garnered through some type of insight or understanding that we possess by dint of being competent users of words or concepts that express the form, and the wider language or practice presupposed by such competence.<sup>96</sup> When Anscombe's delineates the class of

---

<sup>95</sup> See Hursthouse (2000) and Vogler (2016) for detailed exposition of these arguments, and also Ford (2011; 2015) for useful general orientation, at a wider remove from Anscombe's text. For some very early criticism of Anscombe's arguments on this score, see Dennett (1968).

<sup>96</sup> I presume too that Anscombe's talk of 'forms' is not unrelated to her talk of 'categorical' and 'grammatical' distinctions. In her (1981d), Anscombe discusses the affinities she sees between Wittgenstein's conception of the 'grammatical' and

intentional actions as “those actions to which a certain sense of the question ‘Why?’ is given application, [...] in which the answer, if positive, gives a reason for acting” (1963a, p.9) we can read this as an elucidation of the form of descriptions of intentional actions. As she puts it, what is “essential to this form is displayed by the results of our enquiries into the question ‘Why?’” (1963a, p.84). As we saw in the previous chapter, what those results bring out is that when this question is given repeated application, the ensuing answers do not describe actions that hang together as a mere aggregate of events but characterize them as exemplifying a certain explanatory order, Anscombe’s “*A–D* order”.

Although the order displayed by this procedure may be instanced in particular mental episodes of conscious practical reasoning, it need not be. For Anscombe, the *A–D* order is not itself a series of signs or a psychological process, but a logical feature of certain events, something “which is there whenever actions are done with intentions” (1963a, p.80). Of course, this is still far from perfectly clear; but I mean these very general remarks only to help us see that when Anscombe refers to the form of the description of an intentional action, what she is referring to is the *A–D* explanatory order in which we place the action when we grasp it as intentional, and this characterizes the form of the action itself.

### 3. SETIYA’S READING OF ANSCOMBE’S ARGUMENT

Let us now turn to Anscombe’s argument and Setiya’s interpretation of it. Anscombe’s argument occurs in §48 of *Intention*, and we’ve already cited the conclusion, that practical knowledge is the cause of what it

---

what she calls the ‘Platonic-Aristotelian tradition’ on the topic of ‘grammatical’ categories. What Alan Donagan (1987, p.iv) said of Ryle’s *The Concept of Mind* seems apt as a description of Anscombe’s work too: “it is a classic of modern analytic Aristotelianism, dressed in a behaviourism now out of style”. See Thompson (2008) for an important discussion of what he calls ‘analytic Aristotelianism’, the sort of method which I’m vaguely imputing to Anscombe. And see Pouivet (2008) for a more free-ranging discussion that sees a tradition running from Aquinas to Anscombe (and Peter Geach and Anthony Kenny).



understands. Before we consider the full text of the argument from §48, we should say a little about its immediate background. By the start of the preceding §47, Anscombe takes herself to have shown that “there are many descriptions of happenings which are directly dependent upon our possessing the *form* of description of intentional actions” (1963a, p.84). By ‘dependent’ she means that the descriptions are “dependent on the existence of this form for their own sense” (1963a, p.85). A typical example is the description ‘marrying’. But she thinks even some descriptions of kinds of actions which can be unintentional — like ‘going in reverse’ — are also dependent on the existence of the form, as they could not have been used to describe unintentional happenings if they were not also used to describe intentional ones. On the other hand, there are other kinds of descriptions which she says can occur “in the form” of descriptions of intentional actions, but which are not so dependent. Her example here is ‘sliding on ice’. Setting aside descriptions of this latter sort, Anscombe divides the class of descriptions which are dependent on the form of descriptions of intentional action into two columns, which I’ll (partially) reproduce here:

Intruding	Telephoning
Offending	Calling
Coming to possess	Groping
Abandoning	Crouching
Dropping (transitive)	Greeting
Abandoning, leaving alone	Signing, signaling
Switching (on, off)	Paying, selling, buying
Placing, arranging	Marrying, contracting

According to Anscombe, the descriptions in the left-hand column could be used to characterize intentional or unintentional happenings, while those in the right-hand column could only apply to voluntary or intentional events (except that she thinks a few could apply to things one does in one’s sleep).

Now §48, which includes the passage on which we'll focus, starts like this: "We can now see that a great many of our descriptions of events effected by human beings are *formally* descriptions of executed intentions" (1963a, p.87). Anscombe then offers her argument for the Thomistic slogan we cited at the outset:

[W]e can say that where (*a*) the description of an event is of a type to be formally the description of an executed intention (*b*) the event is actually the execution of an intention (by our criteria) then the account given by Aquinas of the nature of practical knowledge holds: Practical knowledge is 'the cause of what it understands', unlike 'speculative' knowledge, which 'is derived from the object known'. (Anscombe 1963a, p.87)

With this much of Anscombe's argument on the table, we can turn to Setiya's interpretation. First, Setiya assumes — and I take it everyone will agree — that Anscombe must think that if an event satisfies condition (*b*), then it is an executed intention. However, Setiya's interpretation really gets going with the thought that Anscombe must also believe that the content of the intention executed according to (*b*) is given by the description mentioned in (*a*). For otherwise, her argument would be invalid, as the content between the two conditions wouldn't necessarily match. But if that much is correct, then it appears as though Anscombe must believe that when an action satisfies condition (*b*), that does not itself entail that the agent's practical knowledge of the action is the cause of what it understands, because otherwise condition (*a*) would be redundant. So condition (*a*), on Setiya's account, doesn't apply to all action-descriptions that could characterize intentional actions, but must rather have a more restricted application. And Setiya suggests there is an interpretation of condition (*a*) that brings out that restriction. Because Anscombe starts §48 in the way just noted, by remarking that many of our descriptions of what we effect are "*formally* descriptions of executed intentions", and then goes on to say (1963a, p.87) that this is obviously true for descriptions of the type in the right hand column, Setiya thinks that Anscombe

takes condition (a) to refer to just *these* descriptions, the ones in the right-hand column.<sup>97</sup> If an agent is telephoning or crouching or paying or marrying, they can't but be doing so intentionally, and their practical knowledge is the cause of what it understands; if they are intentionally intruding or causing offence or dropping a book or sliding on ice (actions under descriptions from the left-hand column, or from neither), then they have practical knowledge of their doing that is not the cause of what it understands, so that the object of the knowledge could exist apart from it. If this is right, the common interpretation of Anscombe's argument is wrong, and the right interpretation of her view about how practical knowledge relates to the particular type of causality under discussion must be more complicated.

#### 4. SETIYA ON CONDITION (A)

But I think Setiya's account is wrong. Let us start with condition (a). Why should we say that Anscombe thinks this condition is only satisfied by descriptions from her right-hand column like 'marrying'? Setiya's thought seems to be that Anscombe believes that only descriptions which entail that the agent is acting intentionally and with practical knowledge of doings under such descriptions warrant the title 'descriptions of events which are formally descriptions of executed intentions'. As it happens, Setiya does think there are other descriptions which also entail that the agent is acting intentionally, and which Anscombe does not explicitly mention in §48: descriptions of the shape 'doing *A* in order to do *B*',

---

<sup>97</sup> This isn't Setiya's final position — he does think more descriptions can fall under (a) — but we set this aside until the next section.

‘doing  $A$  because  $p$ ’ and ‘doing  $A$  for its own sake, for no particular reason’. But many other descriptions of things that can be done intentionally will still not fall under condition (a).<sup>98</sup>

Now, it is right that for Anscombe the description of an event which is formally the description of executed intentions must describe an event that is done intentionally, and with practical knowledge (it is also something that any advocate of  $AP$  will likely say is right, whatever Anscombe’s views). But Setiya puts the idea like this: “where ‘doing  $A$ ’ is formally the description of an executed intention, one could not be doing  $A$  in the absence of practical knowledge” (Setiya 2017, p.158, *n.5*). And this is misleading when combined with Setiya’s claim that Anscombe takes the descriptions in the right-hand column *to be* formally descriptions of executed intentions (because she says they can only be intentional or voluntary, unlike the items in the left-hand column, which she thinks can be intentional or unintentional). For in condition (a) Anscombe does not delineate as a class ‘descriptions of events which *are* formally descriptions of executed intentions’. Rather, she delineates the class of ‘descriptions of events which *are of a type to be* formally the description of an executed intention’. The difference matters, because if we plug Anscombe’s exact words back into Setiya’s remarks — i.e. “where ‘doing  $A$ ’ is a description *of a type to be* formally the description of an executed intention, one could not be doing  $A$  in the absence of practical knowledge” — then we have, I argue, a claim that is false by Anscombe’s lights.<sup>99</sup>

Anscombe allows descriptions that satisfy condition (a) to characterize events that are not intentional in at least two ways: they can describe an action which is unintentional, and they can describe

---

<sup>98</sup> Setiya’s extension of condition (a) to include such cases is part of his own response to the puzzle. As he sees it, when practical knowledge isn’t the cause of what it understands, it still rests on knowledge which is, by dint of also falling into such schemas as ‘doing  $A$  in order to do  $B$ ’. If the account that follows is right, then we don’t need to follow Setiya here.

<sup>99</sup> Setiya (2017, p.159-60) acknowledges the possibility that he is misreading condition (a) in this way; but his response addresses a different position from the one I’ll etch out.

actions which are voluntary but not intentional.<sup>100</sup> The first case is, I think, the most obvious to see from the texts at hand, though it isn't easy to make the point without begging the question against Setiya at this point. The rather more natural reading of condition (a) is just that Anscombe means to include under the condition descriptions of the type that occur in the left-hand columns too, like 'dropping a book', which can also describe unintentional actions.

I will try to support this interpretation below. But first, let us consider the second possibility that descriptions in the right-hand column may describe actions which are voluntary but not intentional. Recall, to begin, that Anscombe describes the items in the right-hand column as intentional *or voluntary*. Now Anscombe goes on to say that every intentional action is also a voluntary action (1963a, p.90), but, importantly for us, she doesn't accept the converse claim. She says: "Something is voluntary though not intentional if it is the antecedently known concomitant result of one's intentional action, so that one could have prevented it if one would have given up the action; but it is not intentional" (1963a, p.89). The sort of case she has in mind here is illustrated by one iteration of her example of the man pumping poisoned water into the house. In the case where replenishing the house water supply is the man's ordinary job, and certain other conditions are in place, Anscombe thinks that the man can reasonably refuse to answer the question "Why are you pumping poisoned water into the house?"; even though "he knows concerning an intentional act of his — for it, namely replenishing the house water-supply, is intentional by our criteria — that it is *also* an act of replenishing the house water-supply with *poisoned* water, it would be incorrect, by our criteria, to say that his act of replenishing the house water supply

---

<sup>100</sup> As noted earlier, Anscombe also mentions that some of the items in her right-hand column could apply to things that we do in our sleep, but I think we should set this case aside. When used in this way, the descriptions take on a different sense, and I think are not in that sense descriptions that satisfy condition (a). Or, in the terminology I'll later adopt, when they have this use they are not 'human-action descriptions'. I think Anscombe suggests this much elsewhere, when she says, in the case of 'walking' that "I say that this is *normally* a human-action description to allow exclusion of sleepwalking" (1982, p.210).

with poisoned water was intentional” (1963a, p.42).<sup>101</sup> So descriptions like ‘poisoning people’ need not describe intentional actions, but merely voluntary ones.

This the only example of this kind that Anscombe provides in *Intention*, but it seems reasonable to surmise she may say same about other descriptions in the right-hand column too. In fact, if we look outside of *Intention*, then we even find Anscombe giving examples of voluntary but non-intentional actions of which the agent *lacks* knowledge. At the start of chapter 1, I cited two remarks from work contemporaneous with *Intention* wherein Anscombe claims that intentional action requires knowledge. In both places, though, she means to distinguish intentional action from voluntary action on that ground. In ‘Two Kinds of Error in Action’, she tells us that “one cannot intentionally, but can voluntarily, do something without knowing one is doing it” (1963, p.7, n.2; also p.8). And in ‘On Being in Good Faith’, she says that although if one is “charged with a certain intentional action it is an adequate defense that you did not know you were doing that”, in this respect “the limits of the voluntary are far wider than the limits of the intentional” (2008a, p.105). But this is not to say that Anscombe must think that every description in the right-hand column could describe actions done in ignorance. In the ‘Two Kinds’ paper, Anscombe suggests her position is more nuanced:

There are some descriptions ‘X’ of things done that cannot hold unless the subject knows he is doing X; for example, making a contract. There are some that can hold though the subject does not know he is doing X; for example signing a transfer. But there are, further, some interesting intermediate cases. (1963a, p.4)

The intermediate cases she mentions are descriptions like ‘murder’, ‘bigamy’ and ‘adultery’. Anscombe argues that such descriptions can characterize an agent’s doing as a voluntary action even when the agent doesn’t know that she is committing such an act, but they could and should have known, in the sense that the agent is culpable for their ignorance. The case of the pilot who sinks the ship in the previous

---

<sup>101</sup> See Hanser (1998) for a discussion of this case.

chapter is another example with this sort of structure: he ‘sinks the boat’, but he doesn’t do so intentionally, but does so voluntarily because he could and should have known that doing what he did (e.g. getting drunk and not paying attention to the lighthouse) *is* to sink the boat, or cause that to happen.<sup>102</sup>

Notice too that Anscombe here includes ‘signing a transfer’ as something which we could do without knowing that we are doing so, though ‘signing’ is in her right-hand column. That brings out a possible complaint about my argument so far. For might there not be a difference to be drawn here between ‘signing’ and ‘signing a transfer’, in that the latter, but not the former, is a description that can characterize an unintentional action, and so belongs in the left-hand column, not the right-hand one, as I just claimed? The same might be said about Anscombe’s understanding of descriptions like ‘murder’, ‘bigamy’ and ‘adultery’ which I’ve drawn from texts outside *Intention*. It is an assumption that Anscombe would say that these belong to the right-hand column. Why not say they belong in the left-hand column instead? I think it is implausible to say that Anscombe would not place descriptions like ‘poisoning someone’ in her right-hand column, as she clearly says the actions they describe can be intentional or voluntary. But the general point here is that my argument turns on introducing concepts of *voluntary* and *involuntary* action that are unexplained, and it would help the discussion to try to say more about them before we continue. That is the job of the next section.

## 5. HUMAN-ACTION DESCRIPTIONS

As it happens, Anscombe’s discussion of voluntary action has received very little attention. This is probably because most think there is little to attend to. Thus, in a recent attempt to resuscitate interest in

---

<sup>102</sup> To be clear, I’m giving a highly simplified and compressed presentation of Anscombe’s account of these sorts of cases.

the concept, John Hyman (2013, p.683) remarks that in Anscombe's view "voluntariness plays a relatively minor role in our thought about human action, compared to the concept of acting intentionally or acting for a reason, and does not raise any interesting problems of its own, once the nature of intentional action has been explained". The first half of that remark seems right; but the second half is too strong.<sup>103</sup> It also rests on a mistake: it is wrong to ask after Anscombe's account of voluntary action *simpliciter*. As we'll see, she seems rather to be working with different (though intersecting) *senses* or *uses* of 'voluntary' and 'involuntary' and related terms.<sup>104</sup> What I'll try to do here is single out the sense that I think is in play in her remark that items in the right-hand column can be intentional or voluntary, and say how this helps us grasp condition (a). What I'll suggest is that for Anscombe condition (a) refers to what she calls 'human-action descriptions' (1982, p.210), in a sense I'll try to make clear.

Let us start with some background to Anscombe's passage in §48. Recall that the section begins: "We can now see that a great many of our descriptions of events effected by human beings are *formally* descriptions of executed intentions" (1963a, p.87). That's what we've seen, presumably, in the immediately previous sections. Now §47 itself begins: "Thus there are many descriptions of happenings which are directly dependent on our possessing the *form* of descriptions of intentional actions" (1963a, p.84). Because of the 'thus', we are presumably supposed to have already seen in §46 why many descriptions of happenings are dependent on our possessing this form. Now when we turn to that

---

<sup>103</sup> It is true that voluntariness is not discussed explicitly at great length in *Intention*, but it does come up for extensive reflection in various other works — see Anscombe (1963c; 1982; 2008a; 2008b) — and the character of Anscombe's discussion is usually exploratory, lacking the finality we might expect if she found the topic philosophically thin. The only concerted discussion of this material which I'm aware of is by Roger Teichmann (2008; 2014), who brings out some intriguing views of Anscombe's concerning culpability, doxastic voluntarism and epistemic agency.

<sup>104</sup> Tenenbaum (2007, pp.92, 95) also misses this. To be clear, I don't think, and I don't mean the following discussion to suggest, that for Anscombe the relation between concepts of *intentional*, *unintentional*, *voluntary*, *involuntary* and *non-voluntary* actions is at all straightforward or clear cut. What follows is only meant to pick out what I think are some of the main threads in how Anscombe understands these concepts.



section, we find Anscombe concerned with the seemingly inscrutable question of *why* we describe certain happenings in the way we do — why we use descriptions like ‘building a house’ or ‘writing a sentence on a blackboard’ or ‘poisoning some people’ — as opposed to having descriptions for all sorts of other happenings and processes in the world, like the minute movements of the leaves of a tree being hit by a breeze. She says:

Of course we have a special interest in *human actions*; but what is it that we have a special interest in here? It is not that we have a special interest in the movement of these molecules — namely, the ones in a human being; or even in the movements of certain bodies — namely human ones. The description of what we are interested in is a type of description that would not exist if our question ‘Why?’ did not. It is not that certain things, namely the movements of humans, are for some undiscovered reason subject to the question ‘Why?’ ... [T]he description of something as a *human action* could not occur prior to the existence of the question ‘Why?’, simply as a kind of utterance by which we were then obscurely prompted to address the question. (1963a, p.83; my emphasis)

Anscombe is here using ‘human’ as a quasi technical term. That much is at least suggested by the fact she is, in some way, contrasting *the movement of a human* with a *human action*. (Are the former not just instances of the latter?) If we again look beyond *Intention* we can see that this is a technical term. In her unpublished later manuscript ‘Good and Bad Human Action’ she explains her terminology:

I am so using ‘human action’ that nothing is a human action unless it is a voluntary action on the part of a human agent. Otherwise, like digesting your food and breathing and sweating, your acts are the acts of a human agent but are not what I call human actions. (Anscombe 2005a, p.203)

Elsewhere, in ‘Action, Intention, and “Double Effect”’, she says more:

[There is] the category called ‘human action’ by some scholastics: they made a contrast between ‘human action’ — *actus humanus* — and ‘act of a human being’ — *actus hominis*. Again, this is usually explained by examples. Idly stroking one’s beard, or idly scratching one’s head, may be an ‘act of a human being’ without being a ‘human act’. And I expect falling over or tripping up is so too; at any rate I lay it down that I will use the term ‘act of a human being’ in such a way that those things come under that heading. Also, that ‘act of a human being’ is a wider notion, which includes ‘human action’. (Anscombe 1982, p.208)

Read in concert, the two remarks appear initially a little puzzling. The first remark tells us that human actions are necessarily voluntary, unlike the involuntary acts of a human such as digesting. Then, in the second remark, human actions are contrasted not just with acts of a human like falling and tripping over but also with idly scratching one’s head. Given that idly scratching one’s head is something we would probably want to be able to describe as a voluntary action, the natural thought is that when Anscombe says that human actions are necessarily voluntary, she is using a notion of ‘voluntary’ that can include actions like idly stroking one’s beard, so that mere acts of a human may be either voluntary or involuntary.<sup>105</sup> And that is just what she does go on to say:

We might say that human action = voluntary action. But that raises a question of meaning ... . We are speaking of voluntary action *not* in a merely physiological sense; not in the sense in which idly stroking your beard *is* a voluntary action. ... Aristotle says that beasts and babies have the voluntary, but we would not want to say so in the sense of ‘voluntary’ that we are trying to introduce. Aristotle too introduced a restricted sense of ‘action’ — *praxis*, which beasts and babies don’t have. But it is a bit too limited for us. It wouldn’t include omissions unless calculated, or sudden impulsive actions. (Anscombe 1982, p.208-9)

It seems clear that the widest category of action of activity that Anscombe countenances and that subsumes everything else is ‘act of a human agent’. Within that category, she distinguishes acts that are

---

<sup>105</sup> There is more to say about this sort of example of idle voluntary action. I’ll return to the topic in §9.

‘physiologically voluntary’ from acts that are ‘physiologically involuntary’, and then distinguishes the narrower category of voluntary acts that are ‘human actions’, which are necessarily physiologically voluntary. And we can also see here the claim from the previous section that the intentional is narrower than the voluntary, though now in this new jargon: “the extension of ‘human action’ is wider than that of ‘intentional human action’” (Anscombe 1982, p.213).

But how are we to grasp the differences between these concepts? Anscombe goes on to say that acts that are physiologically voluntary are “*under the command of reason*”, in that “reason *can* intervene to forbid” (1982, p.208). In yet another place, she spells this out further:

[W]hen we are using ‘voluntary’ to speak of human acts as having to be voluntary — this being part of what we mean by a ‘human act’ or ‘human action’ — we do not say that acts done under threat are not voluntary. An act is voluntary in the way we are speaking of here if it is in the power of the agent to do it or not — he can resist the threat, that is: he can refuse to do what he is being ‘made’ to do, if possible, by threat. This ‘can’ is first and foremost a matter of physical possibility. (Anscombe 2008b, p.127)

We might say that for Anscombe an act or action is physiologically voluntary when it is the act of a ‘two-way’ power to do or not do what is within the agent’s physical constraints.<sup>106</sup> Anscombe is here also flagging a sense of ‘involuntary’ in which acting under (say) a threat of violence would be ‘involuntary’ though still physiologically voluntary. But this sense of the word is of no interest to us here, and I set it aside.<sup>107</sup>

---

<sup>106</sup> See Kenny (1975, Ch.7) and Alvarez (2013) for discussion of the Aristotelian notions of one-way and two-way powers. Again, there is a lot more that could be said here.

<sup>107</sup> Anscombe discusses these kinds of cases in §49 of *Intention*: “From another point of view, however, such things [physiologically voluntary actions] can be called involuntary, if one regrets them very much, but feels ‘compelled’ to persist in the intentional actions in spite of that”; and “intentional actions can also be described as involuntary from

So we have some sense of how Anscombe distinguishes the physiologically voluntary from the physiologically involuntary. But what about the difference between the physiologically voluntary and the category of human action? And between the latter and intentional actions? I believe she would say that it is a bad question to ask how to distinguish the voluntary from the wider notion of the physiologically voluntary. As in the case of intentional action, Anscombe does not want to look for a *property* here, a feature that distinguishes human actions from the wider class of merely physiologically voluntary actions, or all acts of humans.<sup>108</sup> Indeed, as should become clear, the notion is too tied up with the intentional for this to really be a distinct project from searching after the ‘extra feature’ in that case too.

To distinguish intentional from human actions, on the other hand, may seem very easy: can’t we just use Anscombe’s definition of intentional actions? No: if the question is ‘given application’, that doesn’t tell us that the action is intentional but not voluntary, and if it is ‘refused application’ that doesn’t tell us that it is voluntary as opposed to unintentional or involuntary. (I’ll go on to suggest in §9 that the question ‘Why?’ does not, even when given application, cleanly delineate the narrower category of intentional actions from the wider notions of voluntary action at all.) Let us instead consider what Anscombe says about the kinds of descriptions involved in her discussion, rather than the kinds of events described. Anscombe, so far as I can tell, nowhere talks about the type of description ‘act of a human being’. (Such a classification would presumably not exclude much.) Nor does she talk about the type of descriptions of ‘physiologically voluntary action’. One might be tempted to say that ‘descriptions of bodily movements’ would fit here, but that can’t work, since a description of a human action *is* a description of a physiologically voluntary action, and these can go beyond mere bodily movement: e.g. ‘writing’. However — and this is the point I’ve been working up to — we have seen that in her remarks

---

another point of view, as when one regrets ‘having’ to do them. But ‘reluctant’ would be the more commonly used word” (1963a, p.89-90). See Audi (1986, p.93) for related discussion.

<sup>108</sup> See, in this respect, Anscombe (1982, p.209).

in §46 Anscombe *does* talk explicitly about the *type* of description ‘human action’. She is there concerned with the “type of description ... of something as a human action”, which, she says, “could not occur prior to the existence of the question ‘Why?’” (1963a, p.83).<sup>109</sup> And that makes it clear, I think, that Anscombe takes what she elsewhere explicitly calls ‘human-action descriptions’ to be those descriptions which are dependent on the form for their sense. And those descriptions are just the descriptions which can appear in both columns in §47. So we shouldn’t be surprised to see Anscombe describe the items in her right-hand column as possibly voluntary, or the left-hand column as possibly unintentional. For she is here working with a type of description — a human-action description — which can characterize such doings just as well as intentional ones.

What is distinctive of human-action description? Condition (*a*) tells us that these are descriptions that can (but need not) describe actions which are formally executed intentions. That isn’t wrong, but I would rather say that the sense of such descriptions is bound up with *A–D* order. The *potentiality* to represent such an order is inscribed into their content. That is, when we employ such descriptions, we characterize events in such a way that we are prompted (or can be prompted) to ask ‘Why?’ or ‘How?’ about the event at hand. Of course, we can ask ‘Why?’ and ‘How?’ questions about all manner of events; but the sense of the question ‘Why?’ or ‘How?’ at issue is that bound up with a calculative order of reasons, one which can characterize a distinctive kind of event. Anscombe has more to say about these descriptions; but let us first turn to Setiya’s account of condition (*b*).

## 6. SETIYA ON CONDITION (*B*)

If this interpretation of condition (*a*) is correct, then we can see why Anscombe obviously needs condition (*b*) to have an explanation for when and why practical knowledge is the cause of what it

---

<sup>109</sup> She explicitly uses the term ‘human-action description’ in ‘Action, Intention, and “Double Effect”’.

understands: condition (a) does not in itself choose between event which are intentional, unintentional, or voluntary but not intentional. But before I say more about my interpretation of condition (a), or consider objections to it, I want to also consider Setiya's treatment of condition (b). Recall Setiya's problem: "If practical knowledge were the cause of what it understands whenever one acts on one's intention, so that condition (b) is met, condition (a) would be redundant. But Anscombe includes it" (2017, p.159). In a footnote to this passage, he points out, rightly, that this argument assumes that "the description cited in (a), the one of which you have practical knowledge, figures in the intention cited in (b)." And he offers an argument in support of that assumption: "If 'doing A' figures in the content of the executed intention, as I [Setiya] assume, condition (a) will be redundant. If it does not figure in the content of that intention, Anscombe's argument is invalid" (Setiya 2017, pp.159-60).

But now let us recall condition (b): "the event [described by condition (a)] is actually the execution of an intention (by our criteria)". Again, we can note a small mismatch between Anscombe and Setiya's wording, but one which matters: Setiya talks of "*the* intention cited in (b)" and "*the* intention executed in (b)", but Anscombe only talks of an event which is the execution of *an* intention. Nothing about condition (b) itself licenses the definite article, and so Setiya's assumption that "the description cited in (a), the one of which you have practical knowledge, figures in the intention cited in (b)" may suffer from presupposition failure.

This is not a semantic quibble. It follows from one of the primary doctrines of *Intention* that we can't extract an intention which is *the* executed intention in an event merely by being told that the event is an executed intention. That, in one way, is the point of the famous example of the man pumping poisoned water into the house. His action is an executed intention by "our criteria" — i.e. he can answer the question 'Why?'<sup>110</sup> — so it falls under condition (b). If we ask him "Why are you moving your arm

---

<sup>110</sup> See Anscombe (1963a, pp.vi, 42) for her other uses of this 'criteria' locution to pick out the question 'Why?' and its full range of possible answers.

up and down?” he can satisfy the criteria by offering a positive answer that provides his reason for doing so, perhaps by saying that he is moving his arm up and down to pump water through the pipes. But, for Anscombe, there is no answer to the question of what *the* intention is in the event that is his act. *One* intention may be given by that reason, to pump water through the pipes. *Other* intentions in the *same* act could be given by further answers to the question ‘Why?’ about the ensuing action: he is pumping water to fill the cistern with poisoned water, doing that to poison the people in the house, doing that to get paid, etc. It is the *same* event which is an executed intention no matter which description or intention you use to single it out.

It is true that in her discussion of this example Anscombe does say that we can talk of *one* intention in an act if we focus on the intention *with which* the intention is done — the intention which, as she puts it, “swallows up” the other intentions in the series that precede it. In the case as I just described it, that would be the intention to get paid. But how do we identify *that* intention as *the* intention with which the act is done, in this technical sense? We can say, as Anscombe does, that this intention (say, to *D*) is the last one in this series. But the problem is that the point of the example is also to bring out that the one act may also contain *many such series*. As well as pumping water to get paid, the agent may also, to use Anscombe’s example, be pumping water to beat out the rhythm of God Save the Queen (1963a, p.47). So the definite description ‘the intention with which the act is done’ may also be unwarranted. For every *ABCD* series in an event, there might be an *ABCE* series existing alongside it.

So it not surprising that Anscombe (1979b, p.193) says that there “is a mistake in speaking of ‘*the* intention’ in an act”. To talk of ‘the intention’ involved in Anscombe’s argument must then involve an explicit appeal to the description given by condition (*a*). But for just that reason, we can see why condition (*b*) is itself insufficient for practical knowledge, because we need to have condition (*a*) in view to provide us with the description which in turn provides us with *the* intention under which the agent has practical knowledge, something that condition (*b*) can’t do on its own.

What about Setiya's dilemma? We've seen that we can skirt the first horn because it involves illicit talk of 'the' executed intention. Are we then impaled on the second horn, that "If [the description cited in (a)] does not figure in the content of that intention, Anscombe's argument is invalid"? No, because we can still say it does figure in condition (b), thus securing the validity of Anscombe's account. All we need to do is tweak the wording of Setiya's assumption to obtain a claim that Anscombe can accept: "the description cited in (a), the one of which you have practical knowledge, figures in the content of *one of the intentions* in the event mentioned in (b)."

## 7. SLIDING ON THE ICE LIKE A BOSS

Let me now turn to a potential objection to my interpretation of condition (a). If Anscombe construes condition (a) in the manner I've suggested, then an obvious potential problem arises in the case of descriptions like 'sliding on ice', which we earlier saw Anscombe does not take to be dependent upon the form of intentional actions. That is, in the language I've introduced, such descriptions are not human-action descriptions. But Anscombe also says that such descriptions can appear "in" the form, by which I take it that she means an agent can use such descriptions in explanations of her actions, as in "I slid on the ice because I felt cheerful" (1963a, p.85). And if the agent can offer such explanations, then it seems natural to say that by Anscombe's lights the description functions as formally a description of an executed intention, and so denotes an object of practical knowledge. Since my account seems to exclude such descriptions from the scope of condition (a), the reader may wonder why my story is not just like Setiya's in bifurcating practical knowledge into two sorts, with the only difference perhaps lying in a more generous sense of when practical knowledge is the cause of what it understands.

But I don't think the case of 'sliding on ice' poses a serious problem for my interpretation. Let us begin by considering what *type* of description 'sliding on ice' is. It is not, I've said, a human-action



description. I think Anscombe would say that it is what she elsewhere calls a “purely physical description”. For instance, she gives examples of certain physiologically involuntary bodily movements given by purely physical descriptions such as “the odd sort of jerk or jump that one’s whole body sometimes gives when one is falling asleep”, as well as “tics, reflex kicks from the knee, [and] the lift of the arm from one’s side after one has leaned heavily with it up against a wall” (1963a, p.13). These descriptions all contrast with descriptions of human-action descriptions in that the latter go “beyond physics”; they are all “at least basically animal”, in that they either characterize movements “with a normal role in the sensitive, and therefore appetitive, life of animals”, like ‘kicking’, or “they suggest backgrounds in which characteristic things are done” in the lives of animals, like ‘intruding’ (1963a, p.86). Purely physical descriptions are not limited to just these limb movements: another of Anscombe’s examples is ‘tremble’, which she thinks cannot characterize a voluntary action (1963a, p.33). And we can surely talk of purely physical descriptions which don’t inherently involve acts of a human at all. Any textbook on geology or metallurgy or physics would supply us with plenty of examples, but let us stick with Anscombe’s own examples, which can characterize acts of a human: ‘falling through the air’, ‘floating in water’, ‘rolling down a hill’, ‘sliding down a hill’, and, of course, ‘sliding on ice’. These latter descriptions could also characterize the acts of animals, and inanimate objects too.

Now what Anscombe says about ‘sliding on ice’ shows that she is happy to say that at least some purely physical descriptions can occur in the form of description of executed intentions. So I think Setiya is right when he describes Anscombe’s position in the following way:

Not all descriptions of intentional action belong on the list: some are independent of intention. Anscombe gives the example of ‘sliding on ice’, though even this appears in compound descriptions — ‘I slid on the ice because I felt cheerful’, where this gives my reason for acting — that imply intentional action (Anscombe 1963: 85). (Setiya 2017, p.158)

That is, the fact that purely physical descriptions can occur in the form of executed intentions means two things: first, that we can use such descriptions in our ordinary explanations of our actions, and that when we do so we are committed to viewing the explained action as an intentional action. However, these two claims don't together entail that we must take the intentional action to be intentional (and practically known) under *that* description. This may sound like a clean break with Anscombe's view in *Intention*, but, so far as I can see, nothing she says requires us to accept the latter claim. It requires, I think, a deference to ordinary language of the sort that Anscombe would eschew to think that we must say that the verb that we happen to use in our explanation must itself properly capture the intention under which we act and the object that characterizes our practical knowledge. In any particular case, the only real pressures to think we must say that would arise if it were not possible to paraphrase the agent's action or knowledge in another way which the agent would be happy to accept as the right description of what is going on. And this, I think, is just what Anscombe will say is possible about a case like 'sliding on ice'.

To see what I have in mind here, consider the following passage from Anscombe's later paper 'Good and Bad Action', where she spells out her position on this score (Anscombe has just remarked that it is a necessary feature of human actions that they be physiologically voluntary):

If, for example, you lie floating in water and you move your hands above and below you, and then repeat this the other way round, you turn over towards one side and the other in a way doubtless to be explained by Newtonian mechanics; but this motion of yours is not a human action, just the act of a body which happens to be a human being. So though it may well be a human act to *cause* it, and then you do voluntarily bring it about, it is not itself a human action. Of course you may deliberately kill someone by bringing about such an act of the human being that you are, as when you deliberately make yourself fall off a branch of a tree on to someone's neck; but that does not make the falling itself a human action any more than the falling of a log would be, which you deliberately brought about for the same purpose. The turn of your body is voluntary on your part; so is the falling of the log, if you voluntarily brought them about; but they are not themselves voluntary actions of yours; only

the bringing about of them was that. Still, their happening can be called your doing, as when we say: 'Look at what resulted! That was your doing'. (Anscombe, 2005a, p.203)

Anscombe is describing, in part, physiologically involuntary acts. Once you're falling through the air, reason "can't intervene", though it might have stopped you falling in the first place. But there is nothing wrong with saying things like "I fell through the air because I wanted to land on his head". And there is an intentional (or voluntary) action in the offing here, by dint of such an explanation: it is *causing myself to fall onto his head*. But the existence of that action, Anscombe suggests, does not mean we have to say that the action is voluntary under the description 'falling through the air', even if we happen to employ those words to explain the action. Although Anscombe doesn't talk here of practical knowledge, I think she would also deny that the action under such a description is an object of practical knowledge.

We can offer a similar story about her example of sliding on the ice. Notice that she describes the example in the 'because' form, though not in the shape of a naïve rationalization, but rather as offering a mental cause (a feeling) involved in the etiology of the action. And here too we can say that the explanation entails the existence of a voluntary or intentional action, one that can be an object of practical knowledge, but it is better denoted by a description like 'my causing myself to slide on the ice' rather than 'my sliding on the ice'. We could even say it like that, stilted as it would sound: "I caused myself to slide on the ice because I felt cheerful". Of course, putting 'sliding on ice' in the same basket as 'falling through the air' may seem wrong in so far as we're inclined to say that sliding on ice *is* something we can do intentionally, under that description. (For every YouTube video of some poor cat slipping and sliding away on the ice, there's another with a title like "Man Slides on the Ice Like a Boss".) Perhaps Anscombe wasn't very good at sliding on ice. In any case, if what I've said about it seems too off key, we might say that Anscombe is just wrong to think that this description is not itself dependent upon the form for its meaning. But the fact that we can readily offer alternative descriptions of the object of intention and knowledge when purely physical descriptions are employed in explanations of human action ought to make us wary of placing too much weight on that example. Take 'trembling' as another

example. As just noted, Anscombe thinks this can't describe a voluntary action, and so can't describe an intentional one. But there doesn't seem anything wrong with saying things like "I trembled because I knew everyone at the party would get a laugh out of it". (Perhaps I've taught myself how to do this, by recalling some horrific image or memory.) The trembling here isn't the object of practical knowledge; making myself tremble is. Or take 'going to sleep'. From a purely physical perspective, such a description refers to a process whose last phases can't be voluntary or intentional. But of course "I'm lying down because I'm going to sleep" can be a perfectly fine thing to say. Again, we don't have to say that the last phase of what is described is something the agent does intentionally, or with practical knowledge, just because we talk like that, since an intelligible (if long and convoluted) paraphrase is readily available.

This is why I don't think the case of 'sliding on ice' causes trouble for my interpretation of condition (a). On my account, the description doesn't fall under condition (a) because it isn't a human-action description. But then it doesn't fall under, or fall out of, condition (b) either, as it can't be a description of an executed intention, even if it aptly describes something that the agent brings about. So I don't have Anscombe proffering two kinds of practical knowledge. As we don't have another good all-purpose description that we could invoke here to say what we really mean after a mild bit of metaphysical reflection — something to capture (in Anscombe's phrase) 'the turn of the body' involved in sliding — we have to introduce something awkward like 'causing myself to slide on the ice' as the relevant description that properly captures the content of my intention and knowledge. But that description *is* a human-action description and so falls under condition (a). It is the kind of description about which we can ask after the 'Why?' or 'How?' involved, just as we can ask after the 'Why?' or 'How?' involved when someone (tragically) throws themselves out of a window, though we can't ask the same about what is happening as they sail through the air to their end.

## 8. THE FORMAL DEPENDENCE OF ACTION AND KNOWLEDGE

At this point, we are close to bringing together the pieces of an alternative account of what Anscombe means by her Thomistic slogan that practical knowledge is the cause of what it understands. To that end, let us consider a little further some of Anscombe's remarks from outside *Intention* that illustrate the particular manner which she thinks intentional action is constitutively tied to practical thought. In 'The Two Kinds of Error in Action' (1963c), Anscombe has an imaginary opponent offer the following line:

What bearing can what the agent thinks have on the true description of what he does? Someone may want to say: if what she does is a happening, a physical event, something 'in the external world', then that happening must be something that takes place, whatever the agent thinks. If you give a description of it, for the truth of which it matters what the agent thinks, such as "He got married", "He swore an oath", "He murdered his father", then your description ought to be broken down into the descriptions of thoughts and of purely physical happenings. (Anscombe 1963c, p.4)

In 'On Promising and its Justice' (1969), she considers what many would take to be the most obvious way to try to develop such a contrary position, one wherein the agent's thought serves as the 'efficient' cause of what takes place:

Even if my thinking can be a part of the cause of something's coming to be — which is often the case, since I bring things about by planning them and then carrying out the plan — all the same the thinking is only an efficient cause: the substantive fact that is caused partly by my thinking — assuming that what I plan is not itself a thought or set of thoughts, e.g. a calculation — can be described without mentioning what I think. (Anscombe 1969, p.13)

Anscombe argues this view is false by counterexample: actions like *marriage*, *making a contract*, *giving a present*, and *swearing an oath* cannot be broken up into descriptions of thoughts and descriptions of 'purely physical happenings'. As an example, Anscombe asks us to imagine a scenario where all the conditions of

a marriage are in place — the guests have arrived, the celebrant is on hand, the betrothed are about to make their vows — but the groom thinks the whole thing is just a rehearsal. In that case, Anscombe says, there will be no marriage (*cf.* Anscombe 1963a, p.87). But if the agent must think that she is getting married for the event to take place, how could we describe such an event or fact independently of the agent’s thought?

At a first pass, the causal approach looks like a non-starter.<sup>111</sup> The important point for our purposes though is not whether the causalist has a way out, but a detail of Anscombe’s argument that can be overlooked if we merely describe it by saying that it is supposed to show that the agent must think or believe or know that they are doing *A* to be intentionally doing *A*. In the first place, the agent must surely know that she is doing *A* to be doing so intentionally, but more importantly, even if we add knowledge we still leave out the specific *way* in which Anscombe thinks the agent’s ignorance can invalidate certain descriptions of behavior. Her thought is that it is because certain descriptions are descriptions of actions which are done *in* doing other things that ignorance can invalidate their application. As she says in ‘The Two Kinds of Error in Action’ about various human-action descriptions:

It is characteristic of these and many other descriptions ‘*X*’ of human actions that one can ask: “In doing what was he *X*-ing?” For example, by pronouncing such and such words in such and such circumstances one gets married. *That is why it is possible for ignorance or lack of intent to vitiate consent or nullify an action of a certain description.* For example, one thought one was

---

<sup>111</sup> The argument here is a version of what is sometimes called ‘Hume’s Circle’, which pops up in a number of Anscombe’s papers. According to Ford (2011, p.102), Anscombe follows Hume in restricting the application of the circle to actions bound up with “human convention and practice”. Ford himself goes on to suggest the wider view that intentional action *as such* may fall within the circle. I think this wider view is Anscombe’s view. Although I won’t pursue the point here, I think the opening remarks of §46 of *Intention* (“But who says that what is going on is the building of a house ... We all do ...”) are putting in place the thought that human actions are *all* bound up with convention — or, as Hume would say, are all ‘naturally unintelligible’. For some suggestive discussion in this regard, see Anscombe’s posthumously published paper “Was Wittgenstein a Conventionalist?” (2011, p.223).

rehearsing a wedding ceremony. (Anscombe 1963c, p.4, emphasis added; cf. Anscombe 1969, pp.13-4)

To then say that an agent ‘does not know’ or ‘does not think’ or ‘does not intend’ to be (say) getting married elides how the absence of such facts can vitiate consent or nullify the description ‘getting married’. The full story will rather go something like this: the agent does not have the idea that in saying certain words she *is* making her vows; she thinks that in doing so she is only practicing her vows; so she does not intend to make her vows by uttering those words; but for her to be getting married would require that she did have such beliefs and intentions, because the human-action description ‘marrying’ is characteristically a description of an event which is done by an agent in some way, through knowingly pursuing some means — attending a certain ceremony, uttering certain words — with the intention of getting married.<sup>112</sup>

As an aside, notice that Anscombe doesn’t think this kind of dependence is found at work throughout everything the agent does. Again, from ‘On Promising and its Justice’:

When an event, or action, is exterior, then there is always *something* that happens, whose happening is not disproved by any evidence as to the thoughts of the people involved. Perhaps they did not get married, but they went through certain movements and uttered certain sounds; perhaps he did not *say* anything, but word-like noises did come out of his mouth. (Anscombe 1969, p.17)

I don’t think it is accidental that Anscombe describes the events which the agent brings about which are not dependent upon practical thought with purely physical descriptions — ‘going through certain movements’, ‘uttering certain sounds’ — rather than with human-action descriptions. For the application

---

<sup>112</sup> For similar remarks, see also ‘On Promising and its Justice’, where Anscombe (1969, p.17) talks about how a person marrying must have “appropriate *expectations* and *calculations*” in connection with the activity in question. This is also the main example Anscombe uses in §48 to illustrate the character of descriptions that satisfy condition (a) (1963a, p.87).

of such descriptions don't seem essentially tied to practical thought. And this comports with what I said about the examples involving purely physical descriptions discussed in the previous section. If I fall off a branch with the intention to land on someone below, then my causing myself to fall is dependent upon my practical thought. Perhaps I didn't fall intentionally. Perhaps I thought about it and got so nervous that I slipped. Or perhaps someone pushed me. In any case, once I am falling through the air, my thought has no further bearing, and can't determine whether it is true that what I am doing can be described by the description 'falling through the air'.

Now I earlier suggested that we should say that the descriptions that satisfy condition (a) are of a type to be formally descriptions of executed intentions in the sense that they *can* apply to intentional actions — that is, they have a sense with the potential to characterize events in the *A–D* order, something that comes out in how their application prompts us to ask 'Why?' or 'How?'. But the application of such descriptions does not entail, as Setiya's reading would have it, that we have an intentional action in view. So now let us turn to consider Setiya's detailed rendering of how Anscombe's argument runs:

The content of one's knowledge matters to its causality. It is only when one knows that one is doing A and 'doing A' is formally the description of an executed intention, that practical knowledge is the cause of what it understands. It is the *formal* cause in that it forms the essence of its object. According to Anscombe, you cannot be paying, hiring, or marrying unless you are doing so intentionally. But in order to act intentionally, one must have practical knowledge of what one is doing. Thus, part of what is involved in an instance of paying, hiring, or marrying is the practical knowledge of its agent. Such knowledge is both necessary and sufficient for — in fact, constitutes — the action it represents. Where the content of one's knowledge is not formally the description of an executed intention, this argument does not apply. The object of practical knowledge when I know that I am sliding on ice, or dropping the ball, is something that could happen without being the object of such knowledge, which is not its formal cause. (Setiya 2017, p.159)



As Setiya reads it, Anscombe's slogan that practical knowledge is the cause of what it understands purports to record a (necessarily true) biconditional claim that an agent  $X$  is doing  $A$  if and only if  $X$  has practical knowledge of doing  $A$ . The sufficiency half is a triviality, and Anscombe doesn't need to mention it explicitly. What the combination of condition (a) and condition (b) are then meant to put in place on this account is the necessity half: the first condition says that if we have a description of the type 'doing  $A$ ' then we have an intentional action of doing  $A$ , and the second tells us the intention is executed, and so an object of practical knowledge. But if we have a description like 'dropping the ball' which doesn't satisfy condition (a), the necessity claim doesn't hold.

Although this promises to give us a very clear rendering of Anscombe's slogan, we've seen that it rests on a faulty reading of condition (a), and so can't be right. But there is in fact a further deficiency in Setiya's account that should be pointed out, which is that it doesn't explain all the features of practical knowledge as Anscombe envisions it. Immediately after offering her account of practical knowledge, she offers the following clarificatory remark:

This [the slogan about practical knowledge] means more than that practical knowledge is observed to be a necessary condition of the production of various results; or that an idea of doing such-and-such in such-and-such ways is such a condition. It means that without it what happens does not come under the description — execution of intentions — whose characteristics we have been investigating. This can seem a mere *extra* feature of events whose description would otherwise be the same, only if we concentrate on small sections of action and slips which can occur in them. (Anscombe 1963a, pp.87-8)

The sense in which her argument is meant to show *more* than that practical knowledge is a necessary condition on the occurrence of an intentional action is not easy to make out, and I won't pursue the point. The claim to which I want to draw attention is that Anscombe says that "an idea of doing such-and-such in such-and-such ways is such a condition." That is, she clearly says that a necessary condition of practical knowledge being the cause of what it understands is that the agent have an "an idea of doing

such-and-such in such-and-such ways”. Although this claim has occasioned no comment from Anscombe’s readers, it needs to be explained.<sup>113</sup> What does Anscombe mean by such an idea?

I suggest the natural reading of this expression is that for Anscombe it involves the “calculation of means to ends, or of ways of doing what one wants to do” (1963a, p.73). There are a number of reasons to read the expression in this way. Such a reading is suggested by the fact that when Anscombe discusses the notion of calculation operative in action, she insists on the same expression with the prepositional phrase ‘in’, of doing one thing *in* doing another (1963a, pp.1, 31, 34-5, 50, 79, 86). And immediately before telling us that such an idea is necessary for practical knowledge, she says that we “describe what *further* [agents] are doing *in* doing something (the latter description being *more* immediate, nearer to the merely physical)” (1963a, p.86), and gives the example of a cat “stalking a bird *in* crouching and slinking along with its eye fixed on the bird and its whiskers twitching”.<sup>114</sup> Moreover, she has just spent almost a third of the book discussing practical reasoning, on the stated assumption that we can’t understand practical knowledge if we don’t understand practical reasoning (1963a, p.57). As for the talk of an *idea* rather than of belief or knowledge, this too can be explained by the fact that she means to invoke practical reasoning, for she earlier remarks upon the role of ideas in characterizing the particular type of *wanting* that she sees at work in practical reasoning: “where the thing wanted is not even supposed to exist, as when it is a future state of affairs, we have to speak of an idea, rather than of knowledge or opinion” (1963a, p.70).

---

<sup>113</sup> Setiya (2017) ignores it; Moran (2004, p.54) elides it in citation.

<sup>114</sup> She calls this providing an “enlarged description” of an event. (Compare Anscombe’s notion of an ‘enlarged description’ with the so-called ‘accordion effect’ in Davidson (1971) and Bratman (2006).) At the risk of going a bit too far, I note that her expression “an idea of doing such-and-such in such-and-such ways” is in the plural, which picks up immediately on the example she has just given, of the cat stalking in crouching *and* slinking.

Now, as noted, Setiya's account doesn't explain the necessity of such a calculative idea or belief to practical knowledge. But the current account does. For if the descriptions given by condition (a) characterize the agent as doing something *in* doing else, as I'm suggesting, then we have a natural explanation for why such a calculative idea must be involved in practical knowledge, as such cognition will be operative in the agent's action by dint of the sort of process being described. So when Anscombe (1969, p.12) says, as she does in 'On Promising and its Justice', that she is "concerned with facts which are formally dependent on human thought about them", I think she means the particular kind of formal dependency where the agent's practical thought is of doing one thing in doing another.

This brings out a further difference with Setiya's reading. For although one virtue of Setiya's account is that it lets us see how Anscombe's argument could be valid, it also leaves the argument as unsound. For by Setiya's (2017, p.158, *n.4*) lights — and, I've argued, Anscombe's too — it is false that descriptions that satisfy condition (a) must characterize intentional doings. But the current interpretation does better on this score. Or at least it does if the argument of chapter 2 is correct. For on my account Anscombe only needs to say that when the description of an agent's doing characterizes her as doing one thing in doing another then it follows that the agent is acting intentionally, and with practical knowledge. And this is just what the argument of chapter 2 establishes. So if my interpretation is right, I think we can still help ourselves to a version of Setiya's rendering of Anscombe's slogan, that practical knowledge forms the essence of its object in being both necessary and sufficient for it, but now we've also credited her with a sound route to holding it.

## 9. THE NEGLECTED CATEGORY OF 'MERELY' VOLUNTARY ACTION

If this account of Anscombe's argument is right, it raises a further very obvious question: what does she say about actions which don't involve the agent doing one thing in doing another? Does Anscombe

think that we don't have practical knowledge of such actions? Or does she not countenance such actions at all? As it happens, I think Anscombe does discuss a category of action that fits this description in *Intention*. And in fact consideration of this category will bring us full circle and help me fill a lacuna left open in my story. Recall that we earlier saw Anscombe describe a case of idly scratching one's head as a mere act of a human, and not a human action. Mere acts of humans, we saw, break up into the physiologically voluntary and the physiologically involuntary. Throughout, I've relied on the claim that actions that fall into the latter category are not intentional and not objects of practical knowledge. In describing idly scratching one's head as a mere act of a human I took Anscombe to be describing these as physiologically voluntary, and not physiologically involuntary, actions, though ones which fall outside the intentional or the voluntary in the narrow sense of being a human action. At the same time, it is natural to think that Anscombe would allow that a humble scratch of the head could be a human action, indeed an intentional action. So what is Anscombe getting at in labelling some head scratchings as mere acts of a human rather than full-blown intentional or voluntary actions?

To answer this puzzle, let us consider a neglected passage from §49 of *Intention* where Anscombe introduces a further kind (or classification) of voluntary action which I have so far ignored:

Mere physical movements, to whose description our question 'Why?' is applicable, are called voluntary rather than intentional when (i) the answer is e.g. 'I was fiddling', 'It was a casual movement', or even 'I don't know why' (ii) the movements are not considered by the agent, though he can say what they are if he does consider them. (Anscombe 1963a, p.89)

Anscombe gives us no examples to illustrate; but something like idly scratching one's head seems a reasonable candidate, so long as the two conditions are satisfied.<sup>115</sup> For ease of reference, let us say that any "mere physical movement" which satisfies (i) and (ii) is a "mere voluntary action". As in the case of her remarks about idle head scratches, what Anscombe says here is puzzling. If the question 'Why?' is

---

<sup>115</sup> To avoid confusion with the argument we've been so far discussing, I have changed Anscombe's lower case letters in this passage to Roman numerals.

not only applicable to mere voluntary actions but can be given application to it by giving an answer like “It was an idle action” or “I was fiddling”, then why are these not intentional actions? And how does that not contradict Anscombe’s definition of intentional actions as those to which the question ‘Why?’ are given application?

Anscombe’s commentators have had nothing to say about these questions. One reason for the neglect may be the general skepticism flagged at the outset of §5 concerning Anscombe’s remarks about voluntary action, a skepticism which we’ve set aside. A second reason may lie in the remark itself. For notice that Anscombe only says that such actions are “called” voluntary rather than intentional, and that might be taken to show that she is here only making an Austinian point about usage. That, however, seems unconvincing. Anscombe says at the outset of *Intention* that “the object of the whole enquiry is really to delineate such concepts as the voluntary and the intentional” (1963a, p.10), placing the sort of investigation being made into the two on a par. She also routinely uses the locution ‘*x* is called *y*’ where her intent is clearly to make a claim about *x* and *y*, not about our use of ‘*y*’. When this remark in §49 is foreshadowed in §17, she also says that “an action of this sort is voluntary, rather than intentional” (1963a, p.26), dropping the qualification about language. And as a point about usage, does it not just seem false? Who calls these actions “voluntary”?<sup>116</sup>

So I want to hazard an interpretation of the passage. The nutshell version is this: merely voluntary actions are physiologically voluntary actions wherein there is no further ‘Why?’ or prior ‘How?’; and since the intentional requires at least one of the two (because it requires an idea of doing such-and-such in such-and-such ways), mere voluntary actions are not intentional actions. I’ll try to spell

---

<sup>116</sup> We might compare here Anscombe’s remarks about the sense of ‘involuntary’ which I set aside in §5, the usage we apply to intentional or physiologically voluntary actions performed under some kind of compulsion or duress. Part of the reason I set that case aside earlier on is that it does seem to mark out a way we talk — and Anscombe describes this case very much in terms of our usage or linguistic practice — but it also seems right that we do talk this way, and that we want some word to mark it. As Anscombe herself says, the more typical word here would just be ‘reluctant’ or ‘compelled’, and the availability of those terms is also part of the ground for not feeling the need to press deeper here.

this out at greater length in reverse, but first, let us consider how the possibility of merely voluntary actions relates to Anscombe's famous definition. I think there is actually no problem here, but we need to go beyond the most familiar form of the definition to see why.

First, let us recall the definition yet again: intentional actions are "actions to which a certain sense of the question 'Why?' is given application [...] in which the answer, if positive, gives a reason for acting" (1963a, p.9). That obviously doesn't sound like a full definition, as it clearly looks forward to the possibility of 'negative' answers (though Anscombe doesn't use that phrase). She ends up telling us what sorts of answers mean that the question is "refused application", as when the agent doesn't know she is doing *A*, or only knows through observation, or does know but the action is involuntary. Such answers show that the question 'Why?' is not applicable, and could be called 'negative' answers in one sense. The trickier category, though, is when the question *is* applicable. One kind of answer that shows that it is applicable is the positive answer, one that gives a reason for action, such as a forward-looking motive, or an intention. But there are also what we might call 'neutral' answers. Such answers show that the agent is not acting for any further reason. But they come in two sorts. The first sort shows that the question 'Why?' is applicable — like positive answers, such answers 'give the question application'. In §17, Anscombe says that neutral answers of this sort can include "I just thought I would", "It was an impulse", "For no particular reason", or "It was an idle action — I was just doodling" (1963a, p.25). The second sort, however, neither makes the question out to be applicable or inapplicable — but they don't show the question is 'refused application'. Answers of this latter sort show that the 'Why?' question "has and yet has not application: it has application in that it is admitted as an appropriate question; it lacks it in the sense that there is no answer" (1963a, p.26). The only neutral answer of the second kind she gives

in §17 is “I don’t know why I did it”.<sup>117</sup> Such an answer is probably also better described as a negative answer.

Now, looking back on §17, Anscombe says that her discussion there showed that it is wrong to think that “an action cannot be called voluntary or intentional unless the agent has an end in view” (1963a, p.34). And it is clear already in §17 that she means to allow for the possibility that in some cases a neutral answer can reveal an action to be voluntary rather than intentional (1963a, p.26). That means that if we then plug in the possibility of neutral answers to further fill out Anscombe’s definition, we shouldn’t expect the result to sharply delineate the category of intentional action at all. And this is in fact just what we find when Anscombe offers her later and final elucidation of the question ‘Why?’ in light of the possibility of the cases discussed in §17:

The answers to the question ‘Why?’ which give it an application are, then, more extensive in range than the answers which give reasons for action [i.e. they include neutral answers that don’t offer reasons for action]. This question ‘Why?’ can now be defined as the question expecting an answer in this range [i.e. the wider range, including neutral answers]. And with this we have *roughly* outlined the area of intentional actions. (Anscombe 1963a, p.28; my emphasis).

Borrowing a line from Michael Thompson (2011), we might say this early definition is rough because it is only a nominal definition; the real definition comes later, in the form of the argument that has been my main quarry in this chapter. But I think we can see that it is rough in a more straightforward sense, in that the question ‘Why?’ can be given application with neutral answers which don’t decide between whether the action is intentional or only voluntary. And it surely is right to not want a story about this region that does anything more than draw rough lines. Once we start to consider particular instances of

---

<sup>117</sup> Anscombe also includes in this second category the answer given by the man who is pumping poison but is just doing his normal job, but that case is a distraction here.

smiling, head scratching, lip biting, foot tapping, thumb twiddling and eyebrow raising, any attempt to mark sharp divisions between the intentional and the voluntary actions or indeed the involuntary will surely seem suspicious.<sup>118</sup> Even so, we can now see why Anscombe can allow for merely voluntary actions which are not intentional to fall under the aegis of the question ‘Why?’ without giving up her definition.

What we still need, though, is an explanation for why Anscombe thinks such actions are not intentional. Here, we need to recall that Anscombe requires practical knowledge to involve a calculative idea or calculative belief. If intentional action entails practical knowledge, then it follows that intentional action also requires such a calculative idea or belief. And this, I think, is why Anscombe thinks merely voluntary actions are just that and not intentional: they don’t involve such a calculative idea or belief.

As I noted earlier, Anscombe doesn’t actually provide examples in the relevant passage, nor does she tell us what sorts of action-descriptions might be involved. Since she is talking specifically about “mere physical movements”, we are presumably not dealing with actions under descriptions like ‘marrying’ but rather descriptions like ‘scratching my head’, ‘twiddling my thumbs’, ‘jiggling my foot’, or whatever sorts of descriptions we use for all manner of bodily movements for which we lack idioms like those just mentioned and which we could only describe with words like ‘moving your arm like that’. Now suppose we have an action under some such description, but it is also one that satisfies condition (i), that the answer to (say) “Why are you jiggling your foot like that?” is something like “It is an idle action” or “I don’t know why”. If such an answer is truthful, then the agent isn’t moving their foot like that with a further intention, and so we can’t see at work in her action a calculative belief of the form “In moving my foot I’m doing such-and-such”.

Now that doesn’t yet mean that there is no such belief at work. In general, it does not follow from the fact that the agent can give the question ‘Why?’ application with a neutral answer that there is

---

<sup>118</sup> For some remarks along these lines and discussions of cases, see Teichmann (2008, p.42).



no such belief at work in what they're doing. Consider, for instance, Gavin Lawrence's (2004, p.268) example of "advanced fiddling" — I'm arranging some matchsticks into a little house while having a conversation and when pressed I explain myself by saying "Oh, I'm just fiddling". It seems absurd to say that there is no practical thought involved here. I am, after all, building a house, just a little one. Or, to take one of Anscombe's examples, if I act for a backward-looking motive like revenge — "I killed him because he killed my brother" — then the explanation of my action does not, by her lights, involve "a piece of reasoning or calculation" (1963a, p.65; also pp.73, 81). But that doesn't mean the action itself doesn't involve a calculative belief. It surely does involve such a belief, as it doesn't seem possible for one person to intentionally kill another without having some grasp of how they are doing so — for example, by using a knife — the explanation just doesn't happen to bring that out. The point generalizes: there are surely many actions which can't but be done by an agent without some consideration of how they are done. And so too, there may be actions that cannot but be done for some further end.<sup>119</sup>

Remember though, that Anscombe is at this point focused only on what she calls "merely physical movements". We're not, for instance, considering actions involving dealing with objects like matchsticks or knives. Since condition (i) rules out the presence of a calculative belief about a further end, the question is whether we have any reason to deny the presence of a calculative belief about a prior means. I think this is just what condition (ii) is meant to do with its restriction that the movements "are not considered by the agent". This talk of what the agent has not 'considered' has a helpful precedent in *Intention*. In §30, Anscombe argues against the claim that practical knowledge can't extend further than our knowledge of our bodily movements by devising an example where I have practical knowledge that goes beyond such movements but doesn't include knowledge of them. She imagines an experimental

---

<sup>119</sup> Anscombe (1963a, p.27) suggests we could not intelligibly grasp someone 'carefully arranging all the green books in the house on the roof' by offering a neutral answer: "his words would be unintelligible unless joking or mystification" (1963a, p.27).

task where the subject has to keep an object level in motion by pumping a handle up and down in the right way. In this case:

My account of what I am doing is that I am keeping the thing level; I don't consider the movement of my arm at all. I am able to give a much more exact account of what I am doing at a distance than of what my arm is doing. So my keeping the thing level is not at all something which I calculate as the effect of what I really and immediately am doing, and therefore directly know in my 'knowledge of my own action'. (1963a, p.54)

The movements are here not considered as calculated means to an end. But I think Anscombe would equally say that if the movement isn't "considered at all", then we haven't considered a means through which to perform the movement. And so this is what condition (ii) puts in place, thus working in concert with condition (i) to close off the possibility that there is a calculative belief at work in merely voluntary movements.

To spell out the idea a little, consider again the case of jiggling my foot. This is something I can do intentionally. And it is something I can do with further intentions: I can jiggle my foot in order to let my confederate know that the assassin is in the room. It is also something I can do in different ways. I can jiggle my foot and consider the movements that I'm going through (or, as it were, putting my foot through) in just the way I can consider anything else I effect through my chosen means. On the other hand, this is the kind of activity I can also find just 'going on', as when I'm sitting talking to someone and I realize that I'm jiggling my foot. It is a familiar experience to find oneself making some such movements, and taking control of them in some way, maybe by stopping them, or picking up the pace, but in any case, certainly 'considering' them. When they aren't in view in that way, there is no practical thought at work in them of the sort which I think Anscombe takes to be essential to intentional action. But the actions are still physiologically voluntary. If that much is right, we can make sense of why Anscombe thinks merely voluntary actions are not intentional.

What, though, would Anscombe say about our knowledge of such merely voluntary movements? If they don't serve as objects of practical knowledge, how do we know them? The first thing to note here is that immediately after the passage describing the defining features of such movements, Anscombe in fact takes the time to justify the claim that we have non-observational knowledge of such actions. If she thought that we do have practical knowledge of such actions, given that she has just completed her account of practical knowledge in the previous section, we would have to wonder why she bothers to say anything here. But she does take the time to say something, and in particular to discourage the thought that when we do consider such movements we make a kind of empirical discovery; rather, she thinks that we can grasp what the movements are by "going through the motions (i.e. of tying a knot) in imagination" (1963a, p.89). So it looks like there is an epistemological story already on the table in *Intention* which we can apply to these actions: they are yet another subclass of the things she thinks we "know without observation".<sup>120</sup>

Does this mean that our knowledge of such actions is a form of speculative knowledge, or knowledge which is 'derived from the objects known'? However, as McDowell (2011a) has pointed out, Anscombe nowhere explicitly says that the division of knowledge into practical and speculative must be exhaustive. Although it would take us too far afield to pursue the details, McDowell offers an account of Anscombe's view of our knowledge of the position of our limbs in which it falls into neither category: on the one hand, our knowledge of limb position does not involve activity, and so is not a species of practical knowledge, but on the other hand, such knowledge is not of a distinct object garnered through sensory mediation, and so it is not speculative. Rather, on McDowell's account of Anscombe's view, our knowledge of limb position is a species of self-knowledge, a variety of non-observational, non-inferential

---

<sup>120</sup> When Anscombe discusses the class of things known without observation in the early sections of *Intention*, she only considers intentional actions, mental causes, limb position, and certain *involuntary* bodily movements as subclasses, not *voluntary* movements. As she has just introduced this subclass (or a subclass of that subclass) in §49, it thus makes sense for her to take the time to explain how they can qualify as a form of knowledge without observation.

knowledge of one's self as a bodily agent.<sup>121</sup> Assuming, as I shall, that we can read Anscombe like this on limb position, can we not say the story will apply to our knowledge of merely voluntary movements?

An obvious rejoinder here will be that even if we accept McDowell's opening of a space between practical and speculative knowledge, the case of mere voluntary movements is different from that of limb positions precisely in that the former is a form of *agency*. That is true enough, but it would be question-begging to simply assume that for Anscombe all forms of agency must consist in practical knowledge. What is the *argument* for why merely voluntary movements must qualify as such a form, rather than falling between the gaps like our knowledge of limb position? One answer that may be tempting is that our knowledge of merely voluntary movements shares the same normative structure as practical knowledge, in being subject to the possibility of a *mistake in performance* which according to Anscombe characterizes practical knowledge, rather than the possibility of a *mistakes in judgement*, to which the varieties of theoretical knowledge are liable.

Now I think that view only needs to be stated to see what is wrong with it. In Anscombe's discussion of the two kinds of mistakes, a man is going around town with a shopping list while being trailed by a detective making her own list of the man's purchases, and if what the man buys doesn't match what is on the detective's list we can say the detective has made a mistake in judgement — perhaps a trick of the light has made the blueberry jam look like a jar of Vegemite — but if the man's shopping basket doesn't match *his* list then, assuming he hasn't made any mistakes in judgement in the lead up to shopping, or changed his mind along the way, the mistake lies not in his judgement but in the what he has done. But what are we to make of the idea that someone could be, say, doodling incorrectly? Is there a wrong way to absentmindedly tap one's fingers on a table?

Of course, if we imbue the tapping involved with some calculative idea — suppose I'm trying to tap out a particular rhythm, or I'm trying to signal my confederate by my tapping — then we can make

---

<sup>121</sup> McDowell offers this not so much as an interpretation of Anscombe as opposed to a “charitable reconstruction” (2011, p.143). For some engagement with McDowell on this score, see Bermúdez (2015).

perfect sense of a mistake in performance, in Anscombe's sense: I know how the rhythm goes, I have the skill to do it, I haven't changed my mind about tapping it out, but I just got the rhythm wrong, or I tapped too quietly for my confederate to hear. That case would be just like Anscombe's shopper. But this is just to say that we've dropped one or both of conditions (i) and (ii) in our description of the case, and now have an intentional action in view. Keeping the two conditions in place and focusing on mere physical movements, it is hard to see what could count as getting an action right or wrong, *qua* performance. The question doesn't seem far off asking whether someone can make a mistake in performance while they're falling out of a tree. The only mistake that seems to be available here is the idea that making the movement at all is mistaken — perhaps the agent wants to stop being so fidgety — but Anscombe's notion of a mistake in performance seems something of a different order from whether an agent lacks self-control.

As it happens, I think Anscombe is explicit in ruling out that such cases could be subject to mistakes in performance. This isn't clear from *Intention*, but in another place where she discusses the concept of a mistake she appears to characterize mistakes in performance as restricted to actions that involve a calculative idea. In this text, Anscombe is discussing Wittgenstein's conception of a mistake, but in a way that makes clear she is running with the view herself, and her discussion helps to clarify her own construal of the notion in various ways.<sup>122</sup> For our purposes, the following remark is crucial:

---

<sup>122</sup> For instance, it is a long-standing objection to Anscombe's description of the shopper that we want to say there is a mistake in both the performance (which fails to match the intention) and the expressed judgement (which fails to match the performance). For this complaint, see Pears (1964, p.29; 1975, p.51), Houlgate (1966, p.259), Moran (2004, p.61), Setiya (2008, p.396, n.28), Grünbaum (2009, pp.49-52), and Roessler (2010, p.239). It is clear in her remarks on Wittgenstein that she is using his notion of a mistake in a way that it requires more than falsity (and correctness more than truth): "It runs through Wittgenstein's thought that you haven't got a *mistake* just because you have as a complete utterance a string of words contrary to one in which some truth is expressed. I can be accused of *making a mistake* when I know what it is for a given proposition (say) to be true, and things aren't like that but I suppose that they are" (Anscombe 1976, p.124). That is the beginning of a reply to the long-standing objection.

The conception of ‘mistake’ which Wittgenstein has is obviously something like this: in order to say that someone thinks that  $p$ , and it is a mistake on his part, it has to be clear what it is for  $p$  to be true (or correct); he must be acquainted with this. E.g. the situation must be such that *if* he had noticed the relevant features, he would have agreed that  $p$  was false. Or, if it should be a calculation, he must be ready to take back a slip or false step. (Anscombe 2011a, p.222)

Anscombe is here recalling the two kinds of mistakes. Mistakes in judgement are those where the agent ought to agree that the claim is false, if she notices the relevant features of the situation. Mistakes in performance are those where she must be disposed to take back a slip or false step. But the latter are explicitly pitched as involving *calculation* (or as ‘being’ a calculation<sup>123</sup>). For Anscombe, mistakes in performance thus seem to be restricted in their scope to intentional and voluntary actions proper, not mere voluntary movements that involve no calculation. At the same time, just as Anscombe (1963a, p.14) thinks it can be perfectly “clear” that someone has made a mistake in judgement when they say their leg is bent when it is not, there doesn’t seem to be any problem with the thought that I just make a mistake if I say I’m scratching my chin when I’m not. The same normative assessment seems readily available for claims about other merely voluntary movements.

So what does this all mean for Anscombe’s view? I’ve here taken up a neglected challenge of trying to make out Anscombe’s puzzling remark about merely voluntary movements. The interpretation I’ve proposed is that Anscombe takes intentional actions to always involve various descriptions linked as means to end, and that this comports with her account of the distinctive type of normative assessment that governs intentional actions (and voluntary actions, when they are not mere physical movements).<sup>124</sup>

---

<sup>123</sup> We can, I think, make sense of this form of expression if we recall that for Anscombe calculation or practical reasoning concludes in an action (or a decision).

<sup>124</sup> I should stress “suggested”. I haven’t tried to establish why a mistake in performance must involve calculation, but only suggested that there is an intuitive sense in which it isn’t clear how actions devoid of such a structure could involve such mistakes.

Why think otherwise? One reason may be historical. It is sometimes said that Anscombe has been too often read through a Davidsonian lens.<sup>125</sup> As my first epigraph recalls, Davidson popularized the idea that an intentional action must be intentional at least under *a* description. But as my second epigraph suggests, Anscombe perhaps goes a little further: for her, an intentional action has “several [descriptions] under which it is intentional”. In a later lecture, she is emphatic on this point:

Whenever you do anything, there are, as you might put it, lots of things you do. Take an example: you endorse a cheque ... . What are you doing? You are depositing ink on a bit of paper, you are writing your name, you are signing a cheque on its back, you are perhaps thereby paying a bill for someone ... All these might be comprised in your one act. ... And there is also your *objective*. You do the kindness, perhaps, in order to improve family relations. *That* lies in the future: it is a further intention, perhaps *the* main one ... So we can ask this question about what you do: ‘What is the character, or rather, what are the characters, of the act, as intentional, at the time when it is done?’ What are you doing, and doing on purpose? And there is also the question: ‘What are the further intentions with which your act is done? What is its goal?’ These are different questions. The first *always* has several answers. The second may have only one. (Anscombe 1979b, p.193)

If I understand her, she doesn’t mean to say that an action has several descriptions of any sort under which it is intentional, but several descriptions linked in the *A–D* order. Now Anscombe may here slip about her own doctrine, as it is clear from *Intention* that the possibility of neutral answers like “For no particular reason” shows that she allows for intentional actions which have no further intention. Does she also slip in her first claim, that intentional actions must always be comprised of more than one intention in action? Or is she perhaps just describing the central case of intentional action? The assumption I have been working with is that this is indeed the central case of intentional action. If the interpretation of Anscombe’s puzzling remark which I’ve offered here is right, then it is one that can at least be made on good authority.

---

<sup>125</sup> See Stoutland (2011), Moran & Stone (2011), and Hornsby (2011).

## APPENDIX

### IN WHAT SENSE WAS ANSCOMBE AN ANTI-CAUSALIST?

“I was reading ... Anscombe ... and all these other Red Book philosophers.”

— Davidson, *Problems of Rationality*, p.252.

Most philosophers would agree with Davidson’s characterization of Anscombe as a fellow traveler alongside (several of) the authors of the so-called ‘little red books’ that emerged in the post-war period which urged various forms of anti-causalism.<sup>126</sup> I want to question this interpretation. The thought behind this reading of Anscombe seems to be that in her various negative arguments against certain causalist positions, she evinces a commitment to just the sort of prejudices about causation and causal explanation that characterized some of her Oxbridge contemporaries, prejudices that in effect force her

---

<sup>126</sup> For a few recent examples of descriptions of Anscombe as an anti-causalist, see Alvarez (2007, p.103), Stout (2007, p.141), Bratman (2008, p.3), Driver (2011), Hornsby (2011, p.125), Marcus (2012, pp.150-1), Stoecker (2013, p.17) and Hyman (2014, p.97). I suppose this reading begins with Davidson, who, in the opening footnote of his 1963 paper, lumped Anscombe in together with those whom he took to deny that reasons explanations were “a species of causal explanation”, and whom he went on to attack in the paper. Note, though, that unlike all (but one) of the other eight philosophers he listed, Anscombe is nowhere cited in the body of his paper for textual evidence of the sort of views about causation and causal explanation that he went on to dismantle. If what follows is right, this is for good reason. I’ll return to Davidson’s particular worries at the close.



to deny or be blind to the possibility that reasons explanations are causal explanations in *any* sense. Here, for instance, is Rowland Stout, discussing Anscombe's notion of a 'mental cause':

[Anscombe] rejects the idea that explaining an action is done by describing the 'mental cause', and in so doing she appears to reject causalism. And many of the crop of contemporary anti-causalists can be described in the same way. They reject a particular version of causalism — one which perhaps identifies reasons for action with mental causes — without considering other ways that reasons for action may be involved in causal processes. (Stout 2007, p.141)

If the interpretation of 'Causality and Determination' which I sketched earlier is on the right track, we can see why such a line cannot be right. But here I also want to question the implicit thought that Anscombe's discussion of causality in *Intention* and elsewhere shows her to be blind to possible alternatives to the orthodoxies of her day.

Let us begin with the case of 'mental causes', perhaps the most familiar source of Anscombe's supposed theoretical biases about causality. For Anscombe, mental causes are "what someone would describe if he were asked the specific question: what produced this action or thought or feeling on your part: what did you see or hear or feel, or what ideas or images cropped up in your mind and led up to it?" (1963a, pp.17-18).<sup>127</sup> She thinks there may be cases where it is hard to clearly distinguish between a mental cause and a reason for action, as when we act on an order both because we have been given the order and because we've accepted that we are obliged to act on it (1963a, p.23). Thus, "what is commonly said, that reason and cause are everywhere sharply distinct notions, is not true" (1963a, p.24). But she goes on to dismiss mental causes as central to reasons explanations, on the ground that we can act intentionally or voluntarily in their absence. We need not recite her examples: just imagine, as a

---

<sup>127</sup> Anscombe is here grouping mental causes together into a single type, but outside of *Intention* she also says of these kinds of causes, "we could discuss the causality — how many different types there are" (Anscombe 1978, p.106). This also fits with the picture given in Chapter 3.

central kind of case, acting for a reason which doesn't consciously go through your mind in some fashion.

We can make a number of points about mental causes so defined. First, the terminology is, I think, a little unfortunate, as it suggests that Anscombe means to view *any* kind of mental causality or causality at work in the realm of the mind as fitting her definition of 'mental causes'. I don't think she means that. Second, the idea in itself seems harmless: there surely are mental causes in this sense. Third, her negative argument seems sound: we can act (intentionally or voluntarily, or both) without mental causes.

Anscombe's notion of mental causality can thus only look controversial if we assume that the negative argument which employs it is not directed at a very narrow view — a view that says all reasons explanations work by citing consciously experienced mental causes — and if we assume, as Stout does, that Anscombe must think that what is true for mental causes is true for causality simpliciter. Thus, John Hyman (2014, p.97) argues that Anscombe's discussion of mental causes reveals a blind spot against states being causes, rather than events. But Anscombe (1978, p.106) is happy to countenance states (of emotion) as mental causes. And Kieran Setiya (2007, pp.47, 56; 2009, p.130) argues that Anscombe's discussion of mental causality is both hampered by the presuppositions that mental causes must precede their effects (rather than also or only being concomitant with them) and that such causes must be consciously experienced, thus leaving out by stipulation the many contemporary causalist theories which require neither assumption. But Anscombe does not require that mental causes precede their effect as opposed to being concomitant with them (*cf.* Anscombe 1982, p.95). In fact, she may not have even thought that such causes *must* be consciously experienced.

This last point may seem surprising. The matter is not entirely clear, and, of course, turns in part on ambiguities about what we could mean by 'conscious experience'. Strictly speaking, for the purpose of showing that the line sketched by Stout is wrong, we don't need to show that she did construe mental causes in this way, but only that her account of what *she* calls mental causes doesn't require her to say

that *anything* that could be called a mental cause must be consciously experienced. Still, I think we can at least make a case for interpretation that Anscombe thinks that mental causes need not be consciously experienced, at least if that means something akin to ‘phenomenal consciousness’, in the contemporary jargon that we owe to Ned Block. For while Anscombe’s definition of mental causes cited above suggests they must be conscious in this sense, and her examples tend to have an ‘experiential’ character (e.g. I’m signing the document because “[t]he thought ‘It is my duty’ kept hammering away in my mind” (1963a, p.11)), it is important to note that for Anscombe the question ‘Why?’, when used to ask after a mental cause, is not — to adapt her jargon — ‘refused application’ in the case where the agent is unable to give a definite account of what “went on in [one’s] mind and issued in the action”, and can instead only “shrug or say ‘I don’t know that there was any definite history of the kind you mean’ or ‘It merely occurred to me ...’”. (1963a, p.17). This is for two reasons. First, because she thinks it is possible for a mental state (such as a state of fear) to do duty as a cause even if there is no answer to when the agent, as it were, entered or exited the state (so no ‘definite history’).<sup>128</sup> And second, because she thinks that mental events (such as occurrent thoughts) can do duty as causes even if they are not consciously experienced.

This last claim takes a bit of work to bring out. Notice first that Anscombe’s definition (1963a, p.17) of mental causes says that they are either “a mental event [or] something perceived” by the agent. This allows for thoughts to be mental causes, as her category of mental events includes thoughts. It is also a running theme of several of her papers that thoughts need not be consciously experienced to be *bona fide* events in the mind. For Anscombe, an (occurrent) thought “is an event, but not an experience; what we call its content is given by words which do not describe an inner experience (unless that is what the thought was about)” (1963b, p.63). In grasping such a thought, she thinks introspection — whatever mechanism that may involve — may not help; rather, it may be that “we have to consider something that

---

<sup>128</sup> For discussion of this possibility, see again Anscombe (1978, p. 106).

need not have occurred when the thought occurred, namely the words and the explanation of the words, which give the thought” (1963b, p.62). So Anscombe’s tendency to *describe* a thought which serves as a mental cause using a sentence token does not imply that the token must, as it were, ‘go through one’s head’. Now — to join the dots — her standard example of a thought that isn’t necessarily an experience is ‘being struck by a thought’, or ‘having a sudden thought’, and it seems as good a way as any to report such a thought by saying “It merely occurred to me ...”, which, as we just saw, is the form of words that Anscombe uses to illustrate one way the ‘Why?’ question is not necessarily ‘refused application’ when it asks after a mental cause even though the agent can report no relevant experience. Admittedly, though, the interpretative matters here are not at all clear.<sup>129</sup>

So much, then, for Anscombe’s discussion of mental causes.<sup>130</sup> Anscombe’s only other anti-causalist arguments occur in her 1983 paper ‘The Causation of Action’, which considers causalism in the by-then more familiar guise wherein reasons explanations work by pointing to the antecedent or contemporaneous states that bring about the action, but where these causes need not exist in conscious experience. In effect, she considers two versions of such theories. According to the first, the mental state (or states) that causes (or cause) an action must be “something holding of its subject here and now, or over a period of time, without reference to anything outside that of which it holds or the time at which it holds” (1983, p.99). Anscombe argues, on what we would now call ‘externalist’ considerations, that there

---

<sup>129</sup> For a forceful contrary reading of these parts of Anscombe’s text, see Ayers (1972).

<sup>130</sup> Notice though that mental causes are in fact very important for helping Anscombe delineate the sense of the question ‘Why?’ in which she is interested, because they provide her with a non-question-begging entry point in the topic, one that doesn’t presuppose concepts like *intentional action* or *voluntary action*. (The exact details of how they allow her to do this need not detain us.) But their importance should not be overstated. For instance, as part of her provisional elaboration of that sense, Anscombe (1963a, p.24) goes on to say that intentional actions are “marked off [...] by being subject to mental causality”. But this does not mean, as Candlish & Damjanovic (2013, p.698) think, that Anscombe *defines* the class of intentional actions as “those which have [i.e. each in fact have] a mental cause, and which we can explain by doing *more* than citing a mental cause.”

are no such purely ‘internal’ states, at least in so far as they purport to be ordinary psychological states with intentional contents. Suppose, for instance, that an agent has a belief that a bank in Pittsburgh will be open at 9 a.m. tomorrow. If the agent’s belief is construed as a particular internal state in the manner suggested, then we can envision a counterfactual circumstance in which this latter state arises in the agent even though banks, clocks, and Pittsburgh don’t exist. Anscombe finds it absurd to say that in such a scenario the agent could have the belief that a bank in Pittsburgh will be open at 9am tomorrow. Her conclusion is that there simply are no such purely ‘internal’ states, in so far as they purport to be ordinary psychological states with intentional content, and so they can’t be there to do causal work.

This is a familiar enough (if very controversial) type of argument, and one type of familiar response is to insist that such arguments show that causal explanations of action must involve ‘external’ or ‘wide’ mental states, and not the ‘internal’ states that serve as Anscombe’s target.<sup>131</sup> A position of this sort would presumably evade Anscombe’s ‘externalist’ objection. In fact, I think Anscombe accepts that such ‘wide’ states can be causes of actions. She distinguishes between mental states defined ‘narrowly’ as above, and what she calls a “state of a human being” (1983, p.101), by which she means to refer to intentional states construed as having ‘wide’ contents. And she allows that such states can cause actions:

Indeed they may, and the effect of an intention may even be an action in execution of that intention! E.g. suppose I have a standing intention of never talking to the Press. Why, someone asks, did I refuse to see the representative of *Time* magazine? — and he is told of that long-standing resolution. ‘It makes her reject such approaches without thinking about the particular case.’ This is ‘causal’ because it says ‘It makes her ...’: it derives the action from a previous state. (Anscombe 1983, p.95)

---

<sup>131</sup> See, for instance, Burge (1993), Yablo (1997) and Williamson (2000). I don’t mean to identify Anscombe’s position with that of other philosophers who think that ‘wide’ mental states are crucial to the interpretation of action. Unlike Anscombe, many are still happy to talk about the existence of ‘narrow’ states, and also take them to have some explanatory function.

We see here again the language of ‘deriving’ and ‘making’ from our earlier discussion, and similar comments recur through Anscombe’s papers in reference to the potential causal role of intentions and other states (Anscombe 1983, pp.98, 100-1; 1989b, pp.110-1). Although Anscombe allows such states to be causes, she again denies that reasons explanations *must* cite such cases. For on this view, it is still the case that the cause of an action cited by an explanation “has to be thought of [by us philosophers, not by the agent] as a distinct thing [from the action it effects]” (1983, p.96). Her argument against this position is that there need not always be such states. If, for instance, the putative causal state is supposed to be an intention, she claims the explanation won’t work in cases where “the intention is, so to speak, embodied in the action” (1983, p.101), since the cause can’t be metaphysically pried apart from the effect, and she thinks the same possibility arises for other intentional states that might be taken to serve as causes, such as beliefs and desires.<sup>132</sup> So again, we have a very contentious argument in play, but perhaps not the one we would expect if the standard reading of Anscombe were right.

That completes my survey of Anscombe’s anti-causalist brief. In effect, she has considered and dismissed three causal theories: that the explanation of action must cite mental causes, that it must cite ‘internal’ mental states, and that it must cite ‘external’ mental states. Her argument against the first and third views is that such causes are not always involved in reasons explanations; her argument against the second is that there are no such causes. As I’ve said, aside from the first argument against mental causes, Anscombe’s arguments are contentious. But my point here in speeding through them has not been to assess their soundness, but only to bring out that they don’t appear to require Anscombe to presuppose much about the nature of causality *per se*. Rather, where we find more familiar and substantive

---

<sup>132</sup> For analogous remarks about the case involving wanting or aiming, see her (1983, p.96); for the case of belief, see her (2000, p.6).

philosophical theses about causality emerging in her discussions, the theses seem to have the status of being the assumptions of her *opponent*, which she has taken up for the sake of argument.

A case in point is the assumption attacked in the last argument, that cause and effect must be ‘distinct existences’. This is explicitly ascribed by Anscombe (1983, p.106) to her opponent as part of what she calls the ‘modern’ view of causality. It is open to her to deny it, which she does (Anscombe 1974b, pp.148-51). And, to come back to Davidson and the historical narrative with which we started, we can push a similar line about the items on the charge sheet that he nailed to the anti-causalist’s door in ‘Actions, Reasons, and Causes’ (1963, pp.12-19). Nothing here sticks to Anscombe. To run down the list: she does not require causes to be events, and does not rule out states playing such a role, as we’ve already seen; she does not employ Hume’s dictum that causes must be ‘logically’ or ‘conceptually’ distinct from their effects, as we also just saw; she does not argue against the possibility of lawful psychophysical generalizations;<sup>133</sup> and her negative arguments do not rely on the thought that the epistemology of reasons is necessarily different from the epistemology of causes.<sup>134</sup>

None of this conclusively establishes that more substantive theses about causality are not at work deeper down in Anscombe’s argument. Her argument against narrow mental states strikes me as

---

<sup>133</sup> In a way, she even allows that we may devise them for mental causes. ‘In a way’, because she is happy to envision “a Humeian account of [mental causality] as far as concerns the outside observer’s recognition of the cause”; she denies that such a form of explanation could encompass the *agent’s* knowledge, since mental causes are ‘known without observation’ and so can’t be grounded in empirical generalizations (1963a, pp.10, 16; see also her 1948, pp.229-230; 1983, p.96).

<sup>134</sup> Though in this case this isn’t to say that there is no distinct epistemology involved. For one, in so far as Anscombe’s notion of ‘practical knowledge’ encompasses knowledge of one’s reasons, knowledge of reasons will involve a distinct type of knowledge from knowledge of mental causes. Still, it is worth stressing that Anscombe nevertheless allows mental causes to be known without observation or inference (see the previous note), and this is a far cry from what we might have expected if we read her with the expectation that the epistemology of causality would be central to understanding their role in reasons explanations *as opposed to* causal explanations. In any case, this particular epistemic feature of reasons plays no discernible role in the negative arguments I mention below.

purely metaphysical or purely semantical; but debates in that regions are routinely bound up with heavy premises about causality — unsurprisingly so, since the topic is the character of certain explanations.<sup>135</sup> Perhaps a more careful look will bring out that there is more at work in Anscombe’s claims about intentional states. Similarly, it may well be that a closer look at the metaphysical view which Anscombe helps herself to in her argument against the universality of wide causation — the view on which, as it were, the intending can *be* the doing — will itself bring to the surface substantive claims about causality.<sup>136</sup> In any case, whatever is at work in Anscombe’s anti-causalist arguments, I conjecture that it won’t be anything that requires us to step outside the minimal picture of causal explanation sketched earlier, and thus won’t be anything that stops us from seeing reasons explanations as a species of causal explanation. If the causal pluralist picture that I’ve drawn from Anscombe is a good interpretation of her view, then that much is just what we should expect. Anscombe’s own characterization of her negative arguments is that they show that the explanation of action by intention or reasons “does not properly come under my title ‘the causation of action’ as moderns, rather than Aristotelians, have understood the term ‘causation’” (1983, p.106). If we drop the thought that Anscombe’s discussion of causal explanation is hindered by stubborn and outmoded prejudices, and warm instead to the causal pluralist picture, then perhaps we can read *Intention* as speaking to just that topic, the causation of action.

---

<sup>135</sup> For instance, see Fodor’s (1991) classic argument for narrow content based on what “Uncle Hume taught us”, that causal powers can only explain their effects when contingently tied to them.

<sup>136</sup> Or about related concepts, such as the nature of change. This, I think is lurking in the background and foreground to various degrees in Rödl (2007, Ch.2), Hornsby (2012), and Ford (2014).



## BIBLIOGRAPHY

- Alvarez, M. 2001. "Letting Happen, Omissions and Causation", *Grazer Philosophische Studien*, 61: 63-81.
- . 2007. "The Causalist/Anti-Causalist Debate in the Theory of Action: What it is and why it matters", in Leist: 103-123.
- . 2012. "Agency and Two-Way Powers", *Proceedings of the Aristotelian Society*, 113: 101-21.
- Aguilar, J.H. & A.A.Buchareff (eds.), 2010. *Causing Human Actions: New Perspectives on the Causal Theory of Mind*, Cambridge, Mass.: The MIT Press.
- Anscombe, G.E.M. 1948. "A Reply to Mr. C.S. Lewis's Argument that 'Naturalism' is Self-Refuting", repr. in her 1981b: 224-232.
- . 1963a. *Intention*, Cambridge, Mass.: Harvard University Press, 2000, 2nd edition.
- . 1963b. "Events in the Mind", repr. in her 1981b: 57-63.
- . 1963c. "The Two Kinds of Error in Action", repr. in her 1981c: 3-9.
- . 1969. "On Promising and its Justice, and Whether it Needs be Respected *In Foro Interno*", in her 1981c: 10-21.
- . 1971. "Causality and Determination", repr. in her 1981b: 133-147.
- . 1974a. "Memory, 'Experience', and Causation", repr. in her 1981b: 120-130.
- . 1974b. "Times, Beginnings, and Causes", repr. in her 1981b: 148-162.
- . 1976. "The Question of Linguistic Idealism", repr. in her 1981a: 112-133.
- . 1978. "Will and Emotion", repr. in her 1981a: 100-107.
- . 1979a. "Chisholm on Action", repr. in Geach & Gormally (eds.), 2005: 77-87.
- . 1979b. "On *Humanae Vitae*", in Geach & Gormally (eds.), 2008: 192-98.
- . 1981a. *The Collected Philosophical Papers of G.E.M. Anscombe, Vol.I*, Oxford: Basil Blackwell.
- . 1981b. *The Collected Philosophical Papers of G.E.M. Anscombe, Vol.II*, Oxford: Basil Blackwell.

- . 1981c. *The Collected Philosophical Papers of G.E.M. Anscombe, Vol.III*, Oxford: Basil Blackwell.
- . 1981d. “A Theory of Language?”, in Geach & Gormally (eds.), 2011: 193-204.
- . 1982. “Action, Intention, and ‘Double Effect’”, repr. in Geach & Gormally (eds.), 2005: 207-226.
- . 1983. “The Causation of Action”, repr. in Geach & Gormally (eds.), 2005: 89-108.
- . 1989b. “Practical Inference”, repr. in Geach & Gormally (eds.), 2005: 109-147.
- . 2005. “Good and Bad Human Action”, in Geach & Gormally (eds.), 2005: 195-206.
- . 2008a. “On Being in Good Faith”, in Geach & Gormally (eds.), 2008: 101-112.
- . 2008b. “Sin: The McGivney Lectures”, in Geach & Gormally (eds.), 2008: 117-156.
- . 2011a. “Was Wittgenstein a Conventionalist?”, in Geach & Gormally (eds.), 2011: 217-230.
- . 2011b. “On Piety, or Plato’s *Euthyphro*”, in Geach & Gormally (eds.), 2011: 11-24.
- . 2015. “Belief and Thought”, in Geach & Gormally, 2015: 149-181.
- Audi, R. 1986. “Acting for Reasons”, repr. in A. Mele (ed.), 1997. *The Philosophy of Action*, Oxford: Oxford University Press, 75-105.
- Austin, J.L. 1966. “Three Ways of Spilling Ink”, repr. in J.O. Urmson & G.J. Warnock (eds.), *Philosophical Papers*, Second Edition, 1970, Oxford: Oxford University Press, 272-87.
- Ayers, M. 1972. “Some Thoughts”, *Proceedings of the Aristotelian Society*, 73: 69-86.
- Beebe, H. 2003. “Seeing Causing”, *Proceedings of the Aristotelian Society*, 103: 257-280.
- . 2009. “Causation and Observation”, in Beebe *et al.* (eds.), 2009: 471-497
- Beebe, H., C. Hitchcock & P. Menzies, (eds.), 2009. *The Oxford Handbook of Causation*, Oxford: Oxford University Press.
- Bermúdez, J.L. 2015. “Bodily Ownership, Bodily Awareness, and Knowledge Without Observation”, *Analysis*, 75: 37-45.
- Bishop, J. 1989. *Natural Agency*, Cambridge: Cambridge University Press.
- . 2011. “Review of Michael Thompson’s *Life and Action*”, *Ethics*, 122: 168-75.
- Bogen, J. 2008. “Causally Productive Activities”, *Studies in History and Philosophy of Science*, 39: 112-123.
- Boyle, M.. 2012. “Essentially Rational Animals”, in G. Abel & J. Conant (eds.), *Rethinking Epistemology*, Vol. 2, Berlin: DeGruyter, 395-427.
- Bratman, M. 1999. *Intentions, Plans, and Practical Reasons*, Stanford: CSLI Publications.

- . 2006. “What is the Accordion Effect?”, *Journal of Ethics*, 10: 5-19.
- . 2009. “Reflections on the Philosophy of Action”, in J. Aguilar & A. Buckareff (eds.), *Philosophy of Action: 5 Questions*, London: Automatic Press/VIP, 17-23.
- Broncano-Berrocal, F. 2015. “Luck as Risk and the Lack of Control Account of Luck”, *Metaphilosophy*, 46: 1-25.
- Burge, T. 1993. “Content Preservation”, *Philosophical Review*, 102: 457–488.
- . 2011. “Postscript: ‘Content Preservation’”, in his *Cognition Through Understanding, Philosophical Essays Volume 3*, 2013, Oxford: Oxford University Press, 254-284.
- Burgess, J. 2008. “When is Circularity in Definitions Benign?”, *Philosophical Quarterly*, 58: 213-43.
- Candlish, S. & N. Damnjanovic 2013. “Reasons, Actions and the Will: The Fall and Rise of Causalism”, in M. Beaney (ed.), *The Oxford History of Analytic Philosophy*, Oxford: Oxford University Press, 689-708.
- Carroll, J.W. 2009. “Anti-Reductionism”, in Beebe *et al.*, 2009: 279-298.
- Cartwright, N. 2004. “Causation: One Word, Many Things”, *Philosophy of Science*, 71: 805-819.
- Child, W. 1994. *Causality, Interpretation and the Mind*, Oxford: Clarendon Press.
- Coliva, A. 2012. “One Variety of Self-Knowledge: Constitutivism as Constructivism”, in A. Coliva (ed.), *The Self & Self-Knowledge*, Oxford: Oxford University Press, 212-42.
- Dancy, J. 1995. “Arguments from Illusion”, *Philosophical Quarterly*, 45: 421-38.
- . 2000. *Practical Reality*, Oxford: Oxford University Press.
- . 2008. “On How to Act — Disjunctively”, in A. Haddock & F. Macpherson (eds.), *Disjunctivism: Perception, Action, Knowledge*, Oxford: OUP. 262–79.
- . 2011. “Acting in Ignorance”, *Frontiers of Philosophy in China*, 6: 345-57.
- Danks, D., Rose, D., & E. Machery. (2013). “Demoralizing Causation”, *Philosophical Studies*, 2: 1-27.
- Danto, A. 1963. “What We Can Do”, *Journal of Philosophy*, 60: 435–445.
- . 1976. “Action, Knowledge, and Representation”, in M. Brand & D. Walton (eds.), *Action Theory*, Dordrecht: D. Reidel, 11-25.
- Davidson, D. 1963. “Actions, Reasons, and Causes”, repr. in his 1980: 3-19.
- . 1971. “Agency”, repr. in his 1980: 43-61.
- . 1973. “Freedom to Act”, repr. in his 1980: 63-81.
- . 1978. “Intending”, repr. in his 1980: 83-102.

- . 1980. *Essays on Actions and Events*, Oxford: Clarendon Press.
- . 1985. “Reply to David Pears”, in B. Vermazen & M.B. Hintikka (eds.), *Essays on Davidson*, Oxford: Clarendon Press, 211-215.
- Davis, W. 2005. “Reasons and Psychological Causes”, *Philosophical Studies*, 122: 51-101.
- Dennett, D. 1968. “Features of Intentional Actions”, *Philosophy and Phenomenological Research*, 29: 232-44.
- Donagan, A. 1987. *Choice*, Routledge & Kegan Paul: London.
- Donnellan, K. 1963. “On Knowing What I’m Doing”, *Journal of Philosophy*, 60: 401-9.
- Dowe, P. 2004. “Causes are Physically Connected to their Effects: Why Preventers and Omissions are not Causes”, in Hitchcock (ed.), 2004: 189-196.
- D’Oro, G. 2013. “From Anti-causalism to Causalism and Back: A History of the Reasons/Causes Debate”, in D’Oro, G. & C. Sandis. (eds.), 2013. *Reasons and Causes: Causalism and Anti-Causalism in the Philosophy of Action*, London: Palgrave Macmillan, 7-48.
- Driver, J. 2011. “Gertrude Elizabeth Margaret Anscombe”, *The Stanford Encyclopedia of Philosophy*, Winter 2014.
- Falvey, K. 2000. “Knowledge in Intention”, *Philosophical Studies*, 99: 21-44.
- Fodor 1991. “A Modal Argument for Narrow Content”, *Journal of Philosophy*, 88: 5-26.
- Ford, A. 2011. “Action and Generality”, in Ford *et al.* (eds.), 2011: 76-104.
- . 2014. “Action and Passion”, *Philosophical Topics*, 42: 13-42.
- . 2015. “The Arithmetic of *Intention*”, *American Philosophical Quarterly*, 52: 129-143.
- . 2016. “On What Is in Front of Your Nose”, *Philosophical Topics*, 44: 141–161.
- Ford, A., J. Hornsby & F. Stoutland. (eds). 2011. *Essays on Anscombe’s Intention*, Cambridge, Mass.: Harvard University Press.
- Geach, M. & L. Gormally (eds.). 2005. *Human Life, Action and Ethics*, Exeter: Imprint Academic.
- . 2008. *Faith in a Hard Ground: Essays on Religion, Philosophy and Ethics*, Exeter: Imprint Academic.
- . 2011. *From Plato to Wittgenstein*, Exeter: Imprint Academic.
- . 2015. *Logic, Truth and Meaning*, Exeter: Imprint Academic.
- Gettier, E. 1963. “Is Justified True Belief Knowledge?”, *Analysis*, 23: 121-3.
- Gibbons, J. 2001. “Knowledge in Action”, *Philosophy and Phenomenological Research*, 62: 579-600.

- . 2006. “Knowledge: The Priority Of”, *Encyclopedia of Philosophy*, Vol. 2, Detroit: Macmillan Reference, 86-90.
- . 2010. “Seeing What You’re Doing”, *Oxford Studies in Epistemology*, 3: 63-85.
- Ginet, C. 2002. “Reasons Explanations of Action: Causalist versus Noncausalist Accounts”, in R. Kane (ed.), *The Oxford Handbook of Free Will*, Oxford: Oxford University Press, 386-405.
- Godfrey-Smith, P. 2009. “Causal Pluralism”, in Beebe *et al.* (eds.), 2009: 326-340.
- Grice, H.P. 1971. “Intention and Uncertainty”, *Proceedings of the British Academy*, 56: 263–279.
- Grünbaum, T. 2009. “Anscombe and Practical Knowledge of What is Happening”, *Grazer Philosophische Studien*, 78: 41-67.
- . 2013. “Seeing What I Am Doing”, *Philosophy and Phenomenological Research*, 86: 1-24.
- Haddock, A. 2011. ““The Knowledge that a Man Has of His Intentional Actions””, in Ford *et al.*: 147-69.
- Hampshire, S. & H.L.A. Hart. 1958. “Decision, Intention, and Certainty”, *Mind*, 67: 1-12.
- Hanser, M. 1998. “Intention and Teleology”, *Mind*, 107: 331-401.
- . 2000. “Intention and Accident”, *Philosophical Studies*, 98: 17-36.
- Hempel, C. 1965. *Aspects of Scientific Explanation*, New York: Free Press.
- Hetherington, S. 2014. “Knowledge Can Be Lucky”, in Steup *et al.*: 164-76.
- Hitchcock, C. (ed.), 2004. *Contemporary Debates in Philosophy of Science*, Oxford: Blackwell.
- Hitchcock, C. & J. Knobe. 2009. “Cause and Norm”, *Journal of Philosophy*, 106: 587-612.
- Holton, R. 2009. *Willing, Wanting, Waiting*, Oxford: Clarendon Press.
- Hornsby, J. 1993. “Agency and Causal Explanation”, repr. in her 1997, *Simple Mindedness*, Cambridge, Mass.: Harvard University Press, 129-153.
- . 1999. “Anomalousness in Action”, in L.E. Hahn (ed.), *The Philosophy of Donald Davidson*, Chicago, Ill.: Open Court, 623-635.
- . 2007. “Knowledge in Action”, in Leist (ed.) 2007: 285-302.
- . 2008. “A Disjunctive Conception of Acting for Reasons”, in A. Haddock & F. Macpherson (eds.), *Disjunctivism: Perception, Action, Knowledge*, Oxford: OUP, 244–61.
- . 2011. “Actions in their Circumstances”, in Ford *et al.*, 2011: 105-27.
- . 2012. “Actions and Activity”, *Philosophical Issues*, 22: 233-45.

- . 2015. "Causality and 'the mental'", *Humana Mente*, 29; 125-140.
- Houlgate, L.D. 1966. "Mistakes in Performance", *Mind*, 75: 257-261.
- Hunter, D. 2015. "Davidson on Practical Knowledge", *Journal of the History of Analytic Philosophy*, 3: 1-19.
- Humberstone, I.L. 1997. "Two Types of Circularity", *Philosophy and Phenomenological Research*, 57: 249-80.
- Hursthouse, R. 1991. "Arational Actions", *Journal of Philosophy*, 88: 57-68.
- . 2000. "Intention", in R. Teichmann (ed), *Logic, Cause and Action*, Cambridge: Cambridge University Press, 83-105.
- Hyman, J. 1999. "How Knowledge Works", *Philosophical Quarterly*, 49: 433-451.
- . 2006 "Three Fallacies about Action", in C. Kanzian & E. Runggaldier, *Cultures: Conflict, Analysis, Dialogue: Proceedings of the 29<sup>th</sup> International Ludwig Wittgenstein Symposium*, Frankfurt: Ontos, 137-164.
- . 2010. "The Road to Larissa", *Ratio*, 23: 333-414.
- . 2011. "Acting for Reasons: Reply to Dancy", *Frontiers of Philosophy in China*, 6: 358-68.
- . 2013. "Voluntariness and Choice", *Philosophical Quarterly*, 63: 683-708.
- . 2014. "Desires, Dispositions, and Causal Chains", *Philosophy*, 89: 83-112.
- Jones, O.R. 1961. "Things Known Without Observation", *Proceedings of the Aristotelian Society*, 61: 129-50.
- Kim, J. 2007. "Causation and Mental Causation", in B.P. McLaughlin & J.D. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*, Oxford: Blackwell, 227-242.
- Kenny, A. 1975. *Will, Freedom and Power*, Barnes & Noble: New York.
- Lavin, D. 2013. "Must There Be Basic Action?", *Noûs*, 47: 273-301.
- . 2016. "Action as a Form of Temporal Unity: on Anscombe's *Intention*", *Canadian Journal of Philosophy*, 45: 609-629.
- Lawrence, G. 2004. "Reason, Intention, and Choice: An Essay in Practical Philosophy", *Royal Institute of Philosophy Supplement*, 54: 265-300.
- Leist, A. (ed.) *Action in Context*, Berlin: Walter de Gruyter.
- Lewis, D. 1986. "Causal Explanation", in his *Philosophical Papers, Vol. 2*: 214-240.
- Ludwig, K. 1992. "Impossible Doings", *Philosophical Studies*, 65: 257-281.
- Makin, S. 2000. "Causality and Determination", in Teichmann (ed.): 59-71.
- Marcus, E. 2012. *Rational Causation*, Cambridge, Mass.: Harvard University Press.

- Martin, C.B. 1971. "Knowledge Without Observation", *Canadian Journal of Philosophy*, 1: 15-24.
- McDowell, J. 1994. *Mind and World*, Cambridge, Mass.: Harvard University Press.
- . 2006. "Response to Dancy", in C. Macdonald & G. Macdonald (eds.), *McDowell and his Critics*, London: Blackwell, 134-141.
- . 2010. "What is the Content of an Intention in Action?", *Ratio*, 23: 415-432.
- . 2011a. "Anscombe on Bodily Self-Knowledge", in Ford *et al.* (eds.), 2011: 128-146.
- . 2011b. "Some Remarks on Intention in Action", The Amherst Lecture in Philosophy.
- . 2013. "Zum Verhältnis von rezeptivem und praktischem Wissen", *Deutsche Zeitschrift für Philosophie*, 3: 387-401.
- . 2013. "Acting in the Light of a Fact", in Bakhurst, D., Hooker, B., & M.O. Little (eds.), *Thinking About Reasons*, Oxford: Oxford University Press, 13-28.
- Mele, A. 2013. "Action, Explanation, and Causes", in G. Oro (ed.), *Reasons and Causes: Causalism and Anti-Causalism in the Philosophy of Action*, Palgrave-Macmillan, 160-74.
- Mele, A.R. & P.K. Moser 1994. "Intentional Action", repr. in A. Mele (ed.), 1997. *The Philosophy of Action*, Oxford: Oxford University Press, 223-55.
- Menzies, P. 2009. "Platitudes and Counterexamples", in Beebe *et al.* (eds.), 2011: 341-367.
- Moore, M. 2009. "The Nature of Singularist Theories of Causation", *The Monist*, 92: 3-22.
- Moran, R. 2001. *Authority and Estrangement*, Princeton, NJ: Princeton University Press.
- . 2004. "Anscombe on 'Practical Knowledge'", in J. Hyman & H. Steward (eds.), *Agency and Action*, Cambridge: Cambridge University Press: 43-68.
- Moran, R. & M. Stone. 2011. "Anscombe on Expressions of Intention: An Exegesis", in Ford *et al.* 2011: 33-75.
- Müller, A. 2011. "Backward-Looking Rationality and the Unity of Practical Reason", in Ford *et al.* 2011: 242-69.
- Newstead, A. 2009. "Interpreting Anscombe's *Intention* §32ff", *Journal of Philosophical Research*, 34: 157-76.
- Neta, R. 2009. "Mature Human Knowledge as a Standing in the Space of Reasons", *Philosophical Topics*, 37: 115-32.
- O'Brien, L. 2003. "On Knowing One's Own Actions", in Roessler & Eilan (eds.), 2003: 358-382.
- . 2007. *Self-Knowing Agents*, Oxford: Oxford University Press.
- O'Shaughnessy, B. 2003. "The Epistemology of Physical Action", in Roessler & Eilan (eds.), 2003: 345-357.

- Özaltun, E. “Practical Knowledge of What Happens: A Reading of §45”, *Klesis*, 35: 52-73.
- Paul, S. 2009a. “How We Know What We’re Doing”, *Philosophers’ Imprint*, 9: 1-24.
- . 2009b. “Intention, Belief, and Wishful Thinking: Setiya on ‘Practical Knowledge’”, *Ethics*, 119: 546-557.
- Peacocke, C. 2003. “Action: Awareness, Ownership, and Knowledge”, in Roessler & Eilan (eds.), 2003: 94-110.
- Pears, D. 1964. “Predicting and Deciding”, in his *Questions in the Philosophy of Mind*, New York: Barnes & Noble: 13-38.
- . 1975. “Wanting and Intending”, in his *Questions in the Philosophy of Mind*, New York: Barnes & Noble: 39-55.
- Pickard, H. 2004. “Knowledge of Action Without Observation”, *Proceedings of the Aristotelian Society*, 105: 205-230.
- Pouivet, R. 2008. *After Wittgenstein, St. Thomas*, tr. M.S. Sherwin OP, Indiana: South Bend.
- Pritchard, D. 2014. “Knowledge Cannot Be Lucky”, in Steup *et al.*: 153-65.
- Psillos, S. 2008. “Causal Pluralism”, in R. Vanderbeeken & B. D’Hooghe (eds.), *Worldviews, Science and Us: Studies of Analytical Metaphysics: A Selection of Topics from a Methodological Perspective*. Singapore: World Scientific Publishers, 131-151.
- Rödl, S. 2007. *Self-Consciousness*, Cambridge, Mass.: Harvard University Press.
- . 2011. “Two Forms of Practical Knowledge and Their Unity”, in Ford *et al.*, 2011: 211-241.
- Roessler, J. 2010. “Agent’s Knowledge”, in T. O’Connor & C. Sandis (eds.), *A Companion to the Philosophy of Action*, West Sussex: Wiley-Blackwell: 236-243.
- Roessler, J. & N. Eilan (eds.), 2003. *Agency and Self-Awareness*, Oxford: Clarendon Press.
- Rota, M. 2009. “An Anti-Reductionist Account of Singular Causation”, *The Monist*, 19: 133-152.
- Schafer, K. 2013. “Knowledge and Two Forms of Non-Accidental Truth”, *Philosophy and Phenomenological Research*, 89: 373-93.
- Schaffer, J. 2004. “Causes Need Not be Physically Connected to their Effects: The Case for Negative Causation”, in Hitchcock (ed.): 197-216.
- Schwenkler, J. 2015. “Understanding ‘Practical Knowledge’”, *Philosophers’ Imprint*, 15: 1-32.
- Sehon, S. 2005. *Teleological Realism*, Cambridge, Mass.: Bradford Books: MIT Press.
- Setiya, K. 2007. *Reasons Without Rationalism*, Princeton, NJ: Princeton University Press.
- . 2008. “Practical Knowledge”, *Ethics*, 118: 388–409.



- . 2009. “Reasons and Causes”, *European Journal of Philosophy*, 19: 129-157.
- . 2015. “Intention”, *The Stanford Encyclopedia of Philosophy*, Summer 2015.
- . 2017. “Anscombe on Practical Knowledge”, in his *Practical Knowledge: Selected Essays*, Oxford: Oxford University Press, 156-68.
- Sliwa, P. 2017. “On Knowing What’s Right and Being Responsible for It”, in P. Robichaud & J.W. Wiland (eds.), *Responsibility: The Epistemic Condition*, Oxford: Oxford University Press, 127-45.
- Slote, M. 1979. “Assertion and Belief”, in J. Dancy (ed.), *Papers in Language and Logic*, University of Keele Library, 177-91.
- Small, W. 2012. “Practical Knowledge and the Structure of Action”, in G. Abel & J. Conant (eds.), *Rethinking Epistemology, Vol. 2*, Berlin: DeGruyter, 133-227.
- Smith, M. 2010. “The Standard Story of Action: An Exchange (1)”, in Aguilar & Buchareff (eds.), 2010: 45-56.
- . 2012. “Four Objections to the Standard Story of Action (and Four Replies)”, in *Philosophical Issues*, 22: 387-401.
- Snowdon, P. 2011. “Perceptual Concepts as Non-Causal Concepts”, in J. Roessler, H. Lerman & N. Eilan (eds.), *Perception, Causation, and Objectivity*, Oxford: Oxford University Press, 121-138.
- Sosa, E. & M. Tooley. 1993. “Introduction”, in their *Causation*, Oxford: Oxford University Press, 1-32.
- Stein, N. 2016. “Causes and Conditions”, *Noûs*, 50: 465-489.
- Steup, M., Turri, J., & E. Sosa (eds.), *Contemporary Debates in Epistemology*, 2<sup>nd</sup> Edition, Wiley Blackwell: West Sussex.
- Steward, H. 1997. *The Ontology of Mind*, Oxford: Clarendon Press.
- Stoecker, R. 2013. “Action Explanation”, in E. Lepore & K. Ludwig (eds.), *A Companion to Donald Davidson*, Chichester: John Wiley & Sons, 15-25.
- Stout, R. 2007. “Two Ways to Understand Causality in Agency”, in Leist, 2007: 137-153.
- . 2010a. “What Are You Causing in Acting?”, in J. Aguilar & A. Buckareff (eds.), *Causing Human Actions*, Cambridge, Mass.: MIT Press, 101-14.
- . 2010b. “Deviant Causal Chains”, in T. O’Connor & C. Sandis (eds.), *A Companion to the Philosophy of Action*, Chichester: John Wiley & Sons, 159-65.
- Stoutland, F. 2011. “Summary of Anscombe’s *Intention*”, in Ford *et al.*, 2011: 23-32.
- Strawson, P.F. 1985. “Causation and Explanation”, in B. Vermazen & M. Hintikka (eds.), *Essays on Davidson: Actions and Events*, Oxford: Oxford University Press, 115-135.

- Tanney, J. 2009. "Reasons as Non-Causal, Context-placing Explanations", in C. Sandis (ed.), *New Essays in the Explanation of Action*, New York: Palgrave Macmillan, 94-111.
- Teichmann, R. (ed.), 2000. *Logic, Cause & Action*, Cambridge: Cambridge University Press.
- Teichmann, R. 2008. *The Philosophy of Elizabeth Anscombe*, Oxford: Oxford University Press.
- . 2013. "The Importance of the Past", *Philosophy*, 88: 115-31.
- . 2014. "The Voluntary and the Involuntary: Themes from Anscombe", *American Catholic Philosophical Quarterly*, 88: 465-86.
- Tenenbaum, S. 2007. *Appearances of the Good*, Cambridge: Cambridge University Press.
- Thomson, J. J. 1987. "Verbs of Action", *Synthese*, 72: 103-22.
- Thompson, M. 1992. *Some Remarks on the Role of Generality in Practice*, Dissertation, UCLA.
- . 2008. *Life and Action*, Cambridge, Mass.: Harvard University Press.
- . 2011. "Anscombe's *Intention* and Practical Knowledge", in Ford *et al.*, 2011: 198-210.
- . 2013. "Forms of nature: 'first', 'second', 'living', 'rational', and 'phronetic'", in G. Hinrichs & A. Honneth (eds.), *Freiheit: Stuttgarter Hegel-Kongress 2011*, Frankfurt: Klostermann, 701-35.
- Tooley, M. 1987. *Causation: A Realist Approach*, Oxford: Clarendon Press.
- Turri, J. 2009. "The Ontology of Epistemic Reasons", *Noûs*, 43: 490-512.
- . 2012. "Reasons, Answers, and Goals", *Journal of Moral Philosophy*, 491-99.
- Unger, P. 1975. *Ignorance, A Case for Scepticism*, Oxford: The Clarendon Press.
- Velleman, D. 2007. *Practical Reflection*, USA: CSLI Publications.
- Vermazen, B. 1993. "The Belief Requirement on Intention", *Analysis*, 53: 239-42.
- Vogler, C. 2016. "Nothing Added: *Intention* §§19 and 20", *American Catholic Philosophical Quarterly*, 90: 229-247.
- Wiland, E. 2012. *Reasons*, London: Continuum.
- Wilson, G. 1989. *The Intentionality of Human Action*, Palo Alto: Stanford University Press.
- . 1997. "Reasons as Causes for Action", in G. Holmström-Hintikka & R. Tuomela (eds.), *Contemporary Action Theory, Vol.1*: 65-82.
- Wilson, J. 2017. "Determinables and Determinates", *The Stanford Encyclopedia of Philosophy*, Spring 2017.
- Williams, M. 2015. "What's So Special About Human Knowledge?", *Episteme*, 12: 249-68.

- Williamson, T. 2000. *Knowledge and its Limits*, Oxford: Oxford University Press.
- Wittgenstein, L. 1963. *Philosophical Investigations*, tr. G.E.M. Anscombe, New York: MacMillan.
- . 1967. *Zettel*, Los Angeles: University of Berkeley Press.
- Wolfson, B. 2012. “Agential Knowledge, Action and Process”, *Theoria*, 78: 326–357.
- Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.
- Yablo 1997. “Wide Causation”, *Noûs*, 31: 251-81.
- Yamada, M. 2011. “Getting it Right by Accident”, *Philosophy and Phenomenological Research*, 83: 72-105.