

An Anscombean Approach to Animal Agency

Luke Keanu Cash

St. John's College

This dissertation is submitted for the degree of Doctor of Philosophy

September, 2018

An Anscombean Approach to Animal Agency

Luke Keanu Cash

Abstract:

The ultimate aim of this thesis is to explain how the theory of action found in Anscombe's *Intention* can be modified to deliver a plausible account of non-human animal agency (henceforth, animal agency). More specifically, it is an attempt to develop her account in a way that respects the Aristotelian insight that animals act in ways that differ, in certain fundamental respects, from the processes of growth and self-maintenance found in plants, on the one hand, and the self-conscious actions characteristic of mature human beings, on the other.

The negative aim is to show that the theory of action that constitutes the received backdrop in the study of animal minds is ill-suited for the task. This is what I call the Standard Approach to Animal Agency and, despite its widespread acceptance in comparative psychology, cognitive ethology, and the philosophy of animal minds, I argue that it faces serious problems.

This thesis divides roughly into two halves corresponding to these respective aims. In the first half I argue against the Standard Approach. Amongst other things, I suggest that the theory suffers from a tendency to take the notion of action for granted. The result is an oversimplified metaphysics that is ill-prepared to account for the fact that the activities characteristic of animal life are instrumentally structured processes.

In the constructive half of the thesis I develop an Anscombean alternative that takes the structure of action as its starting point. On this view, expressions of animal agency are understood as a distinctive kind of material process. After explaining Anscombe's account of intentional action, I adapt and develop these ideas into a theory of animal agency.

Preface

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

It does not exceed the prescribed word limit for the relevant Degree Committee.

Acknowledgements

I would not have been able to complete this thesis without the help of my supervisor, Michael Potter. His willingness to take my ideas seriously enough to eviscerate them has forced me to face philosophical problems with the sense of perplexity that they deserve. And his depth of knowledge, patience, and good humour has set an example that I will always remember.

I am also indebted to my three shadow supervisors, Tim Button, Tim Crane, and Arif Ahmed. All have supported me throughout different stages in my thesis and I count myself privileged to have been taught by each of them. Arif in particular kickstarted an otherwise stagnant writing process and provided much needed direction.

For financial support, I am grateful to the Cambridge Faculty of Philosophy, The Royal Institute of Philosophy, and especially the benefactors of St John's College Cambridge. I could never have hoped to complete my work without their support.

I have been lucky to know some brilliant philosophers whose willingness to read, comment, and listen to my work has — I am sure — improved it. Thanks in this regard to Cathy Mason, Charlie Temperley, Thomas Schindler, Dan Williams, Raphael Woolf, Jane Heal, Danièle Moyal-Sharrock, Nadav Matalon, Alexander Greenberg, and Lucy Campbell. Thanks also to the amazing graduate secretaries, Zoe Walker-Fagg and Charlie Evans, as well as the legendary Heather Sanderson. A special debt of gratitude is due to Matt Dougherty and Fredrik Nyseth.

Both read and commented on my work, but, more importantly, became two of my best friends over the course of this process.

This reminds me of the other wonderful friends that I have made in the UK: all of the above are included, but also Becca, Lucy, German Stefan #1, German Stefan #2, Adam, Luke, Helen, Stephen, Robin, Ming, Frazier, Robert, Claudia, and coffee-bug Julian and Anne. The UK is a foreign country for me, but they have made it feel like home.

Finally, to my family I feel as though I owe everything. To Jordan: even though I don't always answer your calls, I hope you know how loved they make me feel. To Marty: the support and encouragement that you give me fills me with confidence and pride. To Adam: I am so happy that I picked up the phone and called you 10 years ago. To Patricia: I miss you so much. To mom and tutu: I don't know how to express in words how much you mean to me; this thesis is dedicated to you both. To Alice: I love you more than anything; writing this thesis sucked, but I met you and I wouldn't trade that for the world.

Table of Contents

1	Introduction	
1.0	Forms of agency	1
1.1	Recent work	5
1.2	The Causal Theory of Action	8
1.3	The Standard Approach to animal agency	12
1.4	An Anscombean alternative	17
2	Objective Representation and Propositional Content	
2.0	Preliminaries	22
2.1	The problem of stimulus determination	24
2.2	Ways to triangulate	27
2.3	Propositional content	33
2.4	Practical inference	37
2.5	Summary	40
3	Taking Action For Granted	
3.0	Preliminaries	41
3.1	<i>Intention</i> , deflated	43
3.2	The objects of desire	47
3.3	On the significance of the foregoing argument	51
3.4	The consequences of taking action for granted	54
3.5	Summary	60
	Appendix: The argument of <i>Intention</i> §19	60
4	Action and Structure	
4.0	Preliminaries	66
4.1	The temporal unity of events	70
4.2	Process and law	73
4.3	Anscombe on practical knowledge	76
4.4	The material influence of practical knowledge	82
4.5	Solving the problem of temporal unity	85
4.6	Summary	92
5	Animal Activity	
5.0	Preliminaries	93

5.1	Facts, reasons, and painful experiences.....	95
5.2	Emotions, incentives, and spatially significant auditory information.....	100
5.3	Instinct and skill.....	105
5.4	Some applications.....	108
5.5	A theory of animal activity	119
5.6	Summary	124

Introduction

“We have many vocabularies for describing nature when we regard it as mindless, and we have a mentalistic vocabulary for describing thought and intentional action; what we lack is a way of describing what is in between.”¹

1.0 Forms of agency

It is tempting to treat agency as a uniquely human phenomenon — or at any rate to restrict the notion in such a way that it is difficult to make sense of as applying to anything other than enlanguaged creatures. Michael Thompson, for instance, hazards the following definition:

An agent, we may say, is something that operates on the strength of practical reasons or thoughts or considerations. It is something that can be viewed as doing one thing for the sake of another, according to concepts, or equivalently, I think, simply as a realizer of concepts. I see a process as a phenomenon of agency, in this sense, when I see the concept through which I describe or represent the process as itself at work in the genesis of the process I describe or represent.²

If this is an accurate description of the kind of agency that human beings enjoy, then I suppose that at least this much is true: most non-human animals (henceforth animals) are not agents in the same sense as human beings are agents. But it would surely be an unnecessary constriction of thought to restrict the notion of agency to its specifically self-conscious form.³ When a wasp gets splattered against the wall it is a mere patient; but when it kills a spider, drags it back to its den, and then lays its eggs inside the corpse, it is the agent of all these deeds. Chemicals carry

¹ Davidson (1999b, p. 11)

² (2004, p. 352)

³ I do not mean to suggest that Thompson is committing this mistake — his definition is just an instance of the perfectly legitimate practice of reserving a term for special use as the situation demands. But I do think that the widespread practice of using ‘agency’ to refer only to its specifically self-conscious form has made it easy for philosophers to *ignore* questions pertaining to animal agency.

within themselves the potential to enact change, and even plants are agents in some sense: sprouting a new shoot is part of an active process in a way that being uprooted (say) is not. We should not hesitate to admit, I think, that agency is a genus that encompasses many different species and sub-species.

My topic in this thesis is not agency in general, of course, but its specifically animal determination. Nevertheless, I think that it will be useful to begin with a brief overview of the hierarchical order of forms that I am envisioning.

* * *

Suppose that an asteroid strikes the earth and destroys a shed; it creates a loud noise which, in turn, scares the neighbourhood cat; the frightened cat runs into the street and causes an accident. No doubt we could continue to expand our description of the proceedings further, but let us stop here and ask: What did the asteroid *do*? Well, it struck the earth, destroyed a shed, made a loud noise, frightened a cat and, perhaps also, caused an accident. It was the ‘agent’ of these events in the perfectly harmless sense that it is the subject an active verb, the reference of which has describable effects. With a deliberate air of paradox, let us say that the asteroid’s role in this chaotic welter of happenings is that of an *inanimate agent*.

An inanimate agent is a *mere* source of change. It has no perspective on what it does, and can neither succeed nor fail; it cannot *attempt* to do anything; it has no goals. As such, questions like ‘Did the asteroid mean to frighten the cat or was that a mistake on its part?’ are absurd: it makes no sense to ask whether destroying the shed was a means to the asteroid’s end of making a loud noise, as opposed to, say, an unintended side-effect of striking the earth. Again, the distinction cannot be drawn. The series of true descriptions that we can give of the event are all on a level — they are *mere* descriptions of consequences, none of which could sensibly be classified as a means, or an end, or an unintended side-effect.

But now consider the humble oak. It grows a new branch, expands its canopy, extends its shadow, and blocks the sun from reaching the neighbouring sycamores. Here it would seem that we *can* distinguish between descriptions of means, ends, and incidental side-effects: the tree grows a new branch, not because it is extending its shadow but, rather, because it is expanding its canopy. The descriptive context is opaque. It is also normative: in expanding its canopy it is successfully fulfilling a biological function.

If this is right, then oaks are active in a way that asteroids are not: not only do they cause things to happen, they do so ‘for reasons’ in the minimal sense that their growth can be explained in terms of potentially unrealised ends. But these reasons are not their own; indeed,

as Dennett nicely puts it, they are nobody's reasons.⁴ For notice that the reasons which explain this particular oak's activity equally explain the activity of every other member of the species. Individual oaks pursue ends only in virtue of the fact that they are a particular kind of organism, one for which there are certain tendencies of normal genesis and operation; they are not reasons which belong to them as individuals who must figure out how to act in relation to the contingencies of the here and now.⁵

Moving upward along the *scala naturae* we reach the animals, my topic. I shall not say too much here, but let me lay out a few general features of animal life that will be significant in what follows.

Animals differ from plants in at least two philosophically significant ways. First, they possess the power of locomotion. Animals can move from one location to another and, importantly, these movements typically exist in the context of a wider pattern of activity — hunting, foraging, playing, fishing, mating, and so on. The second point is related. Unlike a plant which merely responds to proximal stimulation, an animal's perceptual faculties allow it to take up a perspective on the world and to adapt its behaviour in relation to it. And while animals share in the general needs characteristic of all living things — to acquire nutrients, to grow, and to reproduce — the power of perception allows the animal to focus its attention on more specific ends: its desires. As Matthew Boyle puts it,

[An animal's] capacities for perception and desire transform its mode of being alive precisely because they make this possible: they open animal life, not merely to the causal influence of present circumstances in the form of triggering, hindrance, or facilitation, but to the kind of influence that enters into the constitution of what the subject is doing. Thus an animal can try *to get that object*, or do something *in order to avoid this obstacle*.⁶

Boyle's remark usefully emphasises the fact that an animal's desiderative and perceptual faculties allow the particularities of the here and now to enter into the description of what happens when it acts. A muskox might pursue a particular mate, for instance, and a wolf might isolate a specific muskox from the heard and try to kill it. The 'reasons' for which an animal acts are not of the same general kind as those which explain the processes of growth and self-maintenance found in plants, but are, rather, tied to the animal's limited — and fallible — perspective on things.⁷

⁴ See, for example, Dennett (2013, pp. 57-58)

⁵ On this point see also Boyle (2012, p. 411)

⁶ (2012, p. 412)

⁷ This conception of animal life is, of course, ancient in origin. Its roots lie in Aristotle's *De Anima* and *De Motu Animalium*. See especially Aristotle (2016; Book III.9, 432a15-17) and (1984b; 7, 701a1.30-35).

Or so one might think. Though initially compelling, this naïve⁸ description is likely to strike many as overly simplistic. On the one hand, the idea that animals and plants are really that different from plants might be challenged. “We are not surprised”, Dennett writes, “to learn that trees that are able to sense the slow encroachment of green-reflecting rivals shift resources into growing taller faster, because that’s the smart thing for a plant to do under those circumstances.”⁹ Moreover, and from the other direction, many animal activities are notoriously inflexible. Frogs, for example, will snap at anything small-dark-and-moving that comes within range be it a fly, a bee-bee, or any number of other things. And is this really so different from the tropistic sort of response characteristic of the Venus fly-trap?¹⁰

On the other hand, some will take issue with the fact that this brief sketch neglects the fact that many animals are capable of impressive feats of practical reasoning and, hence, overestimates the distance between human and animal agency. Indeed, many would want to argue that *unless* animals are capable of practical thought then the kind of agency that they display is nominal.

These themes will all be prevalent throughout this thesis. Ultimately, my aim is to articulate a framework that avoids collapsing the distinction between the bare forms of responsiveness characteristic of plants and the often intricate activities of animals, whilst at the same time avoiding the temptation to over-intellectualise animal minds. More precisely, I want to describe a way of thinking about the agency of organisms that are capable of perception, appetite, instinct, and imagination, but which do not necessarily possess the power of (conceptual) thought. It is, I suppose, a defence of naïveté. But let me emphasise that it is no part of my project to *deny* that animals think (or have concepts, or can mindread, or are self-conscious, etc.); nor do I want to insist that *every* animal is capable of perception and appetite (as Aristotle seems to have thought).¹¹ Rather, I simply want to make space for a conception of agency that sits somewhere in-between the functional responsiveness characteristic of plants and the self-conscious, intentional actions characteristic of human beings. I call it animal agency because I think that Aristotle was *roughly* correct to draw the boundaries where he did. But if certain species are capable of more or less than my account can explain, then that is ultimately fine by me — extending the account to accommodate their activities will in that case be a task for future work.

I shall describe the structure of my thesis in more detail below (§1.4). First, though, let me take the opportunity to discuss some recent work on animal agency (§1.1), as well as the way in which this topic is more typically broached (§§1.2-1.3).

⁸ I use the term ‘naïve’ throughout this thesis to mean, roughly, pre-theoretical. I do not use it pejoratively.

⁹ (2009, pp. 342-343)

¹⁰ Both of these examples are discussed in the body of the thesis (in sections §2.2 and §5.4).

¹¹ See, e.g., Aristotle (2016; Book III, 10.433b27-29).

1.1 Recent work

The first piece of work that I want to mention is that of Tyler Burge. In his ‘Primitive Agency and Natural Norms’, Burge emphasises, as I have above, the need to accommodate forms of agency that sit in-between the self-conscious agency of mature human beings and the functional responsiveness characteristic of plants. To this end, he spends a good deal of time spelling out distinctions between these lower forms, drawing our attention to distinctions between tropisms, taxes, kineses, and orientations, all of which are types of movement that are available to organisms that lack perceptual faculties.

These distinctions are not themselves relevant for my purposes, but more interesting is Burge’s claim that many of these forms of movement are expressions of agency. Burge’s judgments in this regard are mainly based on intuitions about cases, but he does gesture towards something resembling a principle. An action, he suggests, is an instance of “*functioning, coordinated* behavior by the *whole organism*, issuing from the individual’s *central behavioral capacities*, not purely from sub-systems.”¹²

Unfortunately, Burge does not explain the key notions (aside from *functioning*) deployed in this formula in detail, instead relying on examples to get the idea across. Shivering, coughing, and sneezing are instances of behaviour that are not actions, he says, because they are “reflexive, peripheral, processes.”¹³ By contrast, a fish or a paramecium’s swimming through water does count as action because “A significant contribution to the movement comes from within. And the movement often involves whole-organism coordination between central capabilities and peripheral systems.”¹⁴ Again, though, it is unclear to me what sense of coordination Burge is relying on. One plausible way to understand it, I think, is to interpret it as applying only to cases in which the relation between sensory stimulation and action is not invariant: a perception of a certain type does not produce a single form of response but can be put to multiple uses depending on the animal’s needs, desires, and current activities — action arises out of a coordination between them. But this is apparently not the idea that Burge has in mind. For he writes:

Various types of instinctive behavior are inflexible and chain-reflexive, but still count as action. The male grouse will copulate with a stuffed grouse, male grouse, or dead grouse, if it sees any of these as assuming the relevant female mating position. The male grouse’s copulation activity is released by a single stimulus or single perception. The instinctive behavior does not derive from

¹² (2009, p. 260)

¹³ (2009, p. 261)

¹⁴ (2009, p. 260)

an inability to distinguish visually between the sexes, or between live and dead grouses. It is just that the instinctive behavior overrides these distinctions, once the key stimulus is received.¹⁵

I have quoted this example at length because I think that there is something wrong with thinking of this behaviour as *coordinated*, in any significant sense, by the animal in question. Indeed, there is something especially uncanny about examples like this given that the grouse apparently *is* able to distinguish between these various kinds of thing (a claim which I will simply assume is empirically accurate for the purposes of discussion). For it signals a complete *lack* of control on the part of the animal and — I would have thought — an *inability* to coordinate what it does with what it recognises. Indeed, it would seem that when a male grouse perceives a dead grouse positioned thus and so, its situation is not too dissimilar from that of a driver who *sees* that he is about to crash, but is so paralysed by fright that he is unable to pump the breaks.

In any case, the point that I want to make is just that the notion of coordination which Burge relies on needs considerable development before it can be put to theoretical use. If cases like the grouse are supposed to be paradigmatic *exemplars* of that notion, however, then I think that we should be sceptical as to its suitability to illuminate more complex forms of animal activity.

A more promising approach to animal agency is articulated by Helen Steward. Unlike Burge, Steward offers a detailed characterisation of agency that, she thinks, applies to many animals. According to Steward:

- (i) an agent can move the whole, or at least some parts, of something we are inclined to think of as *its* body;
- (ii) an agent is a centre of some form of subjectivity;
- (iii) an agent is something to which at least some rudimentary types of intentional state (e.g., trying, wanting, perceiving) may be properly attributed;
- (iv) an agent is a settler of matters concerning certain of the movements of its own body [...]; i.e., the actions by means of which those movements are effected cannot be regarded as the inevitable consequence of what has gone before.¹⁶

I think that conditions (ii) and (iii) are plausible. A creature that lacks a perspective on things cannot adapt its behaviour so as to accord with the contingencies of the here and now; if it responds to them, then it does so in the sense in which a sunflower bends towards the sun: it is not in perceptual contact with anything. Similarly, I think that Steward is right to emphasise the importance of primitive forms of intentional states, especially desire. For as the example of the

¹⁵ (2009, p. 265) The example derives from Tinbergen (1969, p. 36).

¹⁶ Each point is taken from Steward (2012, pp. 71-72).

grouse nicely demonstrates, some form of coordination between what the animal has an appetite for, what it perceives, and what it does is essential if we are to regard the animal as in control of its activities. When perception and action are invariant then we have at best limiting cases of animal agency.¹⁷

What of (i)? The issue here is delicate. Everyone will agree, I take it, that an agent can in general move her body and that many animals do this. The moot question is whether the agent's role can be explained in other terms — whether, for example, the claim that 'I looked at my watch' can be understood in terms of some event (such as the onset of a desire to know the time) causing me to look at my watch. Ultimately, I will argue that Steward is correct to resist attempts to explain the agent's role away. But my reason for holding this is different from hers and, indeed, I think that her focus on bodily movement (something that she shares with most philosophers of action) serves to blinker her to the *complexity* of animal activities. In my view, which I shall just indicate here, the capacity to move one's body is a *derivative* phenomenon, one that is (conceptually) less fundamental than the capacity to, say, hunt, play, and groom one's kin.¹⁸

Although this difference may seem minor, it is connected to what is perhaps a deeper divergence, namely the fact that Steward wants to argue that agency is incompatible with determinism. This brings us to (iv). Since, in Steward's view, nothing can be conceived to be genuinely active if its bodily movements are determined by what has come before, her focus is on *denying* that a deterministic account of an animal's bodily movements can be given at *any* level of description.

This is problematic for two reasons. In the first place, it obviously puts her theory at the mercy of the soundness of this denial. But secondly, it narrows her focus into a beam that is directed precisely at the difference between, say, someone raising their arm and their arm rising. This emphasis on an agent's being able to *settle* (as she puts it) matters concerning her own body means that what happens *beyond* the animal's body is left entirely untheorized — it is treated as mere happening like any other. As will become clear in what follows, I dispute this assumption. But, in any case, the crucial point at this stage is just that I shall not take a stand on the question of whether determinism is compatible with agency. Rather, I will argue that there is an active role to be played by the animal *even if* its movements are, at some level of description, settled by what has gone before.

¹⁷ I discuss these matters in greater depth in §§5.4-5.5.

¹⁸ I elaborate upon this cryptic idea in chapter 4.

1.2 The Causal Theory of Action

Burge and Steward are in the minority amongst both philosophers and cognitive scientists in treating animal agency as a topic in its own right. The more usual method is simply to ask whether a certain form of behaviour is caused by the sophisticated representational states (beliefs, desires, and intentions) which, it is supposed, distinguish the intentional actions of human beings from other events in which we partake (slips, trips, reflex actions, etc.). If the behaviour isn't so caused then it is to be understood on the model of a 'mere happening': it is not so much something that the animal does as something that happens to it.

There are two assumptions here. The first is that if animals are agents (in any significant sense) they must be capable of acting in the same general way as human beings act. The second is that the way to think about 'the way that human beings act' is through the lens provided by the Causal Theory of Action (henceforth CTA). In this section, then, I'll give a brief overview of CTA — as I propose to understand it — before turning to its reception in the literature on animal minds in the next.

CTA comes in a variety of forms, but most take their departure from Donald Davidson's work, and it will be helpful to begin by describing the central core of his account. It is refreshingly simple in outline; it consists of the following four theses:

- (1) An action is an event — specifically a bodily movement — which is intentional under some description.
- (2) An event is intentional under some description in virtue of the fact that it can be explained in terms of the agent's reasons (call this form of explanation 'rationalisation').
- (3) An agent's (primary) reasons are a pair of mental states (paradigmatically: beliefs and desires).
- (4) Rationalisation is a type of causal explanation.

While the theory is easy to state, its interpretation is not entirely straightforward. Points (1) and (4) in particular require further comment.

(4) should not be confused, as it tends to be, with the simple claim that reasons *are* causes.¹⁹ Although Davidson does sometimes speak in this way, his considered position is that *events*, not states, enter into causal relations, and from (3) it straightforwardly follows that reasons cannot

¹⁹ This point is nicely emphasised by J. Hornsby (1999, pp. 628-631)

enter into causal relations.²⁰ But the more important point is this: an agent's reasons belong to the realm of the intensional: their force depends precisely on how the events that they are meant to explain are described. Causal relations, by contrast, hold between events independently of our descriptions; their *efficacy* has nothing to do with the way in which the events are represented (according to Davidson, at any rate).²¹

The thought contained in (4), then, is not that reasons are causes, but rather that rationalisation is a form of causal *explanation*. Davidson's idea is that when we rationalise an action (a type of event), the *explanans* are a pair of mental states which are associated with — but not identical to — an event that *is* the cause of the action. This might be the onset of one of the states itself, or an intention formed as a result of a practical inference, or something else entirely. The important point is just that, whatever the underlying cause, the action only has its specific effect as a result of the agent's mental states.

What of (1)? There are two parts to it and they are connected: the first is the claim that we can distinguish a class of events that are actions, the second is a thesis about the members of that class. I take each in turn.

Not every description of an action is a description under which it is intentional. A man may be (i) pumping water intentionally and (ii) pumping poisoned water unintentionally by (iii) operating a pump handle. Now Davidson takes over from Anscombe the idea that there is just one action here that has been described in three ways.²² If this is right — and I will follow the action-theoretic tradition in assuming that it is — then being intentional is not a property of actions *per se*, but only of actions under a given description. And from this it follows straightforwardly that there cannot be a class of intentional actions: if there were, every action both would, and would not, belong to the class, which is obviously absurd.

However, Davidson suggests that we *can* distinguish a class of events that are actions, since every action is, he thinks, intentional under *some* description.²³ This means that “although the

²⁰ See, for example, Davidson (1994, p. 287): “Primary reasons [...] are certainly not events [...] Beliefs and desires are not changes. They are states; and since I don't think that states are entities of any sort, and so are not events, I do not think beliefs and desires are events.”

²¹ See, e.g., Davidson (1967).

²² This is not to say that there are no differences in the way that Anscombe and Davidson understand this claim. For Anscombe the various true descriptions of an intentional action are different descriptions of an instrumentally structured process and no one description is in any sense privileged. As she puts it, “The proper answer to ‘What is the action, which has all these descriptions?’ is to give one of the descriptions. Any one, it does not matter which; or perhaps it would be best to offer a choice, saying ‘Take whichever you prefer.’” (1979, p. 220) For Davidson, as we shall see, it is the agent's bodily movements are especially important: in each case it is they that are re-described. On these two ways of understanding the idea see J. Annas (1976).

²³ Is this true? Davidson's argument is hardly watertight as it is based on an inductive generalisation from cases. But since my aim here is simply to outline the theory, however, I will simply grant it. As shall become clear in what follows I think that animals do many things actively, but in ways that do not meet the standards set by Davidson's theory of intentional action.

criterion of agency is, in the semantic sense, intentional, the *expression* of agency is itself purely extensional. The relation that holds between a person and an event, when the event is an action performed by the person, holds regardless of how the terms are described.”²⁴ A consequence of this is that it makes sense to ask about an action’s extensional properties. We might inquire, for example, about its spatio-temporal location or its causes and effects — properties that events possess independently of how we may describe them (according to Davidson). And from this it follows that the action and its consequences are always distinct: I may move my hand, thereby pushing a door, thereby causing it to close, and thereby causing a loud noise — but these are distinct events; the first ends, after all, just as the second begins. It is therefore mistaken, Davidson says, “to confuse what my action does cause — the closing of the door — with something utterly different — my action of closing the door.”²⁵

What invites the confusion, according to Davidson, is the fact that an action may be re-described in terms of its effects. Suppose that I move my finger, thereby causing a button to depress, thereby causing a torpedo to launch, and thereby causing the *Tirpitz* to sink. It would of course be perfectly true to say, in such a case, that I sank the *Tirpitz*. But it is not as if, after pressing the button, I need to do *something else*. Rather, I can simply allow the effects of my action to unfold. As Davidson put the point, “We may indeed extend responsibility or liability for an action to responsibility or liability for its consequences, but this we do, not by saddling the agent with a new action, but by pointing out that his original action had those results.”²⁶ Popularising the work of Joel Feinberg, Davidson described this as the ‘accordion effect’.

Now the accordion effect shows how we may expand the description of an action by tracing its effects as they ripple through time; if I do A, thereby causing B, C, and D to happen, we may re-describe my action of doing A as my action of doing D. Equivalently, though, it also allows us to contract the description of an event by tracing it to its source: ‘How did you do D?’, you ask. ‘By doing C-and-B-and-A’, I answer. Importantly, for Davidson, this process of constriction cannot go on forever. The chain of effects must have a beginning in something that we do but not by doing anything else. (The thought is: If everything we did had to be done by doing some other thing, we would never get anything done.) And it is here that Davidson famously suggests that the chain of effects comes to an end in bodily movements: “[...] our primitive actions, the ones we do not by doing something else, mere movements of the body — these are all the actions there are. We never do more than move our bodies: the rest is up to nature.”²⁷

²⁴ (1971, p. 47)

²⁵ (1971, p. 56)

²⁶ (1971, p. 59)

²⁷ (1971, p. 59)

If this is right, then what appears in language as a complex activity with many different descriptions (e.g., ‘sinking the *Tirpitz*, ‘firing a torpedo’, ‘pushing a button’), turns out, upon analysis, to divide itself into a single structure-less atom and what merely happens thereby. The *descriptions* of actions may be complex but actions themselves are simple: they are, always and everywhere, mere movements of the body.²⁸

* * *

Such is the core of the theory. On what grounds might one accept it? Davidson’s original motivation derived from the need to explain the force of the ‘because’ in ‘I did it *because* I had such and such a reason.’ Since it is possible to do something intentionally, and to have a reason for doing it, and yet for that reason not to be one’s reason *in* doing it, we need to know what it *is* for something to be one’s reason in doing something. Famously, Davidson suggests that the only illuminating answer is that one’s reasons play a role in causing the action.²⁹

But even more importantly than this, I think, is the fact that the theory recommends a way of connecting the mind to the rest of nature. For note that by points (1), (2) and (3) an event only counts as an action in virtue of the fact that it can, under some description, be rationalised by reference to the agent’s mental states. If this kind of explanation is causal (as in point (4)), however, it follows, first, that there is an event that causes the action independently of how it is described (causality is an extensional relation) but, second, that the descriptions under which its effect counts as an intentional action depend on the content of the agent’s mental states. Linking extensional and intensional modes of explanation in this way is already a significant step, but we can push it further. For if we also assume, with Davidson, that causality is ultimately a matter of nomological laws that strictly relate event types to one another and that only physics contains such strict laws, it follows that there will be physical descriptions of the intentional action itself.³⁰ As Fredrick Stoutland observes, “This shows that rational explanations not only do not conflict with the laws of physics but are linked with them.”³¹

Much more could be said about the subtleties of Davidson’s theory. But my purpose here is not to defend it. Indeed, I think that it represents a radically false picture of human agency and constitutes a hopeless framework for theorising about animal activity. But we need to understand it, since it provides the default backdrop for approaching questions of animal agency

²⁸ On this point see also Small (2012, p. 157).

²⁹ See Davidson (1963, p. 11ff)

³⁰ See again Davidson (1967) and also Davidson (1970).

³¹ Stoutland (2011, p. 8).

in both the philosophy of animal minds and the related sciences. In the next section, I'll provide evidence that there is indeed such a consensus.

1.3 The Standard Approach to animal agency

Neither the notion of agency nor that of intentional action features prominently in the literature on animal minds. Instead we have 'psychologically rational behaviour', 'goal-directed behaviour', 'truly goal-directed actions', and various other descriptions of a similar ilk. Nevertheless, it seems to me that the point of each of these modes of classification is the same, namely, to mark a binary distinction between expressions of agency in an emphatic sense — intentional actions — and *mere* happenings: tropisms, fixed action plans, reflexes, and the like.

Eric Sidel presents the distinction like this:

An organism whose behavior is caused by, and thus appropriately explained by, its beliefs and desires has distinct representations of its goals and means to achieve those goals. These distinct representations allow it to behave in a goal-directed, rather than merely goal-oriented, fashion [...]

Goal-directed behavior is behavior directed toward achieving a goal, whereas goal-oriented behavior is behavior oriented, say by evolution, so that if it is performed in the right environment, it will lead to a goal. A representation of the goal does not play a role in causing goal-oriented behavior.³²

One of Sidel's points is that it does not follow from the fact that we can identify some purpose in a given animal's behaviour, that the purpose is the animal's own. For example, a cow's blinking reflex is broadly functional: it serves a purpose of keeping dust out of the eyes. But this sort of 'behaviour' is not under the animal's control: it is an innate form of response. Indeed, on this view the same point holds for forms of behaviour that amount to more than mere reflexes. For example, José Luis Bermúdez points out that when redshanks switch their foraging strategy from searching for only large worms to searching for both large and small worms, it would be wrong to think that the birds are *trying* to get the maximum amount of resources available. That is the function of the strategic switch, but that function is constituted behind the animal's back, as it were: the purposefulness of the behaviour, like that of the oaks, is a product of evolutionary design.³³

³² (2009, p. 38)

³³ See Bermúdez (2002, p. 245) I should point out that Bermúdez's position is more complex than this sketch suggests. Indeed, he thinks that there is an intermediate position occupied by animals that make decisions based on weighting different affordances. See, in this regard, his (2003; esp chapters 3 and 7). Much of what I say in chapters 4 and 5 is congenial to this idea.

In a similar vein, Nicky Clayton and her colleagues explain the requirements for ‘psychologically rational behaviour’ as follows:

[...] the issue of whether an animal is psychologically rational turns on the nature of the processes causing its behaviour; specifically on whether this behaviour is caused by psychological mechanisms or by intentional processes. The [animal’s] behaviour is psychologically rational to the extent that it is caused by the interaction of a belief and desire in such a way that performance of the behaviour in question fulfils the desire if the belief is true (and fails to do so if the belief is false). Such an account is intentional because it requires that the antecedent mental states, the belief and the desire, have intentional properties, such as truth and fulfilment, because their content represents current or desired states of affairs.³⁴

As they go on to note, “according to this analysis, rationality is not a characteristic of an animal, but of the processes causing its behaviour.”³⁵ Consequently:

The complex sequences of behaviour by which we drive, walk, or cycle to work may well appear to be purposive and goal-directed but, on further investigation, to consist of a chain of mechanistically elicited habits, albeit complex and highly structured ones, triggered by the stimuli along our route.³⁶

Like Saidel and Bermúdez, Clayton and her colleagues want to emphasise that we cannot simply read off the goal-directedness of a given course of action from the mere appearance of purposiveness. And it is commendable, in my view, that they consistently apply the same principle to the case of human beings: it might turn out that many of the habitual behaviours occurring throughout *our* daily lives are not ‘psychologically rational’ in the relevant sense.

This theme is common. It follows from the assumption that actions are defined as the causal product of hidden processes. As Cecilia Heyes and Anthony Dickinson remark,

[...] an intentional account must be translated into claims about what the animal would have done if its circumstances had been different in certain, specifiable respects from those in which the action actually occurred. [...] The central problem with attributing intentionality on the basis of naturalistic observation alone is that such observation seldom provides the opportunity to evaluate these claims.³⁷

Heyes and Dickinson’s point is generally well-taken. We *should* guard against accounting for behaviour in terms of an ‘intentional process’ when that level of complexity is not called for. (Though it is one thing to agree to this point, and quite another to think that actions are

³⁴ Clayton *et al.* (2006, p. 199)

³⁵ Clayton *et al.* (2006, p. 200)

³⁶ Clayton *et al.* (2006, p. 200)

³⁷ (1990, p. 88)

intentional *in virtue of* their being caused by such processes — a point that I will return to in chapter 3.)

Consider, finally, Peter Carruthers:

What does it take to be a minded organism [...]? [...] you need to possess a certain core cognitive architecture. Having a mind means being a subject of perceptual states, where those states are used to inform a set of belief states which guide behavior, and where the belief states in turn interact with a set of desire states in ways that depend upon their contents, to select from amongst an array of action schemata so as to determine the form of the behavior. This sort of belief/desire architecture is one that humans share. [...]

The crucial components of this account are the beliefs and desires. [...] For] engaging in a suite of innately coded action patterns isn't enough to count as having a mind, even if the detailed performance of those patterns is guided by perceptual information. And nor, surely, is the situation any different if the action patterns are not innate ones, but are, rather, acquired habits, learned through some form of conditioning.³⁸

Although Carruthers presents this as an account of what it is for an animal to be 'minded', we might equally take it as an account of what sort of functional architecture must be in place if the animal is to be capable of acting intentionally. And if we do, then I think we can see his view as another instance of a recognisable pattern: the distinction between action proper and *mere* behaviour lies in the nature of the internal processes that cause it.

* * *

This brief survey is by no means comprehensive. Nevertheless, I think that it paints a sufficiently detailed picture of an approach that is both widespread in the literature on animal minds and bears a recognisable similarity to the Davidsonian account discussed above. If we abstract from the minor differences between those that adopt it, then the approach can be characterised in terms of two assumptions.

(Composition) Intentional actions are distinguished from 'mere events' by something external to them. Hence, we can explain what makes a mere event an intentional action by discovering this additional feature.

³⁸ (2004, p. 207)

- (CTA) The additional feature in virtue of which mere events count as intentional actions include the action's being caused by a pair of mental states, one motivational and the other instrumental.

Although CTA implies Composition, it is worth dwelling on the principle of Composition for a moment.

The idea that lies behind Composition is that intentional actions can be understood broadly as the sum total of two essentially separate elements: (i) the animal's overt behaviour and (ii) some additional fact about it. Thus, what makes a certain behavioural pattern count as intentional, on this view, is nothing about the material process itself, but is some additional feature that it may or may not possess. This is why simply noting that the lion is hunting an antelope cannot on its own support the attribution of purposefulness to the lion. For on this view a hunt can take place either intentionally or not, depending on whether the behaviour is the result of a fixed action pattern, or is rather caused by the animal's beliefs and desires. Like sliding on ice, hunting is something that an animal might do, but it is also something that might lie beyond the scope of the animal's control. (Indeed, as Clayton and her colleagues suggest, the same behavioural type might transition from being purposeful to non-purposeful as a result of habituation within the same individual.)

This assumption is not without precedent, of course. It is enshrined in the famous question, posed by Wittgenstein, which proponents of CTA take as their point of departure: "Let us not forget this. When 'I raise my arm', my arm goes up. And now the problem arises: What is left over if I subtract the fact that my arm goes up from the fact that I raise my arm?"³⁹ The issue that Wittgenstein is drawing our attention to arises as a result of the idea that a single sort of event — my arm's going up, say — can take place as the result of any number of things: you might lift it up, for example, or an electric shock might cause it to go up reflexively, or, indeed, *I* might raise it. Wittgenstein's question assumes that in each of these cases what actually happens — *viz.* my arm going up — is the same; as such, what makes it an arm raising, as opposed to a *mere* arm rising, must be something external to the movement itself.⁴⁰ And similarly with intentional action generally. As Douglas Lavin puts it,

A presupposition of Wittgensteinian arithmetic is that an account of the nature of action begins with a not-intrinsically-intentional movement [...], and through the addition of further distinct factors to the equation, comes to characterize what amounts to intentional action [...]. On this

³⁹ (2009, §621)

⁴⁰ I do not of course mean to suggest that Wittgenstein took this question to be a good one. That is something I very much doubt.

view, action consists of a not-intrinsically-intentional happening, a ‘mere happening,’ occurring in a context where certain further facts obtain.⁴¹

This metaphysical assumption leads naturally (though perhaps not inevitably)⁴² to a methodological one. For it suggests that a theory of intentional action must account for these further facts; it must, in effect, solve for x in the equation ‘behaviour + x = intentional action’. And as we have seen, there is broad agreement amongst the authors quoted above that the requisite account consists of at least three components: (i) a desire-like motivational representation of the activity’s purpose and (ii) a belief-like representation of the means for achieving it that (iii) combine to cause the behaviour (in some right way). These representational features, according to Dickinson, are what “distinguishes goal-directed, instrumental action from habits, reflexes, fixed action patterns, and the like.”⁴³ And this brings us back to CTA.

Now, as I indicated in the previous section, I am proposing to understand the Standard Approach as being committed to at least the central core of Davidson’s account. However, it might be thought that this is a bad fit for at least two reasons.

First, it is not evident from the literature that I have cited that the Standard Approach is committed to the idea that the worldly event in the equation must be a bodily movement. Indeed, when we consider the behaviours that form the topic of empirical research, there is hardly ever a single bodily movement to which the descriptions harken back. For example, when Nicky Clayton and her colleagues describe the complex caching strategies of scrub-jays, they describe a form of behaviour that consists of a series of sub-actions, no one of which can be thought of as the cause of the whole (picking up a cricket does not cause it to be cached, after all). As such, the accordion effect leaves us without an account of how these various descriptions get bound together as part of a single whole.

This suggests that it would be better to think of the behavioural part of the equation as a complex event that is somehow made up of all of these various sub-actions. I shall consider the plausibility of this idea in due course (§3.4). But for now the point to stress is that the Standard Approach is committed to their being *some* event fit to stand in a causal relation to the animal’s mental states if the behaviour is to count as intentional. Rejecting Davidson’s answer to this question does not absolve proponents of this approach from specifying their own. Since this will be enough for the objections that I mean to raise, the question of whether or not the Standard Approach is committed to this aspect of the Davidsonian account is ultimately orthogonal to my aims.

⁴¹ (2015, p. 610)

⁴² I discuss this issue in chapter 3.

⁴³ Dickinson (1994, p. 75)

The second reason why this account might be thought to be a bad fit is that it is not obvious that proponents of the Standard Approach are committed to an animal's behaviour being genuinely *intentional* (as opposed, say, to 'goal-directed', 'psychologically rational' etc.). Perhaps 'intentional' is an additional feature that can be added to this basic structure.

As before, however, I think that this issue is ultimately beside the point. All that matters for me is the idea that it is at least partly in virtue of satisfying the criteria that Davidson specifies that what is done by an animal counts as *more* than a 'mere happening'. Having drawn a distinction between mere happenings and 'goal directed' (etc) actions, we can always add further features in order to make a further distinction between this kind of action and intentional action proper, if we like. But since this would appear to be little more than a terminological issue, I shall persist in describing the dichotomy that the Standard Approach is committed to as that between mere behaviours and intentional actions.⁴⁴

1.4 An Anscombean alternative

It will have been noticed that I have not yet had occasion to discuss the philosopher who appears in my title. But having defined my main target, let me say something briefly about the alternative approach that, I believe, Anscombe offers.

The fundamental difference between Anscombe and causal theorists like Davidson lies in her rejection of the principle of Composition. For her, being intentional is not a 'mere extra feature' that some happenings might or might not have:

If one simply attends to the fact that many actions can be either intentional or unintentional, it can be quite natural to think that events which are characterisable as intentional or unintentional are a certain natural class, 'intentional' being an extra property which a philosopher must try to describe.

In fact, the term 'intentional' has reference to a *form* of description of events.⁴⁵

The form of description of events that Anscombe is referring to is the famous A-D order that is displayed by the man pumping water into a cistern. The man is moving his arms up and down (A) because he is pumping poisoned water (B), and that because he is filling a cistern with poisoned water (C), all of which he is doing for the sake of bringing about the kingdom of heaven

⁴⁴ Let me briefly mention one further reason one might think that Davidson's account is a bad fit for the Standard Approach, namely, the fact that Davidson is one of the most avid critics of the idea that animals can adopt the sort of propositional attitudes that his theory requires. But the central core of the account that I have presented is independent of this further claim. As we shall see, Davidson's scepticism turns on the idea that higher-order capacities, such as possessing the concept of objective truth and, ultimately, speaking a rich language, are required for belief. The account that I have described makes no such assumptions, of course — if they are somehow implied by the very concept of belief, then an argument must be produced.

⁴⁵ (1963, p. 84)

on earth (D). According to Anscombe, this order is held in place, so to speak, by what she calls ‘practical knowledge’ — a kind of knowledge that is *productive* of what is known. But this self-conscious ordering of means and ends is *essential* to the process that is taking place; it is not a feature that is external to it — something which it might *inherit* as a result of being thought about. For, according to Anscombe, *it* could not exist in any other form.

This idea is difficult. It will be developed and defended in chapters 3 and 4. I have mentioned it here just to indicate what I take to be a significant difference between the two ways of approaching questions of animal agency that I shall discuss in this thesis.

However, I also think that taking this point to heart means rejecting some of Anscombe’s own pronouncements about her theory’s applicability to animals. For Anscombe defends the idea that animals can act intentionally on the grounds that their activities are subject to the same ‘form of description’:

[...] we certainly ascribe intention to animals. The reason is precisely that we describe what they do in a manner perfectly characteristic of the use of intention concepts; we describe what *further* they are doing *in* doing something (the latter description being *more* immediate, nearer to the merely physical): the cat is stalking a bird *in* crouching and slinking along with its eye fixed on the bird and its whiskers twitching.⁴⁶

Now in one sense I think that this is exactly right: animal activities are also essentially structured instrumental processes. But if, as I suggested above, the subject’s knowledge of this structure is *essential* in the human case, then I do not think that animal activities can in general be understood in the same way. As such, I shall ultimately argue that that action is a genus that speciates into an intentional and animal form.

Defending these claims will be the task of chapters 4 and 5. Prior to that, however, I try to motivate the need for an alternative account by focusing on problems facing the standard approach in chapters 2 and 3. Going into slightly more detail, the structure of my thesis from this point onward can be summarised as follows:

Chapter 2 — Objective Representation and Propositional Content

In this chapter I attempt to do two things. In the first half, I aim to rebut the sceptical challenge that states with objective content — i.e., content that is about the medium-sized physical objects that make up an animal’s environment — cannot be truthfully ascribed to anything other than a language-user. Focusing on Davidson’s ‘triangulation’ argument, I show both where the argument goes wrong and what we can learn from it. This allows me to make some tentative

⁴⁶ (1963, p. 86)

proposals about where the lower border between bare sensitivity and perception should be drawn. If these proposals are on the right track, then objective *perceptual* content can indeed be ascribed to many animals.

In the second half of the chapter I turn to the Standard Approach's assumption that *propositional* content can be safely ascribed to animals. After outlining what this notion amounts to, I argue that it is difficult to understand how animals could possess the sort of instrumental beliefs that are required by the Standard Approach. I conclude the chapter by proposing an alternative (i.e., non-doxastic) way in which the Standard Approach might understand the instrumental representation that it requires.

Chapter 3 — Taking Action for Granted

The conclusion of chapter 2 is in many ways modest. On the one hand, the positive claim that animals are in perceptual contact with their environments in a way that plants are not can seem pleonastic. On the other hand, the limitations of the Standard Approach that I emphasise are based on a particular way of conceiving of propositional content that is not universally accepted. Moreover, the alternative way of conceiving of an animal's instrumental representations that I suggest presents a proponent of this approach with a further dialectical option. Thus, a Standard theorist might either reject the conception of propositional content that I have adopted, or, alternatively, accept the minimalistic proposal that I suggest.

In chapter 3, however, I suggest that the Standard Approach faces a far deeper problem — a problem that is not solved by either of these dialectical moves. Expanding upon Anscombe's cryptic claim that "We do not add anything attaching to the action at the time it is done by describing it as intentional[.]"⁴⁷ I argue that the fundamental problem with the Standard Approach lies in its adherence to the principle of Composition. Drawing on work by Matthew Boyle and Douglas Lavin, I argue that the Standard Approach's tendency to treat actions as mere happenings caused by beliefs and desires results in a circular explanation of what an intentional action is. In the remainder of the chapter, then, I go on to outline the ways in which 'taking the action for granted' misleads us. In particular, I suggest that the Standard Approach leaves the instrumental complexity of animal activity unaccounted for.

4 — The Structure of Activity

⁴⁷ (1963, p. 28)

In this chapter I begin to outline an Anscombean account of animal agency. This sort of account *starts* with what happens when animals act and then works backwards towards the mind's involvement in the process.

To begin, I suggest that animal activities need to be understood as complex processes with instrumental *parts*. A lion that is in the midst of her hunt, for instance, must move her legs, one after the other, in order to walk across the savannah; but putting one leg in front of the other is not a *cause* of the next movement, nor does the latter follow thereby in some other way. Rather, each step requires a separate contribution on the part of the animal. And while this is little more than a commonplace, I suggest that it is fundamentally important because it allows us to *triangulate* an active role for the animal to play.

The rest of the chapter is devoted to explaining how we should understand this role in mature human beings. Drawing upon Anscombe's concept of practical knowledge, I argue that the agent's practical knowledge plays a constitutive role in the structure of the activity itself. What emerges is a conception of intentional action as an essentially self-conscious process. And this in turn presents the obvious problem of seeing how an Anscombean theory of agency can be applied to animals.

5 — *Animal Agency*

In this final chapter I take up this challenge. After admitting that Anscombe's account cannot be simply carried over to animal activities, I proceed to show how an recognisably Anscombean account of animal activity can nevertheless be developed. My strategy for doing so is based on two ideas. The first is that practical knowledge, like knowledge more generally, is internally related to the capacity to use a fact as a reason: an agent that is intentionally doing A can use the fact that she is doing A as a reason for (*inter alia*) doing B (where B is a means to doing A). The second idea is that the ability to use a fact as a reason can be seen to lie on a spectrum: at one end there is the sophisticated, two-way ability to take a fact 'into consideration or account'; in the middle there is the capacity to be 'guided' by a fact; and at the other end is the involuntary susceptibility to a fact's imperious 'control'.

In an effort to spell these metaphors out, I introduce some technical notions. These serve to explain how an animal might express their agency by, for example, coordinating perceptions and actions, and also explain how animals are capable of creating enough distance between their perceptions and emotions to select between various alternatives. I go on to suggest that animals whose activities are explicable in this way are indeed guided by facts in a way that is similar to, but distinct from, the way in which an intentional agent takes facts 'into consideration or account'.

I conclude by highlighting some of the limitations of the Anscombean approach that I have defended, before making some tentative suggestions about the ways in which it might be further developed.

Objective Representation & Propositional Content

“...inasmuch as an animal is capable of appetite it is capable of self-movement; it is not capable of appetite without possessing imagination; and all imagination is either calculative or sensitive. In the latter all animals partake.”⁴⁸

2.0 Preliminaries

As I emphasised in the introduction, the power of perception is built into our naïve conception of animal life: animals see, hear, and smell (etc.) the medium-sized physical objects of their environment, including, for example, one another. Unlike plants which merely respond to sensory stimulation, animals have a perspective on things.

One way to put this naïve thought is to say that an animal’s perceptual states have objective content — content that is *about* the medium-sized physical objects with which they interact in the course of their activities. Of course there are difficult questions to be answered about precisely how to specify this content in such a way as to capture the way in which the world manifests itself to the animal. But at a minimum, it seems almost trivial to say that, however the cat perceives the mouse, what her perception is *of*, is, precisely, the mouse.⁴⁹

Though these commonplaces are deeply entrenched in our common-sense conception of animals, this entrenchment can also, I think, give rise to a certain anxiety. For we do not

⁴⁸ Aristotle (1984a; III.10, 433b27-29)

⁴⁹ This is not an entirely happy way of putting matters since the idea that perceptual states have content is controversial (see, for example, Travis (2004)). But all that I mean to insist on is that there is typically an answer to the question of what her perception is *of* that will pick out something beyond the animal’s sensory surfaces.

typically stop to ask ourselves whether or not we are justified in ascribing objective content to animals. The idea that animals are in touch with their environments forms the *start* of many inquiries, but rarely its end. And this could well raise the spectre that our tendency to think that animals are in touch with their environments is little more than an anthropomorphic prejudice. Mightn't animals, no less than plants, be capable of no more than differential responses to proximal stimulation?

To answer this question affirmatively is to deny that our practice of ascribing objective content to an animal's perceptual, cognitive, and conative states can be taken at face value. It is to deny, for example, that a general statement like 'Sheepdogs see the sheep that they herd' is true — or at any rate, that it is true in the same way that 'Shepherds see the sheep that they herd' is. At best, this way of speaking is a theoretically useful but ultimately groundless basis for explaining behaviour that could with equal right be taken to be a mere response to the proximal stimulation on the organism's sensory surfaces.

Although radical, this position has indeed been defended by a number of philosophers since Descartes. Most arguments in this vicinity rely on the theme that language is in some way necessary for objective representation⁵⁰ and, in the first half of this chapter, I propose to examine the version of this argument that appears in Donald Davidson's work. Although I will ultimately suggest that the argument is unsuccessful, I think that we can learn a lot from seeing precisely where it goes wrong. Indeed, I will use one of the central themes of Davidson's argument to make some tentative proposals about where the lower border between bare sensitivity and perception should be drawn. If these proposals are on the right track, then objective perceptual content can indeed be ascribed to many animals.

In the second half of the chapter I turn to the question of whether the Standard Approach's assumption that the practice of ascribing beliefs to animals is also warranted. After spending some time spelling out exactly what the significance of this question is, I suggest that the instrumental beliefs posited by the Standard Approach are probably not ascribable to any species of animal. If this is right, then the Standard Approach's domain of application is severely limited. However, I conclude on a conciliatory note by proposing an alternative (i.e., non-doxastic) way in which instrumental representations that the Standard Approach requires might be understood. If this proposal can be made to work, then the Standard Approach may be more broadly applicable than my initial claims suggest. (Though this proposal will not, I think, save the approach from the arguments of the next chapter). But this is to get ahead of ourselves. First we need to consider Davidson's argument.

⁵⁰ See, for example, Descartes (1991, pp. 302-304), (1985, p. 140); Rorty (1981, p. 187); Brandom (1994; esp chapter 4)

2.1 The problem of stimulus determination

The considerations that lead Davidson to the conclusion that objective content cannot be ascribed to animals derive from his attempt to answer a question that he inherits from Quine. The question is: “where, in the infinite causal chains that lead to the sense organs, should we locate the elements that give content to our observation sentences and accompanying perceptual beliefs?”⁵¹ This needs answering, he suggests, as soon as we accept the modest externalist premise that our perceptions (and the beliefs based upon them) have the contents that they do partly in virtue of the causal relations that obtain between the subject, her sense organs, and the environment. For doing so forces us to explain why certain links in the causal chains connecting us to our environment are more important than the others — why, in other words, an accurate specification of what is perceived is couched in words that refer to medium-sized physical objects (what I will sometimes call ‘ecological objects’)⁵² rather than any of the intermediate links along the causal chain. Call this the problem of stimulus determination.⁵³

In our own case, of course, the problem is one of explaining something that Davidson does not doubt: our ability to perceive and think about an objective world. But he evidently thinks that we cannot be so confident once we turn to animals. In discussion of a dog that has been conditioned to salivate in response to the ringing of a bell, for example, he asks:

[...W]hy say the stimulus is the ringing of the bell? Why couldn’t it be the vibration of the air close to the ears of the dog — or even the stimulation of its nerve endings. Certainly if the air were made to vibrate in the same way the bell makes it vibrate it would make no difference to the behavior of the dog.⁵⁴

When presented in this way the difficulty appears to be one of justifying our confidence that animals are reacting to ecological objects rather than surface irritations. ‘What fact about their behaviour’, Davidson seems to be asking, ‘justifies our claim that they are responding to ecological objects and events rather than anything in the intermediate causal sequence?’ But the challenge is not intended to be purely epistemic. For Davidson goes on to suggest that, at least in the case of the dog, the cause is not merely underdetermined by the behavioural evidence, but is, in fact, indeterminate: “The problem is not, I should stress, one of verifying what objects or events a creature is responding to; the problem is that without a second creature

⁵¹ Davidson (2003, p. 85)

⁵² Following Susan Carey (2011, p. 155) I take this to mean, approximately, “bounded, coherent, separately moving, spatiotemporally continuous material entity.”

⁵³ I take the name from D. Finkelstein (2007)

⁵⁴ (1992, p. 262)

responding to the first [creature], there can be no answer to the question.”⁵⁵ Davidson’s reference to the second creature is an indication of where he thinks the solution to the problem lies. But, setting that to the side for the moment, let us ask why Davidson thinks that the indeterminacy vitiates our practice of ascribing content to the dog’s perceptual state. Why not conclude, instead, that the content is underdetermined?

The reason is that without a right answer to the question of what the creature is responding to, the notion of error, and therewith the idea of a fallible perspective on an objective world, loses its traction. For we can no longer hope to distinguish between a creature’s making a mistake in responding to a distal object and its responding to something closer in. So, for example, a frog that snaps at passing bee-bees might *appear* to be mistaking them for flies. But if it is indeterminate whether the frog is responding to proximal or distal causes, then we could with equal right claim that it is responding appropriately to a stimulus class determined by the size and colour of the retinal impression (say). There is simply no fact of the matter. And notice that the point does not turn, according to Davidson, on the limited repertoire of behaviours available to simple creatures; it applies to *any* solitary creature, including human beings: “If we consider a single creature by itself, its responses, no matter how complex, cannot show that it is reacting to, or thinking about events a certain distance away rather than, say, on its skin.”⁵⁶

I think that this claim is mistaken and, indeed, incredible. I will argue the point below. But let us first consider how, exactly, the second creature is supposed to help.

Davidson’s idea is that the second creature creates the space necessary for error by affording us some justification for locating the source of the reaction at the point in the causal chain that is *shared* between the creatures. To get the flavour of the idea, consider a simple case in which a pair of animals are drinking from a watering hole when a predator appears on the opposite bank. Suppose that each animal reacts similarly: both flee. Since the predator is the only explanatorily relevant common entity (the creatures do not share photons or neural firings), we have a *prima facie* reason for saying that the animals are reacting to an ecological object. Davidson writes,

With a single creature, it is hard to decide what it is reacting to when a stimulus hits. When the frog sticks out its tongue, is the stimulus a fly or the firing of a certain pattern of receptors in the eye? [...] The slightly complex social situation I have been postulating eases this problem. The stimulus that matters is the nearest mutual cause of the joint reaction. When the triangle is

⁵⁵ (1992, p. 263)

⁵⁶ (1992, p. 263)

working normally, the mutual reactions of the two (or, of course, more) creatures triangulate the relevant stimulus, locating it in a public space.⁵⁷

Since ecological objects lie at the intersection of triangles between social creatures, there seems to be some point in our locating the cause of their joint reactions at some distance from the skin. What makes space for the notion of error, then, is that the reactions may not correlate: one animal might flee upon spotting the approaching predator while the other ignores it and gets eaten. In such a case, Davidson remarks, “We, looking on, will judge that [one of the] creature[s] has erred. The creatures themselves are also in a position to come to the same conclusion. If they do, they have grasped the concept of objective truth.”⁵⁸

These passages seem to suggest that since our judgments about what social creatures are reacting to can be justified, we at least know that *some* animals can be credited with contentful states. The error would appear to be there for the creatures to recognise, even if they are not sophisticated enough to do so. But in other places Davidson resists this. In discussion of the parrot Alex who, when prompted to ‘say’ what certain objects have in common is reliably able to do so (for a limited range of properties), he asks:

What is it that tells us that the stimulus (cause) of Alex’s ‘answer’ to the question ‘What’s the same?’ isn’t the activation of certain rods and cones in his eyes, or the firing of certain optic nerves, or the photons bouncing off surfaces we see as the same color? All of these causes, and endless more, are common to the cases where Alex emitted the sound ‘Color.’ We have no grounds for choosing one of these causes over the others.⁵⁹

Alex, his trainer, and the ecological objects that share a colour constitute a primitive triangle (or rather, a primitive polygon). Accordingly, if shared reactions to a common cause were sufficient for the ascription of objective content, we would expect Davidson to allow that Alex is responding to the object’s colour. But, as with a solitary creature, he seems to think that it hardly matters which among the various links in the causal chain we choose to identify as the cause.⁶⁰ So, although triangulation might make it easier for *us* to decide what the creatures are responding to, our decision should not be mistaken for a discovery of how things are independently of us: we are still without objective grounds for supposing that each animal is perceiving their shared environment. And if this is right, then the difference between social and solitary animals does not run as deep as Davidson’s initial remarks would suggest; the problem of stimulus determination applies to both if it applies to either.

⁵⁷ (1999a, p. 41)

⁵⁸ (1997, pp. 26-27)

⁵⁹ (1999, p. 34)

⁶⁰ See Davidson (1995, p. 234)

In order to solve the problem, Davidson thinks that the relevant creatures must come to an appreciation of their own situation. They must, in other words, come to know that they are taking part in triangulation:

Since the [cause] is identified only by the intersection of two (or more) sets of similarity responses (lines of thought, we might almost say), to have the concept of [the cause] is to recognize the existence of a triangle, one apex of which is oneself, another a creature similar to oneself, and the third an object [...] located in a space thus made common.⁶¹

And it is here, Davidson suggests, that linguistic communication comes in: “For two people to know of each other that they are so related, that their thoughts are so related, requires that they be in communication. Each of them must speak to the other and be understood by the other.”⁶² Thus do we arrive at Davidson’s conclusion that objective representation presupposes linguistic communication.

2.2 Ways to triangulate

I turn now to an assessment of the argument.

On my interpretation, Davidson’s reasoning looks something like this:

P1: Nothing about a creature considered in isolation could provide grounds for judging that it is reacting to distal, as opposed to proximal, stimuli.

P2: In such a case, our content ascriptions cannot be taken as right or wrong; our practice lacks objective validity.⁶³

P3: The capacity to recognise that a second creature is responding to the same objects as oneself is necessary for their being such an answer.

P4: Linguistic communication amongst social creatures is a necessary condition for the capacity to recognise that a second creature is responding to the same objects as oneself

C: Linguistic communication amongst social creatures is necessary for objective representation.

⁶¹ (1992, p. 264)

⁶² (1992, p. 264)

⁶³ The allusion is of course to Kant. Kant’s problematic was that concepts that we cannot help but employ in our thought about the world might not actually apply in it. Our thinking would therefore lack ‘objective validity’. Since Davidson would certainly agree that we cannot but apply intentional states to animals even though such ascriptions are untrue, I take the comparison to be apt.

I suspect that P4 is false, but I propose to grant it for the sake of argument.⁶⁴ I will also grant P2. I mean to dispute P1 and P3.⁶⁵

Consider first P3. It is difficult to see how the higher-order capacity to recognise that one is triangulating with another creature can be a necessary condition for objectification. The reason is that, if it were, it would seem to render the emergence of thought, and indeed perception, unintelligible. For how could the first creature grasp what a second creature has in mind unless she is already able to perceive it?⁶⁶ If one's thoughts and perceptions only acquire objective content by interpreting other creatures, it becomes completely mysterious how one could notice that there was another creature there to be interpreted.

Perhaps Davidson would accept this. "[...T]here is" he says, "a perhaps insuperable problem in giving a full description of the emergence of thought."⁶⁷ And I suppose that that is indeed one possible conclusion that can be drawn from the situation. But a more plausible one is surely that Davidson's claim that a higher order capacity to recognise what other individuals are reacting constitutes a necessary condition for objective representation is wrong.

In response, Davidson would argue that the reason that such higher-order capacities are necessary is that in order to have any belief about the world one must understand the distinction between seeming and being: one must possess the *concept* of objective truth, and it is only by coming to see that a second creature has false beliefs that one could acquire it.⁶⁸ But again: why should we accept that *this* higher-order capacity is required? As far as I can tell, Davidson has just one independent argument that is meant to sustain this idea. It goes something like this:

Suppose I believe there to be a coin in my suit pocket and decide to reach for it. Since I expect to find it, I will be *surprised* if it is not there. From this Davidson concludes, not

⁶⁴ There is an enormous empirical literature on the 'mindreading' capacities of animals, some of which suggests that some certain species are capable of understanding what other creatures perceive, believe, know, and intend. The literature is controversial, however, and I will not rest my objections upon it. For a useful overview, see Andrews (2015; chapter 6).

⁶⁵ Let me briefly mention a line of response that I will not pursue. Some have thought that the way to deal with the problem of stimulus determination is to point out that we have a network of holistically supported beliefs about animal life that are incompatible with the idea that animals are not responding to distal objects. After all, animals flee, play, and forage. Some of them hunt. But as David Finkelstein (2007) and Jason Bridges (2006) both point out, it simply makes no sense to say that a lion (or anything else) is hunting patterns of stimulation. Distributions of light arrays cannot, after all, be caught, killed, or eaten. Nor would it make sense to suppose that when animals flee, forage, or play, their attention is directed towards anything other than a medium sized physical object. Animals flee from predators, forage for truffles, and play with each other. Contrary to Davidson's claim that we 'have no grounds' for preferring distal over proximal causes, Bridges and Finkelstein suggest that we do indeed have such grounds.

I think that this sort of response is in danger of begging the question against Davidson. For Davidson would of course accept that these ways of describing things come naturally to us; his point is that our preference for classifying the contents of an animal's states is no more than that. As Davidson (1991, p. 142) himself remarks, "it begs the question to project our classifications on to nature."

⁶⁶ See also Bridges (2006, p. 301)

⁶⁷ (1999b, p. 11)

⁶⁸ Cf. Frege (1956, p. 302ff)

implausibly, that the capacity for surprise is a necessary condition for the capacity to form beliefs. “But surprise involves a further step”, Davidson insists. “[...] Surprise requires that I be aware of a contrast between what I did believe and what I come to believe. Such awareness, however, is a belief about a belief: if I am surprised, then among other things I come to believe my original belief was false.”⁶⁹ Since having a belief about a belief requires the *concept* of belief and this, in turn, presupposes that the organism understand that a belief is the sort of thing that may be false, it would appear to follow that in order to have any belief at all, an organism must have a conception of objective truth. Thus concludes Davidson’s brief argument.

There are numerous premises that one might take issue with here. In particular, why should we should accept the crucial premise that in order to be surprised I must ‘become aware’ of the contrast between my current and prior belief? Why isn’t it enough to simply *give up* the original belief? Why must one form a reflective belief *about* it? Davidson’s suggestion seems to be that in cases where the higher-order belief does *not* follow, one is merely startled, not surprised. But then, consider the difference between (a) a dog that jumps in response to a loud noise and (b) a dog that looks around quizzically after you pretend to throw its ball (which is still in your hand, behind your back). The difference between these cases, on one natural interpretation, is that in (b) the dog is *surprised* that the ball did not get thrown since he had a positive expectation that something else was going to happen, whereas, in (a), he is merely startled by an unusual experience (he was not expecting anything in particular). At no point in telling this story do we need to invoke a second-order thought to distinguish between these cases, and to simply *assert* that in (b) the dog is not surprised because he lacks a second-order thought is to assume what needs to be proved.⁷⁰

If this is right, then Davidson’s argument fails to establish that objective content can only be ascribed to creatures capable of linguistic communication. Still, refuting the idea that linguistic communication between social creatures is necessary for objective representation does not rule out the weaker claim that *some* kind of (non-linguistic) social context is necessary. Thus it might be replied on Davidson’s behalf that a minimal kind of triangulation between social creatures is needed to give sense to the idea that animals are responding to objects in public space.

But I think that even this weaker claim is false. For although I agree with Davidson triangulation is an essential component of the distinction between bare forms of responsiveness and objective representation, I believe that he is wrong to suppose that triangulation can *only* occur through the joint reactions of *multiple* individuals.

⁶⁹ (1982, p. 326)

⁷⁰ Similar assessments of this argument can be found in Burge (2010a), Glock (2000), and Jamieson (2009).

To see why the second individual is unnecessary, suppose that we observe an animal flee from a predator on two occasions: in the first it is facing it straight on and in the second it is facing it from the side. Naturally, the proximal stimulations are different in these two cases. The response, however, is similar. Since the distal object is the common factor between the two responses, our explanations are streamlined by supposing that the ecological object is the cause of the reaction.⁷¹

This case is simple, of course, but the phenomenon is quite general: the same object might give rise to a wide range of surface irritations depending on the illumination, distance, and perspective of the perceiver, while the behavioural response remains constant. This phenomenon (known as perceptual constancy) is widespread in the animal kingdom; it occurs in all mammals, and has been documented in reptiles, amphibians, birds, and even some insects.⁷² For example, frogs tend to snap at dark moving objects of roughly the same size as a bug within a range of about seven inches: the size of the retinal image varies widely within this range, but the response remains the same. And, importantly, the converse is also true: frogs do *not* snap when larger dark objects are situated at a distance that produces retinal images whose size falls within this range.⁷³ The consensus in perceptual psychology is that this provides compelling evidence that frogs are perceiving the object's real size. On what grounds can Davidson dispute this conclusion?

I think that there are two things that Davidson might say in response. The first is that all that this case shows is that it is easier for *us* to explain the animal's reactions by ascribing perception of a distal object to the frog. We still lack sufficient grounds for supposing that the animal is *really* perceiving distal objects. Thus Davidson writes that "even if the frog [could] save its fire when the target is a large and distant bird or an airplane,^[74] we would be in no position to choose one source of the cause of the reaction over another; the frog would just be habituated to a different, though more complex, pattern of proximal stimuli."⁷⁵ But why should we accept this? The objection is simply *ad hoc* unless Davidson can provide some independent support for it. Yet his typical means of doing so depends on the idea that objective representation requires the higher-order cognitive abilities outlined above. We have rejected this idea, however, and it is unclear how else Davidson might hope to sustain it.

At this point the only remaining option for Davidson, it seems to me, is to go on the offensive and argue that if we do not take higher-order cognitive capacities to be required for objective

⁷¹ See also Bridges (2006, p. 296ff).

⁷² See, e.g., the references collected in Burge (2010a, p. 419fn).

⁷³ See especially D. Ingle (1983, 1998).

⁷⁴ And just to reiterate, there is substantial evidence that frogs *can* do this.

⁷⁵ (2001, p. 141)

representation, this commits us to the absurdity of ascribing objective content to any creature with minimal discriminatory abilities. Davidson remarks:

I assume we don't want to view earthworms and sunflowers as having [propositional attitudes]. This would be a terminological mistake, for it would be to lose track of the fundamental distinction between a mindless disposition to respond differentially to the members of a class of stimuli, and a disposition to respond to those items *as* members of that class.⁷⁶

This passage reveals a gap in Davidson's thinking, however. For without further argument, it is unclear why ascribing perception of ecological objects to an animal should *ipso facto* commit us to holding that it has any beliefs about them.⁷⁷ Perceptual content is not propositional content (or so I will later argue). But nor is it ascribable merely on the basis of the discriminatory abilities evidenced by the simple organisms that Davidson mentions. Perceptual abilities sit somewhere in-between.

I shall later have something to say about the border between perception and propositional thought. But let me briefly comment upon the distinction between bare sensitivity and perception.

Consider (to take an example that I know more about) the Venus flytrap. Such plants are notoriously poor at discriminating between sources of stimulation. They will snap shut around insects, of course, but also around fingers, pencils, and whatever else might come in contact with their mechanosensors. Indeed, such plants are profoundly *insensitive* to the source of the stimulation: Their only discriminatory capacity lies in their ability to respond differentially to the number and frequency of the stimuli themselves.⁷⁸

Of course, this minimal discriminative ability is not insignificant. It means that, in natural environments at least, the plant's responses correlate highly with the presence of live insects: inanimate objects which happen to brush against the plant's sensors tend to do so only once (and so do not elicit the closing response) and there are fewer humans around to skew the averages. In a purely information-theoretic sense, then, the registration of sufficient proximal stimulation carries information about the environment.⁷⁹ But to re-emphasise: *all* that matters

⁷⁶ (2001, pp. 137-138) In the first sentence I have replaced 'concepts' with 'propositional attitudes'. This is justified, I think, since Davidson prefaces this passage with the claim that there is no distinction between having one and having the other.

⁷⁷ Pace Glock (2013). See also Burge (2010a, Part II)

⁷⁸ See, e.g., Böhm *et al.* (2016)

⁷⁹ It is not uncommon to hear such facts being used to justify claims that the plant is discriminating ('representing', 'identifying', 'recognising', etc.) features of the environment. A recent article in *Current Biology*, for example, would have it that "[T]he number of [sensors] a victim triggers while trying to break out of the trap *identifies* the moving prey as a struggling [sodium] rich animal and nutrition for the plant." (2016, p. 286; emphasis mine) For the reasons outlined below, however, I agree with Davidson that these claims are at best misleading.

when it comes to determining the plant's response is the number and frequency of the proximal stimulation itself. We can freely vary the shape, size, texture, or any other property of the object we brush against the sensors with, but, so long as we do so twice within a twenty-second window, the plant will respond in the same way. As Tyler Burge remarks about a similar case involving magnetotactic bacteria, "Nothing in the individual's capacities [...] distinguishes (a) environmental causes that figure functionally in the individual's basic needs and activities from (b) sensory registration (or functional encoding) of proximal causes — from the surface effects of the environmental causes."⁸⁰ Since there is no evidence that the flytrap is capable of responding differentially to the properties of the object *per se*, mentioning them does nothing to advance the explanation — all such references can easily be dropped without loss to the theoretical economy. It is in cases like this, I think, that it is wholly appropriate to describe the plant's reaction as a functional — but mindless — response to proximal stimulation.

At what point do we arrive at object-directed responses — at perceptual sensitivity? I do not pretend to have a developed answer to this question, but I think that the evidence cited above about frogs is compelling. More generally, when responses remain suited to the properties of the object despite wide variation in proximal stimulation and, conversely, alter when disparate properties produce similar proximal stimulations, we have evidence that the animal is perceptually sensitive to the objective features of the situation. As mentioned above, such evidence has been well documented in mammals, birds, reptiles and many insects. Indeed, this evidence becomes even more compelling when it is coupled with signs of anticipatory expectations of object permanence and evidence that the animal in question can keep track of particulars over time and space⁸¹ — evidence which has similarly been documented in a wide range of species.⁸²

Again, I do not want to make strong claims about the necessity or sufficiency of any of these features. They are offered simply as indications of what sort of empirical evidence might be brought to bear on the issue. In any case, I think that I have said enough to show that Davidson has not given us any *a priori* reason for thinking that the evidence cannot be taken at face value.

⁸⁰ (2010a, p. 317)

⁸¹ Cf. Burge (2011, p. 125) "Object [...] perception is constitutively dependent on coordination between the perceptual system and perceptual anticipations of persistence over time; commonly in motion and commonly behind barriers. Therefore, to be a perception as of [an object], a representation must be associated with a tendency to perceptually anticipate certain types of continuity. A perceptual representation can present something as looking like [an object] in current experience, as long as the 'look' is associated with perceptual anticipations of certain types of continuity."

⁸² See, for instance, Barner *et al.* (2008) and Burge (2010a, p. 450ff). For evidence of these capacities in pre-linguistic infants, see especially S. Carey (2009).

2.3 Propositional content

This is of course good news for the Standard Approach since it is predicated on the assumption that animals are capable of representing their respective environments. But the support that it affords is minimal. For the Standard Approach assumes not merely that the animals are in *perceptual* contact with objects, events, and properties, but that they have instrumental *beliefs* about them. Yet states of belief, on most accounts, take propositions as intentional objects — and can anything other than a language user believe a proposition?⁸³

In order to answer this question, we obviously need to know what is involved in believing a proposition. I shall have some things to say about that in a moment. But let me begin by describing one of the cases that is often used to suggest that the ascription of instrumental beliefs to animals is warranted.

In a classic study, Robert Rescorla and J.C. Skucy showed that rats which learn to press a lever for food will cease to do so once the delivery of food is made independent of the pressing. That is, rats appear to be sensitive to the counterfactual connection between a specific action-outcome pair — a sensitivity which is evidenced by their ability to adjust their behaviour once the connection ceases to hold.⁸⁴ According to Heyes and Dickinson, this shows that the rats have a belief with causal content: they believe that pressing the lever causes food to be delivered, but will adjust this belief once the causal dependence is broken.⁸⁵

Such findings are certainly interesting. They *do* make it plausible to suppose that the rats are representing a connection between doing one thing and doing another. But they also present something of a puzzle. For the rats are *not* able to learn to adjust their behaviour in light of a similar contingency holding between approaching their food dish and the delivery of food: they will persist in their approach regardless of whether doing so results in food being delivered (and will lose significant rewards as a result).⁸⁶

Heyes and Dickinson suggest that although this undermines the idea that the rats' approach behaviour is caused by an instrumental belief, it is not especially surprising. For it just reinforces the point that "intentionality is primarily a property of an agent with respect to a particular action rather than of the agent *per se*."⁸⁷ But this result *is* surprising. For consider how odd this result would be if a mature human being were in question: their behaviour is sensitive to the

⁸³ For a recent defence of the idea that the answer is 'No', see Marcus (2012; chapter 3).

⁸⁴ For the original study, see Rescorla and Skucy (1969).

⁸⁵ I will follow Heyes and Dickinson in thinking of the relation recorded in the instrumental belief as causal, but I do not mean anything particularly weighty by that. For example, we can think of the relation as enabling, or as triggering, etc. The minimal claim is that the belief is supposed to track the counterfactual dependence between doing one thing and some outcome's coming about.

⁸⁶ Holland (1979)

⁸⁷ (1990, p. 91)

counterfactual dependence that obtains between doing A (in context C) and outcome O, but it is seemingly insensitive to the counterfactual dependence that obtains between doing D (in context E) and the same outcome. As a one-off occurrence this is of course unremarkable: they simply haven't learned that D and O are causally related. But the situation surely does become incredible when it is added that the subject in question cannot *learn* to appreciate the contingency — especially so given the relative simplicity of the situation in question.

The oddity of this situation will perhaps become clearer if we reflect on the idea that the objects of belief — i.e., propositions or thoughts⁸⁸ — are compositionally structured. As I shall understand it, this is the broad idea that propositions are 'composed' of elements which the subject puts together in acts of judgment. Now there are obviously important questions concerning what these elements are and what this putting together amounts to (I shall address them in a moment). But at this stage we can leave it at an intuitive level: the idea is just that, if someone is able to form the thought that raspberries are sweet (say), then they are able to think *of* raspberries, that they are sweet.⁸⁹

This compositional conception of thoughts is not universally accepted, of course.⁹⁰ But I think that we can see the motivation for it as lying at the intersection of two ideas. The first relates to the systematic and productive features of thought, and the second relates to its inferential power. I'll discuss the first point in this section before turning to the second in the subsequent section.

Start with the naïve idea that the thought that *a* is *F* has something in common with the thoughts that *b* is *F*, that *c* is *F*... etc., on the one hand, and with the thoughts that *a* is *G*, that *a* is *H*...etc, on the other.⁹¹ As Frege teaches: the first series of thoughts exemplifies the logical form $F(\xi)$, the second exemplifies the form $\phi(a)$, and both exemplify the form $\phi(\xi)$.

Now someone opposed to a compositional conception of thoughts will protest that this structural similarity is just a reflection of the fact that the sentences used to express these thoughts have words in common.⁹² But there is more to the idea than that.

For note that just as anyone that understands the sentences 'Fa' and 'Gb' will also understand the sentences 'Fb' and 'Ga', anyone that is able to *think* that *a* is *F* and that *b* is *G* will also be able to entertain the thought *a* is *G* and the thought that *b* is *F*.⁹³ Why should this be? On a

⁸⁸ I shall use 'thought' as equivalent to 'proposition' throughout.

⁸⁹ Cf. Heck (2007, p. 127ff)

⁹⁰ It is not, for example, by Glock. See his (2012, p. 580ff) for the objection that the idea that thoughts are structured in any significant sense is unmotivated at best and incoherent at worst. I hope that the remarks that follow will alleviate this worry.

⁹¹ See Evans (1982, p. 100ff).

⁹² For an objection along these lines, see Glock (2012, p. 581).

⁹³ Provided, that is, that the predications are categorically appropriate. Thus it is no part of the view that, just because someone can understand the sentences 'eleven is prime' and 'John is loud' they must also be able to

strong version of the language of thought hypothesis the explanation is that the neurological mechanisms that are supposed to implement these thoughts have a sentence-like syntactic structure; talk of ‘composition’ can be taken quite literally, on this view. But we do not need to commit ourselves to this (nor do we need to deny it). Instead, we can adopt Gareth Evans’s suitably broad suggestion that thinking that *a* is *F* involves the exercise of a pair of general-purpose abilities: the ability to think, first, of a certain object, *a*, and, second, the ability to think that something is *F*.

Evans’s idea is simple in outline: the reason that anyone who can think that *a* is *F* (and that *b* is *G*) can also think that *a* is *G*, is that these thoughts are underwritten by the same ability: the ability to think of the object *a*. And similarly, someone who can think that *a* is *F* (and that *b* is *G*) will *ipso facto* be able to think that *b* is *F* because entertaining these thoughts involves the exercise of the ability to think of some arbitrary thing that it is *F*.

Evans’s proposal is attractive because it gives us a way to understand talk of propositional structure as an abstraction from psychological reality: propositions decompose into isolable ‘parts’ just in the sense that acts of *judgement* involve the exercise of distinct — but always conjoined — abilities. These abilities enable their possessor to think of a particular object in indefinitely many ways, on the one hand, and to think of a property (or relation) in such a way that it is independent of the objects believed to instantiate it.

It is in light of these considerations that Evans proposes his famous ‘Generality Constraint’:

If a subject can be credited with the thought that *a* is *F*, then he must have the conceptual resources for entertaining the thought that *a* is *G* for every property of being *G* of which he has a conception.⁹⁴

I take it that having a conception of a property *G* is another way of possessing the concept of that property, and this of course raises the question of what it is to possess a concept. But, setting that question to the side, I think that we have come far enough to appreciate what is odd about the juxtaposition of (i) the idea that the rats believe that depressing a lever causes the release of food with (ii) the fact that the rats cannot learn to ‘detach’ the concept of causality from these specific events and apply it to others. For this shows that the rats’ representation (to adopt the broadest term available) of the causal relationship between lever pressing and food-delivery cannot be understood as the exercise of the separable abilities to think, of lever pressing and food-delivery, on the one hand, and the ability to think of two arbitrary events that they are causally related, on the other. For if the rats did possess these abilities, then we would expect

understand the sentences ‘eleven is loud’ or ‘John is prime’. Similarly in the case of thoughts. (I will take this qualification as read in what follows.)

⁹⁴ (1982, p.104)

them to be able to deploy them in new contexts, to new types of event. What the experimental results show, however, is that their representation of the causal relation between two events is bound up with specific event-types: the rats represent, as it were, that-depressing-a-lever-causes-food-pellets-to-be-released, but this mode of apprehension is not compositional; it does not involve the exercise of the general-purpose abilities that Evans describes.

This interpretation might be challenged. After all, it has not been shown that the rats ‘have a conception’ of their approach behaviour and, hence, that the Generality Constraint has been violated. But think again of what we would expect of a paradigmatic believer (i.e., a person) in a similar experimental set-up. He already knows that pressing a lever produces some desired outcome, and he is then exposed (repeatedly) to another action-outcome pair. We would surely expect him to be able to *use* the conceptual ability associated with the concept of causality — that is, the ability to think of any two arbitrary events that they are causally related⁹⁵ — in this new context, and, hence, to learn about the contingency which holds between the novel action-outcome dependence. If we want to treat the rats as believers *in the same sense* then they should be subject to the same kind of expectations.

Examples of contextual limitations of this nature are widespread in the literature. In another notorious experiment, Sarah Boysen’s chimp Shiba was presented with two bowls of candy, one of which had noticeably more candy than the other. Shiba would be given the opposite bowl from the one that she indicated, the one she indicated going to another chimp. Shiba tended to indicate the bowl containing more candy (and hence to receive the lesser reward) well-above chance. However, when the candies were replaced with numerals, Shiba mastered the task quickly and without further training.⁹⁶

This study is interesting for a number of reasons, but for our purposes the relevance is just that Shiba’s grasp of the dependence between what she does and what further that affords seems to again be limited to specific event-types. That is, she seems capable of representing the causal relation that obtains between indicating the bowl with the lesser numeral and receiving a large reward of candy, but *not* of representing the causal relation that holds between indicating the bowl with the lesser number of candies and receiving a large reward of candies.

As before, this interpretation might be challenged. Boysen and her colleagues suggest that it is not so much Shiba’s failure to understand the relation between reaching for the larger amount

⁹⁵ For Evans, possessing the concept and having this ability are not two different things: “Someone who thinks that John is happy must have the idea of a *happy man* — a situation instantiated in the case of John (he thinks), but in no way tied to John for its instantiation.” (1982, p. 103) Quite generally, I think, Evans’s point is that one cannot think of something as *F* unless one knows what it is to be *F*, but, equally, one cannot know what it is to be *F* unless one can think of any arbitrary object that it is *F*: the knowledge and the ability are coeval.

⁹⁶ See especially Boysen and Bernston (1995) and Boysen *et al.* (1999). There is a nice discussion of this experiment in relation to animal agency in Hurley (2003, p. 245ff).

of candy and receiving the lesser reward, but rather the fact that she experiences an irresistible incentive which overrides this understanding.⁹⁷ In any case, though, my minimal suggestion is that an animal's inability to learn that causal relations hold between other action-outcome pairs provides *evidence* that the form of representation is not compositional. And if we want to preserve the idea (i) that beliefs are propositional attitudes and (ii) that propositions are compositionally structured, then this kind of representation is not a belief.

2.4 Practical inference

In a way, of course, this is simply a terminological dispute. After all, we can call the rats' representations of the instrumental relationship whatever we like — 'proto-beliefs', 'apprehensions', 'quasi-beliefs', etc. Accordingly, the natural response on behalf of the Standard Approach is simply to adopt one of these less weighty terms — and to chastise me for being pedantic.

However, I think that this terminological point masks a deeper issue. For the Standard Approach does assume that the animal is able to make simple practical inferences: the animal combines its beliefs and desires (or rather, its beliefs and desires combine) to cause the action — or so we are told.⁹⁸ But if the object of an animal's belief is unstructured, how are its states supposed to interact so as to produce the desired end? The logical relations that hold between the contents are supposed to be mirrored, I take it, by the causal relations between the animal's psychological states (to some extent, at least). But then it seems obvious that the object of the animal's *desire* must also occur in the context of its instrumental belief: after all, an animal that wants to do A and believes that doing B will afford it the opportunity⁹⁹ to do A, must appreciate the fact that doing B leads to what it wants to do. If this is correct, however, then it seems as if the Standard Approach *is* committed to the idea that an animal's instrumental representations are structured.

I think that this point will perhaps become clearer if we consider its more familiar application in the case of deductive inference. To this end, consider the following passage from Frege:

⁹⁷ (1999, p. 229)

⁹⁸ Admittedly, I did not build this requirement into CTA. The reason is that I do not think that every causal theorist would hold that a practical inference is everywhere required in order to constitute behaviour as intentional. Rather, I take the idea to be the Davidsonian one that "had [the animal] been aware of them and had he the time, he *could* have reasoned that his act was desirable[.]" Davidson (1978, pp. 85-86). In any case, though, the idea that animals do engage in practical inferences is built into many accounts adopted by cognitive ethologists. See, for example, Clayton *et al.* (2006, p. 199) and Heyes and Dickinson (1990, pp. 88-90).

⁹⁹ I take this to be a type causal relation (the enabling kind). When 'do-ables' (i.e., things people and animals do) are in question, it is borderline unintelligible to think that doing A will cause *the animal* to do B. Hence, my terminological preference.

The task of our vernacular language is essentially fulfilled if people engaged in communication with one another connect the same thought, or approximately the same thought, with the same sentence. For this it is not at all necessary that the individual words should have a sense and meaning of their own, provided only that the whole sentence has a sense. Where inferences are to be drawn the case is different: for this it is essential that the same expression should occur in two sentences and should have exactly the same meaning in both cases. It must therefore have a meaning of its own, independent of the other parts of the sentence.¹⁰⁰

If we transpose Frege's point from the level of words and sentences to the level of sense, then the idea is that the inferential transitions between thoughts depend on their structural elements occurring in different propositional contexts. The proposition that *Socrates is mortal*, for example, is entailed by the proposition that *Socrates is a man* and the proposition that *every man is mortal*. And while it took the genius of Frege to formalise the idea, it is intuitively clear that the entailment depends (*inter alia*) on the fact that that what is predicated of Socrates in the first proposition is the same as that which is predicated of all men in the third.

Frege makes this point with respect to deductive inference, of course, but my suggestion is that a similar point applies in the case of practical inference — at least in so far as the practical inference is thought to involve an instrumental belief with 'causal content'. For if the animal wants to figure out how to do A, it must be capable of representing the enabling relation between doing something else, say B, and doing A.¹⁰¹

Now, I have suggested that Evans's account offers a nice way of cashing this out. For we can understand the structure of thought as an abstraction from the agent's cognitive abilities: in this case, the ability to conceive of something do-able, A, and of something else do-able, B, and the ability to think of two arbitrary do-ables that they satisfy the affording relation. But as we have seen, the possession of these abilities carries expectations that threaten to vitiate the idea that animal representations are structured in any significant sense. In particular, the general-purpose nature of these abilities requires that the animal's thoughts do not violate the Generality Constraint. And this in turn requires the flexible capacity to learn new affording relations in

¹⁰⁰ (1980, p. 115) The passage is from a letter from Frege to Peano.

¹⁰¹ Would it help if we instead represented the practical inference in propositional logic? For example:

Goal: p

Belief: if q , then p

so: do q

I suggest not. In the first place, and as I will argue in detail later (§3.2), the objects of the desire (i.e. goals) are not propositions. For another, one cannot do something propositional, such as q . Finally, the capacity to represent logical constants is no less controversial than the capacity to entertain structured contents (in the context of animal minds, that is). On this point see and cf. Bermúdez (2003; esp chapters 8 and 9), Burge (2010b), Mody and Carey (2016).

different contexts. There can be no ‘conceptual barrier’, as Evans puts it,¹⁰² to the animal’s coming to form the appropriate instrumental belief in these new contexts.

I am sceptical that any species of animal is capable of satisfying the Generality Constraint. I would be more than happy to be proved wrong about this, of course, but in the absence of compelling evidence to the contrary, I shall proceed under the assumption that few — if any — species of animal are capable of forming instrumental beliefs.

However, let me suggest a way in which we might think about a non-propositional *analogue* of instrumental belief. If this proposal is viable, then it suggests a way in which an animal that lacks the abilities associated with propositional content might nevertheless be capable of making practical inferences.

My idea is that some animals may be capable of reasoning purely with representations of do-ables. Thus, we might imagine an animal combining a directive representation of the form ‘Do A’ with the instrumental representation ‘To do A, do B’ — ‘To get a reward, press a lever’, for example.

There are several intriguing features about this proposal that make it *prima facie* attractive. The first is that instrumental representations of this kind do not mention the enabling relation between the do-ables *per se*. They do not *say* that the performance of one action will enable another; rather, they provide a form of instruction that is of use to any animal with the relevant desire. This distinction is subtle, but not, I think, insignificant. For as long as there is no representation of the enabling relation itself, the animal does not need to exercise a *general* ability to think of any two do-ables (of which it has a conception) that they stand in that relation in order to think: to do A, do B. Accordingly, crediting an animal with this sort of representation does not *ipso facto* carry with it the implication that the animal will be capable of learning new affording relations that obtain between other do-ables in different contexts.

The second point worth mentioning is that putting these representations to use in figuring out or deciding what to do does not require the supplementary representation of logical constants. Accordingly, it does not require the animal to be capable of mastering any of the general inference rules that serve to define the logical constants.

As far as I can tell this sort of representation would suffice to explain the rats’ behaviour, as well as Sheba’s: for it provides them both with an instruction about how to achieve a desired end which can be stored away and updated in light of future contingencies. Of course, whether we *need* to invoke an instrumental representation in either case is a further question that I shall not attempt to answer here. But I remind the reader that my purpose is not to defend the

¹⁰² Evans (1982, p. 102)

ascription of instrumental representations to animals: if this proposal cannot be made to work, then so much the worse for the Standard Approach.

2.5 Summary

My purpose in this chapter have been two-fold. In the first place, I have tried to defend the practice of ascribing objective content to animals. I have argued against Davidson's contention that this practice is misguided and, along the way, I made some suggestions about what sort of positive empirical evidence might be brought to bear on the issue. This evidence suggests that objectifying *perceptual* capacities are indeed widespread throughout the animal kingdom.

In the second half of the chapter, I raised concerns about the idea that animals are not merely in perceptual, but also doxastic, contact with the world. This necessitated some groundwork relating to the idea of propositional content and, to this end, I adopted Evans's suggestion according to which propositional content is compositional in a sense that is abstracted from the systematic, productive, and inferential abilities characteristic of mature human beings.

I proceeded to suggest that this raises considerable challenges for the ascription of beliefs to animals. In particular, ascribing instrumental beliefs with causal content seems to presuppose a general grasp of the concept of causality. But this is not likely to be available to animals which show widespread contextual limitations in the range of causal dependencies that they are capable of appreciating. I ended by making a tentative proposal on behalf of the Standard Approach: if we give up on the idea that instrumental representations must include a causal element, we can isolate a non-propositional form of practical reasoning. Needless to say, this proposal is tentative in the extreme. But if it could be made to work, then it offers the Standard Approach a way of circumventing my objection that instrumental beliefs are (probably) not ascribable to animals.

Taking Action for Granted

“All this conspires to make us think [...] that if we wish to understand what intention is, we must be investigating something whose existence is purely in the sphere of the mind; and that although intention issues in actions, and the way this happens also presents interesting questions, still what physically takes place, i.e. what a man actually does, is the very last thing we need to consider in our enquiry. Whereas I wish to say that it is the first.”¹⁰³

3.0 Preliminaries

In the previous chapter I was concerned with questions of content. This was the inevitable consequence, I think, of adopting a ‘mind-first’ approach in the context of animal agency — a concession that I made to the Standard Approach in order to assess it on its own terms. It is therefore unsurprising, if a little odd, that there hasn’t yet been much discussion of what actually happens when animals act — of the sort of thing that animals *do*. In keeping with the Standard Approach, I set these matters to the side on the supposition that the characteristic mark that distinguishes actions from mere events lies in the animal’s psychology. What is done by the animal — which, according to the Standard Approach, *is* a mere event — has thus far been taken for granted.

Interestingly, it is this sort of approach that Anscombe cautions against. In discussion of the problem posed by deviant causal chains, for example, she suggests that the difficulty arises as a result of

[...] the standard approach by which we first distinguish between ‘action’ and what merely happens, and then specify that we are talking about ‘actions’. So what we are considering is

¹⁰³ G.E.M. Anscombe (1963, p. 9)

already given as — in a special sense — an action, and not just any old thing which we do, such as making an involuntary gesture. Such a gesture might be caused, for example, by realizing something (the ‘onset of a belief’) when we are in a certain state of desire. Something I do is not made into an intentional action by being caused by a belief and desire, even if the descriptions fit.¹⁰⁴

It is clear from the context of this passage that Anscombe has in mind Davidson’s example of the mountaineer who *wants* to be free of the weight of the second man on the rope, *realises* that he can free himself by releasing his grip, and is so unnerved by the thought that he shudders, loosens his grip, and thereby brings about the object of his desire. And while the passage is terse, her point seems to be that the difficulty that causal theorists face with respect to examples such as this lies in their assumption that actions are distinguished from ‘mere events’ by some specifiable additional feature (e.g. causation by beliefs and desires). For if ‘sudden starts’ are categorically distinct from intentional actions — if the material process is itself of a different nature — then it doesn’t matter what causes them: the *mere event* of loosening one’s grip as a result of being unnerved by the onset of a belief shares but a name with the *action* of loosening one’s grip intentionally.¹⁰⁵

More generally, in *Intention* Anscombe remarks that

If one simply attends to the fact that many actions can be either intentional or unintentional, it can be quite natural to think that events which are characterisable as intentional or unintentional are a certain natural class, ‘intentional’ being an extra property which a philosopher must try to describe.¹⁰⁶

[But t]his 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.¹⁰⁷

Recall Wittgenstein’s question about what is left over if one subtracts the fact that something happens from the fact that you *do* it. As I said in §1.3, this question presupposes that what happens when someone raises their arm, say, is the same sort of thing as what happens when

¹⁰⁴ Anscombe (1989)

¹⁰⁵ Candace Vogler (2016, p. 239) reads Anscombe in this way. As she puts it, “The whole tendency to worry about wayward or deviant causation of actions presupposes that a single sort of event—for example, an action (e.g., crossing the street) comes about in an unusual way (e.g., as a result of a seizure). For Anscombe, there is no such thing as an action, *A*-ing (e.g., crossing the street), that could happen in any number of ways — through intention, during a seizure episode, as a result of a mighty wind, through elaborate electrical means involving manipulation of the pedestrian’s muscles having paralyzed him with a poison dart gun just as he began to cross the street, etc.— such that we might need to explain away many street-crossings as actions that were not properly intentional.”

¹⁰⁶ (1963, p. 84)

¹⁰⁷ (1963, p. 88) The feature that Anscombe is considering in this passage is specifically ‘practical knowledge’. But is clear that she thinks that *no* feature can play the role of an ‘extra feature’ distinguishing intentional actions from ‘mere events’.

their arm is raised by something else; it presupposes, that is to say, that intentional actions are metaphysical compound made up of mere events and their distinctive source in the mind. As before, however, the central message of the passages that I have just quoted is that this assumption is false; the principle of composition — the idea (i) that intentional actions are distinguished from ‘mere events’ by something external to them and hence (ii) that we can understand what an intentional action is by discovering this additional feature — cannot be maintained.

In this chapter, my aim is to provide independent support for Anscombe’s conclusion. Doing so serves two functions in the overall structure of my thesis. In the first place, it constitutes a more general argument against the Standard Approach than the concerns that I have raised regarding the applicability of propositional content to animals. The point will be to show that *even if* propositional content can be applied more broadly than I have claimed, the idea that the Standard Approach gives us an understanding of animal agency will still be undermined (the argument will also show, I think, that adopting the tentative proposal muted at the end of the last chapter is not enough to save the approach). The second function that the argument serves is to demonstrate just how radically different an Anscombean approach to animal agency will have to be. And this serves to prepare the way for the positive account to be developed in the remaining chapters.

I want to begin, however, by discussing a different interpretation of Anscombe’s position. On this reading, there *is* an additional feature in virtue of which actions count as intentional, namely, the fact that the agent grants the application to ‘the special sense’ of the question ‘Why?’ Although this reading is both philosophically and exegetically implausible, it will serve as a useful introduction to some of the problems to be raised later in the chapter.

3.1 *Intention*, deflated

According to what I will call a the ‘deflationary’ reading of Anscombe’s *Intention*, (a reading championed, at one point, by J. David Velleman) Anscombe is a neo-behaviourist who “tried to identify an overt language-game in terms of which intention [and intentional action] could be understood.”¹⁰⁸ What defines the language-game are the various ways of granting (and rejecting) application to the question ‘Why?’ as it is addressed to the agent in the midst of her deed. And it is no doubt true that Anscombe does isolate her subject matter — intentional action — by describing various ways in which the question ‘Why?’ is granted and refused

¹⁰⁸ Velleman (2004, p. 228). The context of Velleman’s remark makes it clear that his reading extends to Anscombe’s treatment of intentional action as well.

application.¹⁰⁹ This allows her to define a ‘special sense’ of the question and then to equate intentional actions with those actions for which the question is granted application in the sense thus defined. For example, she says that the question is *denied* application by the answer ‘I didn’t know I was doing that’ whereas it is *granted* application by a response that cites a ‘reason for acting’.¹¹⁰ Thus we might imagine the following exchange:

A: ‘Why are you breaking those duck eggs?’

B: ‘Am I? I didn’t realise that they were duck eggs.’

A: ‘Well, why are you breaking eggs anyways?’

B: ‘I am baking a cake.’

In this case the explanation as to why B is *not* breaking duck eggs intentionally, but is breaking eggs intentionally, is that the question ‘Why?’ is granted application under one description but not the other. And so on more generally: intentional actions are those events for which the question ‘Why?’ receives (or at least would receive) an answer falling within a certain range. This, in a nutshell, is Anscombe’s theory of intentional action — or so the deflationary reading would have it.

Notice, however, that isolating a subject-matter is not the same as offering an account of the phenomena that constitute it. After all, it is common enough to *introduce* a topic by circumscribing it. An author might restrict the topic of their book to the ecology of the naturally occurring fauna of the Pacific Northwest, for example. Doing so explains why they are going to be talking about certain species of animal and not others, but it is obvious that it is the account that follows that constitutes the author’s account. Similarly — and this is probably a closer analogy to Anscombe’s project — at the beginning of the *Foundations of Arithmetic* Frege says that his subject matter will be limited to the positive whole integers — to those numbers “which give the answer to the question ‘How many?’”¹¹¹ Like Anscombe, then, Frege distinguishes the sort of thing that he will be discussing from other, related topics (complex and irrational numbers, fractions, etc.), by means of a question. But one will clearly misunderstand Frege’s work if one takes his account of the natural numbers to be given *simply* by reflection on the question ‘How many?’ His theory is to be found in what comes after the first few pages.¹¹²

¹⁰⁹ See, for example, (1963, p. 28)

¹¹⁰ This is not to say that it cannot be granted application in other ways. It can, according to Anscombe. See (1963, pp. 15-28).

¹¹¹ Frege (1953, p. 5n1)

¹¹² The parallels between Anscombe’s and Frege’s respective questions is usefully discussed by Ford (2015). As he brings out, one of the apparent disanalogies is that Anscombe, unlike Frege, has to speak of a ‘special sense’ of her question. But this dis-analogy is somewhat coincidental, according to Ford: had Frege been writing in a language where a single term is used to ask for both ‘How many?’ and ‘How much?’ (like French)

This means that we cannot infer from Anscombe's talk of isolating her topic by means of the question 'Why?' that, having thus isolated it, her *account* of intentional action is thereby complete. On the contrary, if her project is anything like that of Frege's, then we should expect it to develop from this point onwards.

That is one reason to be sceptical of the deflationary reading of Anscombe. There are many other reasons for rejecting it.¹¹³ But let us set these concerns to one side for the time being and assess the account on its own merits. Can we explain what intentional actions are by means of considerations about what people would say if asked to explain themselves in various circumstances? I think that the answer is clearly 'No', but it will be illuminating to see exactly where the account goes wrong.

Returning to the example given above, notice that the reason that the agent gives in explanation of her act of breaking eggs is expressed as follows: 'I am baking a cake'. If we bring the *explanans* into explicit contact with the *explanandum*, we can derive an example of what Michael Thompson calls 'naïve action explanation': 'I am breaking eggs because I am baking a cake' or, more generally, 'I am doing A because I am doing B'.¹¹⁴

Explanations of this kind work by bringing an action within the scope of a wider project. The agent rationalises her action by re-describing it in such a way as to show the further 'intention with which' (in Anscombe's terminology) it is being done. In so doing, the agent "gives the action a place in a pattern, and in this way the action is explained", as Davidson eloquently put it.¹¹⁵ The trouble that this poses for the reductive ambitions of the account, however, is that the description that appears in the *explanans* would seem to refer to something which is itself an intentional action (in progress), namely, the baking of a cake. But if so, then how are we to understand what is being said without presupposing the notion that we are trying to explain?

The difficulty that I want to draw attention to is perhaps more familiar from other attempts to reduce philosophical concepts to mastery of the relevant linguistic rules. Consider, for instance, the attempt to reduce the notion of self-consciousness to the ability to think *I*-thoughts — thoughts that are 'immune to error through misidentification' — and to explain these in turn in terms of the mastery of the semantic rules governing the use of the first-person pronoun, 'I'.

he, too, would have had to speak of a special sense of the term. Conversely, if Anscombe had been speaking in a language that explicitly separated the 'reason requesting' sense of the question from others, she would have had no use for the qualification.

¹¹³ For example, this reading is presumably not *so* simplistic as to entail that the agent cannot tell a lie. So we must now add to this the idea that the agent's answer is truthful. But now, what is it for the answer to be truthful? Anscombe's own view on this matter is clear: although there exist observable criteria for the truthfulness of someone's answer up to a point, once these give out there is truth in the statement that 'only he can know what he intends' (1963, p. 44). How are we to make sense of this on a deflationary (behaviourist) reading?

¹¹⁴ See Thompson (2008, Part II)

¹¹⁵ (1963, p. 10)

The success of this project obviously hinges on the question of whether we can explain the meaning of ‘I’ without surreptitiously invoking the notion of self-consciousness in so doing. And philosophers have indeed held out hope that this might be done. Thus we might try to explain what ‘I’ means through the use of some rule such as this: ‘Whenever someone uses the word ‘I’, she refers to herself’. Now there is no doubt that this rule is *true*. But does it explain first-person thought?

The well-known difficulty facing proposals of this kind is tied to the role of ‘herself’ in the *explanans*. It is obvious that it cannot be the ordinary reflexive pronoun, since it is possible to make reference to oneself by means of it and yet not to realise that one is doing so. To borrow John Perry’s famous example, Perry believes that the person making a mess on the supermarket floor is a nincompoop, and thus believes something of himself, without realising it. This is possible because he does not know that *he* is the person making a mess. But then it cannot be *this* sense of the reflexive pronoun — i.e., the one that allows for this possibility — that is being used in the rule above. As Anscombe notes, “[...] we are inclined to think that ‘It’s the word each one uses in speaking of himself’ explains what ‘I’ names, or explains ‘I’ as a ‘referring expression’. [But i]t cannot do so if ‘He speaks of himself’ is compatible with ignorance and we are using the reflexive pronoun, in both cases, in the ordinary way.”¹¹⁶ It seems rather that the pronoun is being used in a special sense that rules out the error through misidentification; that is, it is being used in such a way that it is not possible for ‘himself’ to refer to some object, *a*, and for the user to fail to recognise that *I* = *a*. But now we are explaining this special sense of the reflexive pronoun in terms of the capacity to think *I*-thoughts! Thus Anscombe complains, “it is no explanation if that reflexive has in turn to be explained in terms of ‘I’; and if it is the ordinary reflexive, we are back at square one.”¹¹⁷

Now the point of discussing this case is emphatically not to broach the topic of self-consciousness. Nor does anything in my argument turn on whether a deflationary account of first-person thoughts can in fact be made to work. The point is just that a similar *prima facie* circularity is to be found in the attempt to explain the notion of intentional action in terms of overt speech-acts of the form ‘I am doing A because I am doing B’. For although there is a sense of ‘I am doing B’ that does not presuppose the notion of intentional action — just as there is a sense of the reflexive pronoun that does not presuppose the sense of ‘I’ — it is precisely *not* this sense that enters into rationalisation. Or so I claim.

To test this, we need to consider whether it is possible to specify truth conditions for rationalisations of the form ‘I am doing A because I am doing B’ without invoking the notion

¹¹⁶ (1975, p. 22)

¹¹⁷ (1975, p. 23)

of intentional action in question. In what must be the paradigmatic sort of case (on this view), an agent rationalises her action by means of a description that can represent both intentional actions and mere happenings. So, for example, I might be moving my hands thus and so because I am making shadow puppets on the wall, on one occasion, and, on another, making shadow puppets on the wall without meaning to. Since both cases are clearly possible, it would seem that there is no circularity about explaining what sort of circumstances the description ‘I am making shadow puppets’ is true: they are just those circumstances in which I produce some effect which causes them to appear on the wall.

But, of course, it is not simply these descriptions in isolation that we need to understand but rather what is represented when they are conjoined by the rationalising sense of ‘because’. And as the example just given demonstrates, it is certainly not enough that I am doing A and that, in so doing, I am also doing B. The problem here is not one of circularity, but rather one of obvious falsity: I may be moving my hands thus and so and, in so doing, making shadow puppets, without moving my hands thus and so *because* I am making shadow puppets — perhaps I am thumb wrestling and just happen to produce them (e.g.).

Nor will it do to say that I am doing A and, in so doing, ‘furthering my project’ (‘contributing to the realisation of my goal’ etc.) of B-ing. For it is possible to further one’s project unintentionally: Perhaps it is one of my long-term aims to impress the next-door neighbour with my shadow puppet displays. Believing her to be away on vacation, I decide to practice my technique. As it happens, her flight got cancelled and she is watching through the window. In such a case I am furthering my project of impressing her (let us suppose), but I am not doing so intentionally. If it is replied that ‘furthering one’s project’ has a special sense in this context then it behoves the defender of the deflationary approach to explain this sense without illicitly smuggling in the notion we are attempting to explain. In the absence of such an explanation, however, I think that we must conclude that the deflationary account is dead in the water.

3.2 The objects of desire

The rationale behind the forgoing digression may seem oblique. After all, the mistake of reading Anscombe as a neo-behaviourist has been diagnosed on numerous occasions¹¹⁸ and there are arguably more obvious problems facing the deflationary account than the one that I have laboured to bring out. What is to be gained from pointing out the obvious and then attacking a straw-man? And, in any case, what does all of this have to do with the Standard Approach to animal agency?

¹¹⁸ See, for instance, Vogler (2016) and Moran and Stone (2011).

I want to suggest that the problems facing this caricature of Anscombe actually recur, more subtly, for the Causal Theory of Action (CTA). Of course the causal theory is not a behaviouristic theory: as we have seen, its defenders hold that the *differentia* separating ‘mere events’ from intentional actions are to be found, not in what the agent would say, but in the beliefs, desires, and intentions that cause them to act as they do. CTA is nevertheless deflationary in the sense that it attempts to reduce the notion of intentional action to the concepts of event, cause, and belief, desire, and (on some views) intention.¹¹⁹ The moot question is whether moving the *explanans* inside the mind avoids the circularity inherent in the account just discussed.

Now one reason to think that it *does* avoid the circularity is that it would appear to specify truth conditions for sentences like ‘I broke the duck egg because I am making a cake’ that do not invoke the notion of intentional action. For, on this account, the worldly condition that would make this sentence true is one in which (i) my desire to make a cake and (ii) my belief that I could make a cake by cracking the duck egg in the bowl (iii) conspired to *cause* me to break it. And it would not seem that we need to understand what an intentional action is in order to understand the (i), (ii), or (iii).

Picking up on a train of thought recently elaborated by Matthew Boyle and Douglas Lavin,¹²⁰ however, I am going to suggest that the appearance that CTA avoids the circularity inherent in the deflationary reading of Anscombe is illusory. But to get a sense of the issue, we need to turn our attention to the objects of the representational states favoured by causal theorists.

The thing to notice is that the objects of belief and desire differ not simply in terms of content, but also in the form that the content takes. Returning to the cake example, the belief is directed towards a propositional object — *that I can make a cake by* (inter alia) *cracking a duck egg in a bowl* — whereas the object of desire is, *prima facie*, something different; it is not a complete thought, but rather the fragment of one: *to bake a cake*. And, indeed, this is rather typical: although states of desire can be directed towards material objects, like pieces of fruit, or, if forced, proposition-like objects (e.g. that I bake a cake),¹²¹ by far the most natural interpretation of the desires that motivate us *to act* are given in terms of the infinitival phrase ‘To do A’. But now, how are we to explain this phrase? Under what conditions would a desire to do A be *satisfied*? The answer had better not be: ‘When, and only when, the agent does A intentionally’. For a compositional

¹¹⁹ Not all causal theorists are deflationists in this sense. As mentioned in the introduction, Davidson, e.g., is not. I shall consider the tenability of Davidson’s position in the next section.

¹²⁰ See their (2010). These ideas expand upon a line of thought found in M. Thompson (2008, pp. 120-128).

¹²¹ I say ‘propositional-like’ since what follows ‘that’ is in the subjunctive and hence not truth-evaluable.

theorist who wants to deflate the notion of intentional action cannot help herself to a prior understanding of that notion.

This problem is masked, but not resolved, by the widespread tendency of causal theorists to assume that desires can be brought within the general rubric of ‘propositional attitudes’. Neil Sinhababu, for example, says that “the conditions under which the desire is satisfied must be the same as those under which the relevant proposition is made true”¹²²; similarly, Michael Smith describes the ends of desire simply as “ways the world could be[.]”¹²³ As Boyle and Lavin point out, this assumption allows causal theorists to paper over the intuitive thought that the worldly condition set by the sentence ‘I want to do A’ is, precisely, that of an occurrence which will only be realised when a certain intentional action has been completed. Instead, the condition that a desire must ‘fit’ is presented as a generic sort of fact, one which can be brought about in various ways, intentional or otherwise. *If* the worldly condition is realised through intentional action, then that is simply a further fact about it.

But what *is* the worldly condition that would render the proposition that corresponds to a desire to do A true? As before, the question is whether we can specify a state of affairs that can serve as the satisfaction condition of this state *without* appealing to the notion of intentional action itself. The first thing to notice about attempts to do so is that, by insisting on the inclusion of a propositional clause, we are forced to specify an aspect and tense for the verb. To take a new example, it would appear that we have roughly four options for specifying the state of affairs that would satisfy my desire to turn off the lights:

- [1] *I turn off the lights* (present perfective)
- [2] *I am turning off the lights* (present progressive)
- [3] *I have turned off the lights* (past perfect)¹²⁴
- [4] *I was turning off the lights* (past progressive)

It seems pretty clear that neither [2] nor [4] will do, since both are consistent with a state of the world in which the lights never go off — I may be in the midst of turning off the lights when I get interrupted by a heart attack, for example. This is surely not a case in which the desire has been satisfied.

¹²² (2015, p. 166) Curiously, Sinhababu goes on to say that “We also might say of Jenny that she desires to eat chocolate, wear a diamond necklace, keep a goldfish in her aquarium, and make love to Johnny Depp. *One might regard it as a corresponding virtue of Propositionalism that it fits these attributions.*” (2015, p. 168; my emphasis) As will become clear below, there is no meaning-preserving way to transform the infinitival ‘To do A’ into a ‘that-clause’.

¹²³ (2004, p. 165)

¹²⁴ Alternatively, we might use the simple past tense ‘I did A’. Since this would not affect the proceeding considerations, however, I ignore it.

[1] is also clearly wrong. For the combination of the present tense with perfective aspect leads us to a habitual reading of the propositional clause. That is, the state of the world which corresponds to ‘I turn off the lights’ is one in which I am in the *habit* of turning them off. But although this is something which I can certainly desire to develop, it is obviously different from that which is specified by a desire to turn off the lights.¹²⁵

This brings us to [3]. Now it is true, of course, that if my desire is satisfied, then the state of affairs here represented will have come about as the result of something I did. But it matters how. For suppose that I am on my way over to the light switch when you push me into it, thus turning off the lights. In that case, it is true that I turned off the lights; I am the ‘agent’ of this deed in the trivial sense that I am the subject of an active causative verb. But this would presumably not be a case in which my desire has been satisfied, any more than would a case in which *you* turned off the lights for me. Indeed, this sort of case makes it especially clear that an object of desire is *not* the sort of thing that can be brought about anyhow, but is rather sensitive to the contribution of the agent in question.

I think that this point is intuitively clear, but it can be reinforced by changing the example. Thus consider the difference between someone, like me, who wants to finish their PhD, and someone who would certainly like to have a PhD, but does not want it enough to write a thesis. Both of us want a certain state of affairs to obtain. But I have a further desire that he lacks. For I want to complete it *as a result of my own effort* — a desire that would not be satisfied if the university were to award me one as the result of an administrative error (say). If desires are to move us to act, they cannot be indifferent to our own agency.

Perhaps this will prompt the suggestion that the problem can be overcome if we combine [3] and [4] into [5]: *I was turning off the lights and have now turned them off*. But again: I *was* turning off the lights, when you pushed me (‘Why are you getting up from your chair?’ — ‘I am turning off the lights’),¹²⁶ and *have now* turned them off. As Boyle and Lavin bring out, the problem with this suggestion is that the desire to do *A* is not simply to have been in two unrelated states, “but to have arrived in the end state in virtue of having done all of the *A*-ing that was required to *A*.”¹²⁷ And if it is suggested that we can avoid the issue by including the idea that a requisite amount of *A*-ing is necessary, then it needs to be explained how much *A*-ing is required. If I get thrown across the room upon getting up from my chair, for example, I don’t need to have been turning off the lights for very long. But this is still not a case in which my desire has been satisfied, any more than my desire to finish my PhD will be satisfied if I write half of it before the

¹²⁵ This point is familiar. See, e.g., Thompson (2008, p. 125).

¹²⁶ This is an example of what linguists call the ‘broadness’ of progressive aspect. I discuss the phenomenon of aspect in more detail in the next chapter.

¹²⁷ (2010, p. 172)

administrative error occurs. No, I want to do just enough to bring the project to completion. But how are we to explain what ‘enough’ means here without invoking the idea of intentionally realising one’s end?

Perhaps more imaginative alternatives will be proposed. But it is difficult, at least for me, to see what further options are available to a causal theorist at this point. For once we have acknowledged that the desire is not indifferent to the contribution of the agent, then it seems that the only way in which it can be satisfied is if the world comes to conform to it as the result of intentionally bringing it about. As Boyle and Lavin put it, “To represent my doing *A* is to represent, as it were, a kind of state of affairs whose obtaining *is* my having intentionally caused it to be.”¹²⁸ At any rate, in the absence of any further suggestions to the contrary, we are justified in concluding that the concepts of belief, desire, cause, and event, are incapable of explaining the notion of intentional action.¹²⁹

3.3 On the significance of the foregoing argument

I think that this is a significant and interesting conclusion. Nevertheless, the response that I anticipate to this line of thought is a shrug of the shoulders, so to speak: ‘Perhaps it is true that we cannot reduce the *concept* of intentional action to any more fundamental notions. But this lesson was learned a long time ago. Indeed, it was precisely this moral that Davidson himself drew from the problem of deviant causal chains.’¹³⁰ Moreover, it doesn’t much matter. For just as Davidson accepted that it is not possible to specify the ‘right way’ that beliefs, desires, and intentions must conspire to cause intentional actions without giving up his idea that intentional actions *are* events that are caused, in the right way, by such mental states, the argument just given provides no reason to replace this metaphysical picture. The conclusion licensed by the argument, if sound, is one about *explanation*; it cannot support the *metaphysical* claim that there is *no* additional feature in virtue of which mere events differ from intentional actions; all it shows is that we cannot describe this feature in a non-question begging way.’¹³¹

Let me begin my response to this objection by agreeing that the argument does not entail the conclusion that the compositional assumption characteristic of the Standard Approach is false. Indeed, this much is obvious since there are other theories of action that subscribe to the

¹²⁸ (2010, p. 173)

¹²⁹ Would it help if we added the notion of intention in this list? I suggest not: for we can just repeat the argument with respect to the objects of intent which, like the objects of desire, are represented by infinitival clauses. Indeed, there are arguably further problems posed for intention. See Annette Baier (1970, 1977) and Lucy Campbell (2015, 2018).

¹³⁰ See Davidson (1987b, p. 108): “Several clever philosophers have tried to show how to eliminate the deviant causal chains, but I remain convinced that the concepts of event, cause, and intention are inadequate to account for intentional action.”

¹³¹ Thanks to Arif Ahmed for pressing me to consider this objection.

compositional assumption but reject the specific version endorsed by CTA.¹³² The argument poses no threat to such theories. Nevertheless, CTA remains by far the most popular theory in both philosophy and the cognitive sciences, and I am going to assume that it is the most *plausible* of the compositional approaches.

Even if this assumption is granted, however, I accept the objector's point that the inability to explain the *differentia* of intentional action in a non-circular way does not straightforwardly entail that the metaphysical picture of actions as the compound of bodily movements caused in some right way by mental states is false. This only follows (by *modus tollens*) if we make the supplementary supposition that *if* actions are a metaphysical compound *then* they can be analysed in terms of the concepts of their components. And I admit that I do not know how to show this, at least not conclusively. But I do think that there are reasons for endorsing the implication. Let me describe two.

The first is that the position that the Davidsonian objector seems willing to occupy is a distinctly unhappy one. For compare other cases: a sunburn is (i) a burn that is (ii) caused by the sun. There is no problem in explaining what a sunburn is in terms of the concepts of its components, precisely because being caused by the sun is an 'additional feature' that burns may or may not have. More specifically, such explanations work by specifying the *genus* to which sunburns belong (namely burns) and then explaining the 'additional feature' that separates the species from other members of the genus. To take a different example, consider the jumps that one experiences when one drifts to sleep. These 'hypnagogic jerks' are (i) muscle spasms that (ii) occur as one is drifting off to sleep. As before, we can explain the phenomenon in terms of the feature that distinguishes muscular spasms of this kind from other sorts of spasm without any threat of circularity.

These examples show a *prima facie* connection between the compound nature of a phenomenon and its explicability in terms of the concepts of its components. In both of these cases, there is a single sort of thing — a muscle spasm, a burn — that is brought about in a certain way, or in a certain context, but which we could easily imagine being brought about otherwise: the same sort of muscle spasm might be brought about artificially, when one is wide awake, just as an indistinguishable burn might be produced by artificial light. And similarly, the objector would have it, an arm rising (say) is the same sort of thing whether it is brought about intentionally or not. But if these cases are really analogous, then it is puzzling that we cannot explain the 'further feature' that distinguishes intentional actions from mere events, as we could in the other cases.

¹³² I am thinking, for example, of J. Hyman (2015; esp chapters 2 and 3).

The response that I anticipate here is that the cases that I have mentioned are dis-analogous insofar as the concept of intentional action is part of a holistic network of interrelated psychological concepts. It is no surprise that we cannot analyse the former by means of the latter because these concepts are in general interdependent.¹³³

But the holism of the mental is supposed to affect the way in which we understand the *content* of these states, not the states themselves. Holism about the mental is the view that one cannot have a belief about cats, say, without believing various other things about them: that they are animals, that they have fur, etc. And while Davidson certainly thinks that this poses a problem for the idea that mental state concepts can be reduced to *physical* concepts,¹³⁴ it is unclear why this should have any effect on intra-psychic reductions. Why couldn't one accept both that content is determined holistically and that some psychological phenomena can be analysed in terms of others which possess certain further features? J. David Velleman, for example, maintains that intentions are "self-fulfilling expectations that are motivated by a desire for their fulfillment and that represent themselves as such."¹³⁵ If this is right, then intentions are beliefs that are qualified in a certain way. And while this idea may be independently problematic, it is unclear why the holism of the mental should constitute an objection to it. If holism presented a *general* problem for analysing psychological concepts in terms of one another, however, then it surely would.

So without further argument, I do not accept the alleged dis-analogy between the case of intentional action and the others given above. And if they are analogous, then the position that the objector seems willing to accept is a problematic one.

The second reason why I think that genuine compounds should be capable of being analysed in terms of their components has to do with the fact that the *inability* to analyse a certain phenomenon into (i) the genus to which it belongs and (ii) some additional feature that distinguishes it from other species is actually quite characteristic of phenomena that lack a compound character. No one thinks, for example, that the colour red is a compound of colour and some further quality. It is not as if we could *subtract* the fact that an object is red and thereby be left with its colour *simpliciter*. Rather, as A.N. Prior notes "The colour of what is red is its redness; and the colour of what is blue is its blueness; we can say that the red and the blue agree in being coloured, but of their difference we can only say either that their colour is different, or

¹³³ I anticipate this response since, according to Davidson, the holism of the mental explains why the problem of deviant causal chains cannot be solved. In order to solve it, on his view, we would need psychophysical laws — laws which he does not believe exist. See Davidson (1973, p. 80); see also G. Keil (2007) for discussion of this theme in Davidson.

¹³⁴ See especially Davidson (1970).

¹³⁵ (1989, p. 109)

that one is red and the other blue.”¹³⁶ Following tradition, Prior says that colour is a *determinable* concept and that red and blue are *determinates* of it.

The relation between a determinable genus and a determinate species is interesting because, in contradistinction to the cases mentioned earlier, the qualities which distinguish the species from other members of its genus cannot be grasped independently from their application to the species itself. We can say that the redness of a colour distinguishes it from others, if we like, but then we cannot explain what this quality amounts to without reference to precisely what was meant to be distinguished in the first place, i.e., the colour red. But if the properties which distinguish a determinate species from a determinable genus can only be grasped in virtue of their application to the former, then we obviously cannot appeal to them in an effort to *analyse* the species.

In an illuminating discussion of the determinable-determinate relation and its relevance to intentional action, Anton Ford brings out that the issue turns, at least in part, on the fact that certain species are conceptually more fundamental than their respective genera. Something is a tree, for example, because it is an oak, or a pine, or a beech, say, but it is not one of these because it is tree that has certain additional features. Consequently:

If *all* that we know is that A is a genus of which B is a species, it remains to ask: “What is the source of the truth of it that, *as we know*, every B is at the same time also an A? Is it the case that something is a B because it is, among other things, an A? Or is it, perhaps, that something is an A because it is, for instance, a B? Does the species in question transpire from its genus, or the genus from its species?”¹³⁷

If Ford is right, then in cases in which the species is conceptually prior to its genus, we cannot hope to understand the former by means of adducing additional features that the latter may or may not possess. Ford’s suggestion is that it is precisely this relation that holds between the concepts of action and event. I have argued for a weaker version of this claim: we cannot analyse the concept of action into the concept of event plus the additional idea that the event be caused by beliefs and desires. And if actions are *not* metaphysical compounds, then this is exactly the result that we should expect.

3.4 The consequences of taking action for granted

I do not claim that these considerations are conclusive. But I do think that they provide some independent support for Anscombe’s contention that being intentional is not a mere ‘additional

¹³⁶ A. N. Prior (1949, pp. 5-6)

¹³⁷ Ford (2011, p. 87)

feature' of events 'whose description would otherwise be the same', and, hence, for her view that the *action* that occurs someone loosens their grip on the rope (in order to free themselves of the weight) shares but a name with the event that occurs when that happens as the result of becoming unnerved.

Spelling this idea out in a positive way is the challenge of the next chapter. The negative conclusion, however, is that the Standard Approach's practice of taking action for granted — the practice, that is, of treating 'what merely happens' when animals act as something that does not require an account of its own — is illicit. Indeed, the considerations canvassed above suggest that, if anything, the order of priority is the reverse: in order to understand what it is to desire *to do* something we must first have some conception of the things that animals do.

But at this point it might reasonably be demanded that I explain where taking the action for granted misleads us. *What* is so problematic about the assumption that actions are just those bodily movements that are rationalised by the agent's mental states? *Which* facts do we thereby fail to understand?

Before drawing this chapter to a close, then, let me respond to this challenge.

* * *

Consider the activities¹³⁸ characteristic of animal life: hunting, foraging, fishing, mating, building nests, scavenging, swimming, chasing, playing, fighting, dancing, fleeing, stalking, drinking, eating, protecting one's young, hiding... These are all examples of the sort of thing that animals of a diverse range of species do: bees and birds dance, lions and sharks hunt, dolphins and dogs play, ants and pigs forage, and human beings do all of the above. If we abstract from the differences between the various activities, as well as the disparate ways in which they are undertaken by diverse species, then what unites these examples is, first of all, the idea of a process that has parts: first the animal does this, and then that, and then something else, and so on and so forth, until they are interrupted, or give up, or start doing something else,

¹³⁸ 'Activity' is sometimes given a technical meaning that I do not intend. According to Vendler (1957), for example, activities are (i) 'homogenous', in the sense that if S was doing A over some temporal duration, then it was true that S did A over every moment of that duration, and (ii) involve no culmination or anticipated result. Actions, by contrast, are said to be non-homogenous and do involve an anticipated result.

I do not distinguish actions from activities in this way (though I will provide the vocabulary for marking the distinctions that Vendler wants to emphasise in the next chapter). I treat them as equivalent and both as expressing an ambiguity between *what* is done and someone's *doing of it*. Although many authors prefer to distinguish these as 'acts' and 'actions' (e.g. Hornsby (1999, pp. 623-625)), I prefer to let context disambiguate. This may seem infelicitous, but the reason for it is that I have to make constant reference to authors who abide by different conventions or, indeed, none at all (Anscombe is the latter).

or bring the activity to its natural conclusion.¹³⁹ Secondly, these examples are unified by the fact that the stages of an animal activity can be linked in a way that is not dissimilar from the sort of ‘naïve rationalisation’ considered above. Thus we might explain *why* the lioness is following the wildebeest, for example, by pointing out that she is hunting. And similarly: the answer to the question why the polar bear is chasing her cubs might be: ‘Because she is playing with them.’

I shall have much more to say about all of this in the next two chapters. But for present purposes this sketch is enough to raise the question that I want to pursue, namely, how we are supposed to make sense of these commonplaces with the theoretical tools that the Standard Approach has at its disposal.

The first thing to notice is that the fact that animals engage in activities is in *prima facie* tension with the assumption that actions are to be identified with the animal’s singular bodily movements. For the events that correspond to completed activities are complex.¹⁴⁰ The event corresponding to the sentence ‘The scrub-jay cached four crickets’, for example, is not equivalent to any one movement made by the bird. Rather, it corresponds to an event that was brought to completion through a *series* of movements: first she picked them up with her beak, and then she flew to the caching site (itself a complex activity that involved moving her wings, in unison, over and over again), and then she buried them (by pecking at the ground, etc.).

Nevertheless, and somewhat curiously, it is true that by picking up the crickets she *was* caching them, and indeed, that she was picking them up *because* she was caching them. But how are we to appreciate these points on the Standard Approach? In the examples favoured by causal theorists, a person moves a finger, flips a switch, and thereby turns on a light; or they depress a button, launch a torpedo, and sink the *Tirpitz*; or they raise their arm and, in so doing, cast a vote. In such cases the problem is masked since causal theorists can appeal to the accordion effect in order to explain why it is true that someone *is* turning on a light *by* flipping a switch (the explanation, recall, is that bodily movements can be re-described in terms of their effects).¹⁴¹ But our scrub jay did not *cause* the crickets to be cached by picking them up; nor did picking them up thereby cause her to flap her wings. Picking up the crickets was a *part* of caching

¹³⁹ The natural conclusion of an activity is a point before which it is incomplete and after which it is finished. In the case of hunting, for example, the natural conclusion is the catching and killing of prey. Not all of these examples have a natural conclusion. Playing, for instance, does not.

¹⁴⁰ Most causal theorists take singular bodily movements to be basic in the sense that they are not composed of further actions. I shall grant this idea for the time being.

¹⁴¹ Admittedly, raising one’s arm to cast a vote does not fit this model. The relation between raising one’s hand and casting a vote is conventional, not causal. Still, this sort of case seems relatively easy for the causal theorist to accommodate.

them, but getting the whole thing done required doing other things that were not accomplished thereby.¹⁴²

It will not do for the causal theorist to respond by pointing out that complex actions are brought off through a *series* of bodily movements. For not every series of movements constitutes a unity: Picking up the crickets is part of the scrub-jay's activity of caching them (she is picking them up because she is caching them), whereas her subsequent feather ruffling is not; temporal contiguity is not enough. Nor will it do for them to answer that the difference is that collecting crickets contributes to the fulfilment of that end, whereas ruffling does not. For the end *may* be expedited quite accidentally: if the bird drops the cricket mid-flight and the cricket lands on the head of another jay, thus scaring it and causing it to flee, the jay may be able to retrieve the cricket and complete its caching more quickly (and without the watchful eye of a pilferer). Nevertheless, dropping the cricket does not bear the sort of instrumental relation to the activity as a whole that would allow us to say that she dropped it *because* she was caching.

The suggestion at this point will no doubt be that this is precisely why we need to appeal to beliefs and desires (and, perhaps, intentions): it is because she was picking up the crickets *as a result of* her desire to cache them and her belief that doing so would contribute to the furthering this end, that the corresponding naïve rationalisation is true. But, first, I have argued that this solution is only available in a limited range of cases. And even if beliefs can be ascribed to more creatures than I have argued, most causal theorists will not be tempted to appeal to this solution when the example is, say, a caterpillar foraging. Yet the naïve form of rationalisation that I have outlined above does not appear to be limited in these ways.

The more fundamental issue, however, is that if we appeal to causation by the animal's mental states in order to explain how the parts of a complex action are unified into an instrumental unity, we face an unpalatable dilemma. The dilemma is simple: either each individual sub-action is intentional or it is not.

Suppose each sub-action *is* intentional. How are we to make sense of this idea, given the Standard Approach's insistence that this entails causation by beliefs, desires, and, perhaps, intentions? Presumably, the causality that is involved is not constituted by practical inferences prior to every sub-action. For unless one pushes these inference processes down to the sub-animal level, this idea cannot be taken seriously (when one writes the word 'action', one does not reason about how to write each letter); and if one *does* push them down, and is thus willing to construe their operation as something entirely beyond the animal's control (much like

¹⁴² The general point that 'and then' typically signals a different relation from 'and thereby' is usefully discussed in Small (2012, p. 150fn), Hanser (1998, p. 384) and Vogler (2002, pp. 126-146).

digestion), then they must also be willing, I think, to accept that the very idea of animal agency is an illusion.

An alternative suggestion that might seem to avoid these concerns is made by Kieran Setiya:

Suppose I am playing a particular piece of music. This is something I do *by* performing other actions, like playing a passage, or a movement of the piece, and finally performing *basic* actions, like moving my hand in order to play a note [...] What this comes to in the typical case is that I *intend* to be playing the piece, and that the further intentions that motivate the basic actions by which I have been doing so are motivated (in turn) by that original intention[.]¹⁴³

This idea seems to do away with the idea that the agent must engage in an absurd number of practical inferences; rather, she can just let the original intention take its course.

But is this really an improvement? As before, it requires postulating an unpalatable number of mental events. No one would recognise the idea that Setiya expresses in this passage as deliverance of naïve common sense in the human case.¹⁴⁴ And it is no more plausible when we consider its application to animal activities (think of the number of ‘further intentions’ that the theory would require in the case of, say, a lion hunting a gazelle). Since we are not already committed to their existence, we need some independent (i.e., not theoretically motivated) reason for believing in them.

Pursuing what someone like Setiya might say in defence of this idea would take us far afield. But, in any case, I think that most proponents of the Standard Approach would prefer to embrace the other horn of the dilemma. The idea would be that the animal’s beliefs, desires, and intentions initiate an action plan, or habit, that is then automatically carried out. On this view the activity’s parts are not caused by the belief and desire directly, but are rather mechanically elicited habits that are triggered by perceptual stimuli as the deed unfolds. (Thus, recall Clayton *et al.*’s suggestion that “The complex sequences of behaviour by which we drive, walk, or cycle to work may well appear to be purposive and goal-directed but, on further investigation, to consist of a chain of mechanistically elicited habits, albeit complex and highly structured ones, triggered by the stimuli along our route.”¹⁴⁵) Accordingly, although the completed event of catching crickets, say, is an intentional action (since it was initiated by a practical inference), the individual parts are not.

But let us consider this suggestion in more detail. The first thing to notice is that making this move renders the relation between the various sub-actions and the ultimate end *external*. That is

¹⁴³ (2003, pp. 363-364)

¹⁴⁴ McDowell (2011, p. 2)

¹⁴⁵ Clayton *et al.* (2006, pp. 199-200). Note that this quote comes from a paper discussing the complex caching strategies of scrub-jays, so, while it is ostensibly about human habits, I take the implication to be that the same lessons are applicable to the case at issue.

to say, while picking up a cricket may in fact bring the animal's end closer to realisation, it will do so in precisely the same sense in which dropping it on the other jay's head progresses it. Indeed, it would seem that we cannot so much as distinguish between what happens as a happy accident and what was, rather, undertaken as a means to an end. Worse: it would seem as if we cannot so much as *make sense* of the idea of a means. For *ex hypothesi* nothing that the animal does during the process leading up to the consummation of the process is undertaken for the sake of the end (except, perhaps, the very first step; but to think that the initial step undertaken by an animal bears a different relation to her ultimate end than any of the others is surely absurd).

To this it may be objected that, while the individual sub-actions are not, strictly speaking, intentional, they are nevertheless *guided* by the animal's mental states. That is, although they do not 'trigger' each sub-action, they are still operative in the sense that, if something goes wrong, they provide the animal with the means to steer her action back on course.¹⁴⁶

It is difficult to know how to respond to this proposal in the abstract. But one reason for being sceptical is that it would seem to model intentional actions *in general* on the case of a man who, sitting on a sleigh, pushes himself off, and *watches* as the process unfolds: if things go wrong and he is headed for a snow bank, he can intercede in the course of events; but he is otherwise placed in the role of an observer of a process whose progress towards completion is independent of his contribution. As Lavin remarks about the case of basic action:

It looks like performing a basic action is just being the subject of a mindless, automatic process which the subject has somehow initiated, triggered, or launched. [...] It seems the *I am doing A* of basic action is like the *I am going to the moon* of someone strapped to a rocket labeled 'to the moon' who has already (somehow) launched the rocket.¹⁴⁷

Treating the animal's role in the case of complex activities as this proposal recommends means, I think, that the problem that Lavin discerns in the case of basic action is simply writ large. Yet even if these cases (i.e. the sleigh rider and the rocket man) deserve to be counted as intentional actions, they are surely not paradigmatic exemplars of that notion. To base a theory of agency on the opposite assumption, is, to my mind, unacceptable.

No doubt there are further avenues of response that proponents of the Standard Approach will want to pursue. But I think that I have said enough, for now, to help myself to Anscombe's contention that "the assumption that some feature of the moment of acting constitutes actions as intentional leads us into inextricable confusions, and we must give it up."¹⁴⁸

¹⁴⁶ Cf. H. Frankfurt (1978)

¹⁴⁷ (2013, p. 293)

¹⁴⁸ (1963, p. 29)

3.5 Summary

In this chapter I have tried to provide some independent support for Anscombe's conclusion that being intentional is not some property that mere events might have or lack. I have argued for this claim indirectly, drawing on work which shows that CTA's appeal to desire as one of the *differentiae* of intentional action, far from explaining what an intentional action is, takes it for granted. This conclusion is weaker from Anscombe's own: I do not take myself to have shown that there is *no* additional feature in virtue of which mere happenings count as intentional actions, only that one of the most popular candidates is unsuitable for the task.

I then defended this conclusion against the supposition that its focus on explanation does not license a metaphysical conclusion (i.e., the conclusion that intentional actions are not 'composed' of a mere event and some additional fact about it). Though my defence on this score is — I admit — less than conclusive, I do think that it has at least shifted the burden of proof to those who want to defend such an account.

In the last section, then, I have tried to show where taking action for granted misleads us. My central claim was that the Standard Approach lacks the ability to account for the complex, instrumental structure of animal activities in a satisfactory way: either they posit an implausible number of mental events or they reduce the activity to a kind of mindless, automatic, process that places the agent in the role of an observer of events whose occurrence is by and large independent of her contribution.

Appendix: The argument of *Intention* §19

This chapter arose out of an attempt to understand Anscombe's cryptic argument for the striking conclusion that there is no 'additional feature' in virtue of which actions are intentional.¹⁴⁹ And while I have defended a weaker version of this claim in this chapter, I have not discussed Anscombe's own argument. For those interested, I have included my interpretation of it below, along with an explanation of why I do not believe that it is ultimately successful. (As what follows is ultimately irrelevant to the rest of my thesis, I thought it best to relegate it to an appendix.)

* * *

Suppose that we come upon a man — Smith — at his workbench. Smith is sawing a plank. As it happens, he is sawing an oak plank; he is moving his arms, contracting various muscles,

¹⁴⁹ (1963, p. 28)

making a pile of sawdust, and disturbing the neighbours. These are all things that he is *doing*, but only some of them are intentional actions. The question is: Is there some ‘additional feature’ that, in combination with Smith’s bodily movements in this situation, makes it the case that he is sawing a plank *intentionally*? In a brief and obscure section of *Intention*, Anscombe offers an argument that is intended to show that the answer is ‘No’, and, indeed, more generally that “We do not add anything attaching to the action at the time it is done by describing it as intentional.”¹⁵⁰ It is perhaps best to begin by simply stating it.

Assume that there *is* some such additional feature. Call it ‘*x*’.¹⁵¹ Insofar as *x* is going to fulfil its role, we know that it must have certain properties. Since we want it to account for the intuitive idea that acting intentionally is a form of *self*-movement, *x* must presumably be a feature of Smith. And since we know that Smith’s movements are intentional under some descriptions but not others, *x* must make it the case that some true description of the action *is* one under which the action is intentional. For example, if his action is intentional under the description (a) ‘sawing a plank’ but not (b) ‘disturbing the neighbours’ then *x* is the feature that explains why. Our hypothesis then becomes: ‘Smith’s bodily movements (in this context) + *x* = Smith’s action of sawing the plank intentionally’ but ‘Smith’s bodily movements (in the same context) + *x* ≠ Smith’s action of disturbing the neighbours intentionally’. But, Anscombe says, nothing about the man or his bodily movements, “can possibly determine the content of that description; which therefore may be *any* one, if we are merely considering what can be determined about the man by himself in the moment.” That is to say: each of the wider descriptions of Smith’s bodily movement’s — including (a) and (b) — will be underdetermined by *whatever* features of the man or his movements we may choose. But then *no* feature — no *x* — can provide a basis for saying that his action is intentional under one description but not the other. Thus Anscombe concludes that “an action is not called ‘intentional’ in virtue of any extra feature[.]”¹⁵²

Such is the core of Anscombe’s brief *reductio* of the compositional approach. Its interpretation is difficult, but it seems evident that Anscombe is not making the true, but benign point that we do not, as a matter of fact, typically *call* actions intentional in virtue of some ‘additional feature’

¹⁵⁰ (1963, p. 28)

¹⁵¹ I have replaced Anscombe’s ‘*T*’ with my preferred ‘*x*’ in order to mitigate unwarranted assumptions about what values the variable can take. For if her argument is as general as she claims it to be, the variable ranges over *every* possible ‘additional feature’ that one can think of.

¹⁵² Both quotes are from (1963, §19) My presentation of the argument departs from Anscombe’s text in two ways. First, I have replaced her reference to ‘pre-intentional muscle contractions’ with ‘bodily movements’. As far as I can tell, the argument runs in the same way and I prefer this way of phrasing because it connects it to the Standard Story more explicitly. Second, Anscombe introduces *x* as some additional feature present *at the time of acting*. This weakens the argument unnecessarily, making it impotent against most versions of the Standard Story which take the ‘additional feature’ to be a preceding cause. Thus Davidson says that “the intention is not part of the action, but a cause of it. Just as nothing is added to my telephoning my friend when that act becomes a thanking, so nothing is added to my arm going up if that event is caused by an intention.” (1987b, p. 105)

attaching to them at the time that they are done. After all, there need be no readily observable difference between my arm rising as a result of a thin piece of fishing wire yanking it up and me raising my arm, and yet we may have no trouble distinguishing them in light of the ‘wider circumstances and further consequences’ of the movement; on the other hand, even if there is an internal physiological difference, *this* is obviously not the criterion that we actually use in forming our judgments about what was done intentionally.

No, her claim is surely the stronger one that no such feature *could* be discovered since *every* such feature would fail to determine the description under which the action is intentional. But why? Consider a theorist who thinks that Smith is sawing a plank intentionally (but not disturbing the neighbours intentionally) in virtue of the fact that his *intention to saw it is causing him to do so*. They will surely hold that this intention is a feature of Smith which determines a description under which the action is intentional, namely (a). Such a simplistic theory may be incorrect, but not — on the face of it at least — because it fails to determine a description under which the action is intentional.

To understand Anscombe’s position, I think that we need to consider two different ways of cashing out this sort of view. On a naïve understanding of what it would be for an intention to cause an action, we take an intention to be the result of a conscious deliberation that prompts the agent to act. For example, Smith says to himself ‘I need to get a fence built, so I’ll go saw a plank’ — and then he does so *because* of this episode. Now, certain remarks aside,¹⁵³ I do not think that Anscombe’s considered view is that there is anything amiss in the simple idea that intentions of this kind *can* cause intentional actions. Her objection is rather that, if we take episodes of this kind to bear a constitutive relation to intentional actions, we face an unpalatable dilemma. For either this conscious process is itself brought about intentionally or it isn’t. If it isn’t, then it is unclear why its result should be considered intentional; but if it is, then we will be forced into a regress.¹⁵⁴

Quite independently of what we make of this argument, however, I think that the more obvious problem with the simple view is just that conscious episodes of this kind are clearly not necessary conditions of intentional action. I might roll out of bed, make coffee, and collect the newspaper, for example, all the while thinking about other things entirely (or, indeed, nothing at all). And Smith *need not* have engaged in an act of practical reason prior to acting. Indeed, he might be sawing a plank, according to Anscombe, ‘for no reason at all’.

¹⁵³ Cf. e.g. (1989, p. 110) and (1963, p. 24).

¹⁵⁴ At any rate, this is how I interpret Anscombe’s reasoning in the final few sentences of the second paragraph of *Intention* §19. See also Ryle (1949, p. 67ff), Geach (2000, pp. 73-74), and Hyman (2015, pp. 20-24).

So, even if an x of this kind *can* determine a description under which the action is intentional, it does not belong in an account of intentional action. As Anscombe says,

[...] it is one thing to say that a distinct and identifiable state of a human being, namely his having a certain intention, *may* cause various things to happen, even including the doing of what the intention was an intention to do; and quite another to say that *for* an action to be done in fulfilment of a certain intention (which existed *before* the action) is *eo ipso* for it to be caused by that prior intention.¹⁵⁵

This point evinces Wittgenstein's teachings and is all right as far as it goes.¹⁵⁶ But Anscombe is well-aware that most compositional theorists will disavow such views. Instead, they will maintain that in cases like these the agent's beliefs, desires, and intentions have simply failed to reach occurrent consciousness. They are nevertheless features of the person in the sense that there are states of the individual (or perhaps her brain) which correspond to (or are identical with) her attitudes. On this more sophisticated account of things, states of belief, desire, and intention do not need to be conscious in order to serve as the springs of action.¹⁵⁷

I take it that it is *this* sort of view that Anscombe has in mind when she complains that 'nothing about the man can determine the content of the description under which his action is intentional'. And although Anscombe says frustratingly little in *Intention* §19 to support this complaint, in the 'The Causation of Action' she is more forthcoming. Here she objects that such a position has the absurd consequence that a person might have an intention to saw a plank (e.g.) even though they have never heard of planks (or saws, or anything at all). For, on the assumption that a brain-state of an individual could serve as a sufficient condition for a given intention, we can imagine the same state produced in various artificial ways, in circumstances that do not involve the existence of the objects upon which the state is directed. She writes,

We take it that a state is supposed to be something holding of its subject here and now, or over a period of time, without reference to anything outside that of which it holds or the time which it holds: in particular, without reference to the history of the thing whose state it is. If that is how we understand a state, we can suppose the same state of an object in quite difference circumstances and with a completely different history.¹⁵⁸

This sort of argument is familiar from the literature on semantic externalism: since the meaning of a term depends crucially on the environment surrounding its use — my term 'water' does

¹⁵⁵ (1983, p. 101)

¹⁵⁶ But cf. Wittgenstein (1987, §67)

¹⁵⁷ For a particularly clear expression of these ideas, see T. Crane (2002, p. 69). These points are also made by D. Dennett (1968) in his criticism of Anscombe's argument.

¹⁵⁸ (1983, p. 99)

not mean the same thing as my twin-earth doppelganger's term 'water', etc. — “meanings ain't in the head.”¹⁵⁹ Anscombe's thought is an extension of this: intentions and other intentional states are no more in the head than are meanings.¹⁶⁰ And neither are they anywhere else in the individual. For “[...] we are implicitly looking *away* from the individual and into his world if we ascribe *any* [intentional state] to him.”¹⁶¹

Arguments of this kind aim to put pressure on the project of construing the truth conditions for intentional state ascriptions individualistically. They are of course controversial, and Anscombe's version is no less so. But even if we accept it, it is less than clear whether it supports her wholly general conclusion that actions are not intentional in virtue of any feature of the individual. For mightn't someone maintain that although the *content* of an intentional state is determined externally, the state itself is (i) identical to a feature of the agent (or her brain) and (ii) it is in virtue this feature's causing a certain bodily movement that the action thus performed is intentional?

The argument to show that this is *not* possible moves from: (1) The content — hence identity — of an intentional state about (e.g.) planks presupposes that the bearer of the state stands in certain relations to actual planks, and: (2) No feature of an individual presupposes any relation to planks whatsoever, to: (3) No feature of an individual can be identified with an intentional state about planks.¹⁶² But on one reading of (2), the conclusion does not follow, and on the other it is false. To see why, consider a parallel argument: (1') a gunshot wound presupposes a causal relation to a bullet; (2') no feature of an individual presupposes a causal relation to a bullet; *ergo*: (3') a gunshot wound cannot be identified with any feature of an individual. If read straightforwardly, (2') is false: were I to have the misfortune of suffering such a wound, one of my features — namely the wound — *would* presuppose a causal relation to a bullet. Alternatively, (2') is true if it means that it is *possible* to describe me and my features in ways that do not presuppose a causal relation to a bullet. But this reading will not support the general conclusion that my gunshot wound cannot be identified with these physical features, since one and the same feature may be described in various ways. Davidson put the point in this way:

¹⁵⁹ Putnam (1975, p. 227)

¹⁶⁰ R. Hursthouse (2000, p. 89) reads *Intention* §19 in light of Anscombe's argument in 'The Causation of Action' in a similar way. More recently, Vogler (2016) has offered a markedly different interpretation of the argument. On this reading, Anscombe anticipates the problem posed by devious causal chains. Although I find this implausible as a reading of *Intention* §19, I agree with Vogler that her interpretation is compatible with the spirit of Anscombe's conclusion (Vogler admits that the text does not directly support her strong reading).

¹⁶¹ (1983, p. 100)

¹⁶² Cf. Andrew Woodfield (1982, p. viii): “No *de re* state about an object that is external to the person's brain can possibly be identical with a state of that brain, since no brain state presupposes the existence of an external object.”

Individual states and events don't *conceptually* presuppose anything in themselves; some of their *descriptions* may, however. My paternal grandfather didn't presuppose me, but if someone can be described as my paternal grandfather, several people besides my grandfather, including me, must exist.¹⁶³

But now, returning to Anscombe's argument, the claim was that if we identified intentional states with features of the individual, then we would be forced into the absurd conclusion that an individual who had never heard of planks might have an intention about them. But this no more follows than the claim that an artificially created wound which was qualitatively indistinguishable from a gunshot wound would *be* a gunshot wound. For again: a gunshot wound is a feature of me which *can* be described without any presuppositions about its cause, but it is only describable *as a gunshot wound* given a certain aetiology. By parity of reasoning, it is open to a compositional theorist to maintain (i) that features of individuals can only be described as intentional states when they are suitably embedded in a wider context and (ii) it is in virtue of their causing some bodily movement of mine that the movement counts as an intentional action. Indeed, something not far from this would appear to be the position of Davidson himself.¹⁶⁴ And against this sort of view, Anscombe's argument gets no purchase.

¹⁶³ (1987a, p. 452)

¹⁶⁴ See, e.g., Davidson (1963, 1971; 1978).

Action and Structure

“As a person and animal are metaphysically distinctive forms of living thing, so, presumably also *intentional action* and *animal movement* are metaphysically forms of vital operation. The traditional hierarchy of forms of life appears to correspond to a hierarchy of forms of agency.”¹⁶⁵

4.0 Preliminaries

In this chapter and the next I develop an Anscombean alternative to the Standard Approach. As should come as no surprise, this alternative account *starts* with what happens when animals act, before working backwards towards the mind’s involvement in the process.

To this end I want to begin by describing some of the most striking differences between the activities of animals and other kinds of process found throughout nature.

* * *

In the introduction I suggested that just about any material object is an agent in the minimal sense that it is a potential source of change. Bolts of lightning, asteroids, and the sun are agents in this sense (as, indeed, every material object presumably is). With a deliberate air of paradox, I called this *inanimate* agency.

Inanimate agents typically play a correspondingly minimal role in the processes that they bring about. A forest may be burning down as the result of something that a lightning bolt ‘did’ (it struck a shrub), but the lightning itself plays no role in what happens from that point onward. If the sun is warming a stone, or a rock is crushing a snail, the respective agents do play a role

¹⁶⁵ Thompson (2008, p. 44)

in the processes taking place. But it is a limited one, and is the same across time. From beginning to end, it was, precisely, *warming the stone* that the sun was doing. Of course, warming the stone may in turn have further effects; but the sun does not have to do anything *else* to produce them.¹⁶⁶ Similarly with the stone: although it may be killing the snail *by* crushing it, the death is not the result of some further contribution on the part of the stone; the killing of the snail is simply the inevitable consequence of the crushing. I call processes of this kind *aetiological* since their temporal stages either perpetuate each other or are sustained by some external perpetuating force.

Consider now the kind of vital process characteristic of plants — the sprouting of a new shoot of bamboo, the phases of photosynthesis in a daylily, and the flowering of a rosebud in early summer, for example. In each of these cases the respective plants have a role to play in the development of the process in question that is, in some ways, unlike that of the stone or the sun. For although these processes are no doubt effected in part by external forces — the arrival of rain and sunlight, for example — the plants are ‘primed’, as it were, to take advantage of them. As Thompson remarks, in cases like these the arrival of light, or rain, “[...] are also phases of larger processes *in the plant* [...] The [plants] are subjects of processes of their own, which the meteorological or physical events merely complete or continue[.]”¹⁶⁷

What is interesting about these cases, I think, is that there is a minimal kind of *coordination* between what the plant does and what its environment affords it. It is no accident that the warming temperatures of summer correspond to the blossoming of the rosebud, in the way in which there *would* be something accidental, a mere turn of fate, if the rosebud were to be subjected to the withering effects of a child’s magnifying glass. The latter sort of event serves to *interrupt* the processes that were already underway within the plant. So, although the withering of the rosebud and its flowering in the summer are both the effect of sunlight, in the latter sort of case the rose is no mere patient: it is, rather, the locus of processes that depend for their successful completion on the arrival of the light; the two events are united within an aetiological nexus not merely as cause and effect, but, also, as part and whole.

Still, and as everyone knows, this coordination is effected by processes that lie beyond the plant’s control. They are processes which are triggered by the environment in conjunction with mechanisms that belong to every healthy member of whichever species of rose the bud belongs

¹⁶⁶ With respect to the process in question, that is. The sun was of course participating in an indefinite number of other processes (warming the ocean, e.g.) at the same time.

¹⁶⁷ (2008, p. 41)

to. Accordingly, the coordination at issue is that which is provided for by evolutionary history; it is not something that the plant is itself responsible for bringing about.

Animals can play the role of inanimate agents and, indeed, they can initiate and sustain vital processes. A lion may dislodge a stone and thereby trigger a landslide, for instance, whilst a horse may shiver in response to the cold. But animals also engage in *activities*: lions and dolphins play, snakes and spiders hunt, pigs and caterpillars forage, and so on and so forth. As I pointed out in the preceding chapter, doing any of these things requires that the animal do a number of further things as well. A lioness that is in the midst of her hunt, for instance, must move her legs, one after the other, if she is to walk across the savannah; but putting one leg in front of the other is not a *cause* of the next movement, nor does the latter follow thereby in some other way. Rather, each step requires a separate contribution on the part of the animal. Unlike a bolt of lightning whose role as agent is over as soon as it starts the forest fire, the lioness must continuously *act* until the activity concludes. And unlike the rosebud, the lioness must actively coordinate its movements with the contingencies of its environment: most obviously as she adjusts her movements in response to her prey, but also more subtly as she takes advantage of things like tall grass, watering holes, and other features of the locale in which her activity takes place.

I call processes of this kind *agential* because I believe that commonplaces such as these serve to motivate the naïve thought that the agency of animals is different in kind from that found throughout the rest of nature.^{168, 169} For although I have been happy to speak of lightning bolts and plants as sources of change, notice that their role in the explanation of what actually happens can typically be replaced by talk of further events and processes: we can say that the lightning bolt caused the fire to start, if we like, but we could with equal right say that some *event* — such as the lightning’s striking of the shrub — did so. Similarly, we can eliminate the appeal to rose by instead speaking of the processes of flowering, photosynthesis, and whatever else goes into the production of a rosebud. Indeed, the second sort of explanation is arguably more

¹⁶⁸ For extensive discussion of the empirical literature documenting the presence of these naïve judgments in early childhood development, see Steward (2012; chapter 4).

¹⁶⁹ Note that the point here is not to deny that an animal might exercise her agency when she does just one thing — when she raises her paw, say (though, in fact, it is far from obvious whether she really has just done one thing here. See Lavin (2013) and Thompson (2008, pp. 113-119) for arguments that bodily movements of this kind are complex). My point is just that *if* this counts as an agency, it will because the animal is *capable* of putting that sort of movement to further use in the context of broader patterns of activity. My claim is that the capacity to move one’s body in accordance with one’s will (to adopt the broadest term available) is *conceptually* — though obviously not causally — dependent on the capacity to do things like hunt, play, and swim.

perspicuous in both cases.¹⁷⁰ And if this is right, then there is, strictly speaking, no role for these agents *qua* agents to play in the explanation of what happens.

Of course, the Standard Approach would have it that the same is true of *all* agents. According to Davidson, for example, “although we say that the agent caused the death of the victim, that is, that he killed him, this is an elliptical way of saying that some act of the agent — something he did, such as put poison in the grapefruit — caused the death of the victim.”¹⁷¹ But reflection on the part-whole structure of activity suggests that we cannot, without loss, eliminate the agent in the explanation of agential processes. For, unlike the case that Davidson describes, there is no single event or process to which the wider description harkens back. And while it is true that activities are realised through a series of sub-actions, I have argued that the Standard Approach lacks the resources to unify these parts into an integrated whole (see §3.4).

At a minimum, then, this suggests that the role of the animal in these cases is not comparable to that of a mere trigger or sustaining force. More boldly, we might repurpose another Davidsonian theme and argue that the various parts of an activity ‘triangulate’ an agent at its centre, thus locating a common source of motion and change and thereby ensuring a role for agency proper. But describing this role in a way that leaves the basic insight intact is difficult. Certainly, the answers that immediately suggest themselves give rise to misleading comparisons. For instance, it is tempting to describe the agent as a ‘guide’ and then to analogise her role in the activity to a puppet-master directing her play. But this is no better, I think, than the old metaphor of the captain directing his ship. For once we have separated the animal from the process that she is supposed to guide, we are led to contemplate the mysterious way in which they are conjoined — a dead end.

The grain of truth in these metaphors, I think, is that we need to locate the animal *in the midst of* her activity: Just as the strings that animate the puppets trace back to a common origin, so, too, do the parts of an ongoing activity revolve around a locus of motion and change. And I do believe that this provides *some* justification, however obscure, for the naïve thought that we must not collapse the agency of animals into that of rocks and roses. But, again, the metaphor serves at best to isolate the problem and in no way constitutes a solution to it.

It is here, I think, that Anscombe has much to teach us. But in order to appreciate her lessons, we first need to reflect on the temporal nature of events and processes more generally.

¹⁷⁰ On this point see and cf. Davidson (1967) and Galton (2012). For defense of substance causation in cases like these, see Steward (2012; chapter 8).

¹⁷¹ Davidson (1971, p. 49). See also Velleman (1992). Cf. Alvarez and Hyman (1998)

4.1 The temporal unity of events

Let me begin by marshalling some familiar facts about the way that events and processes are represented. Compare the following sets of sentences:

- | | |
|---|--|
| (1a) ‘The sun was warming the stone’ | (1b) ‘The sun warmed the stone’ |
| (2a) ‘The forest was burning down’ | (2b) ‘The forest burnt down’ |
| (3a) ‘The gazelle was walking to the river’ | (3b) ‘The gazelle walked to the river’ |
| (4a) ‘The lion was playing with its cubs’ | (4b) ‘The lion played with its cubs’ |

In each case, the corresponding pairs are formed from the same materials — they each share a subject, a tense, and a predicate — but they differ in aspect: the predicates on the left have taken on a *progressive* aspectual determination while those on the right have taken on a correspondingly *perfective* determination.

This difference of aspect relates to two different perspectives on events. Put simply, the perfective is used to describe completed events, whereas the progressive is used to represent events-in-progress or *processes*.¹⁷² Going into slightly more detail, we might say that the use of the progressive aspect affords us a perspective that is, so to speak, sensitive to the temporal dynamics of the situation in question. To paraphrase Bernard Comrie, the progressive represents a point of view from which we can look back to the beginning of the process’ instantiation and anticipate future change.¹⁷³ By contrast, perfective aspect suggests that the event is something that is over, done with, or otherwise finished. As he summarises, the use of the perfective abstracts away from an event’s temporal structure, instead presenting the

¹⁷² Some remarks about terminology are perhaps appropriate here. I understand events and processes as formal correlates of one another, respectively answering the questions ‘What happened?’ and ‘What is happening?’ An event-or-process form, as expressed by predicates that admit of the progressive aspect, corresponds to the general shape that a particular happening may take; it does not designate a particular occurrence. For example, a predicate such as *walk-to-the-river* can be used to represent an event, in progress or otherwise complete, that might be instantiated any number of times, by any number of subjects, and in any number of locations. As Rödl remarks, “One describes the form of an event in saying what someone did — eating a bowl of ice cream, say. One may further say who did it — Jochen, for example — when he did it — last Monday — where he did it — at the Café San Remo — and how long it took him to do it — two minutes. In saying, Who?, When?, Where?, How Long?, one does not describe an event as regards its form. One describes it as a concrete event. An event-form does not define the subject, the date, the place, and the duration of events which fall under it.” (2002, pp. 326-327) (See also Hornsby (2016, p. 3).)

Technically, then, the sentences that we are interested in (those of the form ‘S is/was doing A’ and ‘S did A’) describe particular process *instantiations* and particular event *occurrences*. Since it would complicate the text unnecessarily to include the ‘particular’ qualifier in front of every use of ‘event’ and ‘process’, however, I only include it when context does not provide the necessary disambiguation.

¹⁷³ See Comrie (1976, p. 4).

situation “as a single unanalysable whole, with beginning, middle, and end rolled into one.”¹⁷⁴ It is for this reason that the sentences on the left, but not on the right, can be cast in the present tense.¹⁷⁵

Now, if we focus our attention on the first and fourth pair, the difference between the perfective and progressive representation of events might be thought to be *merely* perspectival. After all, the truth of (1b) and (4b) implies the corresponding truth of (1a) and (4a) and *vice versa*: If the lion played with her cubs, then she was playing with them over some (perhaps short) temporal interval, and, conversely, if she was playing with them, then she has played with them. Indeed, if we suppose that she was playing with them from t_1 – t_n then it follows that she played with them over every temporal interval in-between; as Zeno Vendler comments, “any part of the process is of the same nature as the whole.”¹⁷⁶

But the other examples demonstrate that this is not in general true. For although it follows from the fact that the gazelle walked to the river that there was a period of time during which she was walking there, the converse implication does not hold. This phenomenon is known as the ‘openness’ of progressive judgments and it records the fact that just because something *is happening* it does not necessarily follow that it *will happen* — some events are such that they may be interrupted, frustrated, or, in the case of voluntary activities, given up. Thus (3a) the gazelle may be *walking* to the river, but — due to an unforeseen lion attack — ($\neg$ 3b) never end up having *walked* there. And similarly, it might be true that (2a) the forest was *burning* down, but also true that the fire brigade stopped it (i.e., that ($\neg$ 2b) it never *burnt* down).

Now as I just noted, (1a) and (4a) are not open in this sense. To understand the difference, we may start with the observation that the invalidity of the inference from both (2a) to (2b) and (3a) to (3b) is underwritten by the fact that the predicates *burn-down* and *walk-to-the-river* each specify a natural terminus: a time before which the process is incomplete and after which it is finished. By contrast, consider the predicates formed by taking away this implied terminus, *walk* and *burn*. In this case the fact that such processes are underway *does* imply that they have taken place: if the forest was burning, then it has burned; if a gazelle is walking, then it has walked. As these examples suggest, we can distinguish between these two types of predicate according to

¹⁷⁴ (1976, p. 3)

¹⁷⁵ I touched on this point in §5.4. Just as ‘I turn off the lights’ does not report, of the present, what ‘I turned off the lights’ says of the past, so, too, ‘The gazelle walks to the river’ does not say of the present what (3b) says of the past. In both cases, the combination of the present tense with perfective aspect supports a habitual reading. For discussion see Comrie (1976, p. 26ff), Galton (1984, pp. 11-12) and Thompson (2008, p. 125).

¹⁷⁶ (1957, p. 101)

the validity of the inference from the present progressive to the past perfect: If the inference is valid, then we are dealing with an ‘atelic’ predicate, and if it isn’t, the predicate is ‘telic’.

Before moving on, however, let me emphasise that grammar is not an infallible guide to the nature of the process being described. Indeed, this is readily apparent in respect of many of the predicates used to describe animal activities. Thus consider the predicate *hunt* (or even *hunt-wildebeest*). By our criterion this is an atelic predicate. However, the *activity* of hunting surely does have a natural endpoint, namely, the killing of prey. And, moreover, a specific occurrence of this sort of activity *can* be frustrated, interrupted, or given up: the lioness may be struck by a poacher’s bullet as she is walking across the savannah, for instance, or her prey may prove to be too fast for her on this occasion; either way, the hunt ends, but does not *complete*. In this sort of case I shall say that although the *predicate* is atelic, the represented *process* is telic.¹⁷⁷

Terminological points aside, it is the notion of completion-yet-unattained, inherent in the progressive representation of many telic processes, that gives rise to the puzzle that I mean to discuss. The puzzle might be put like this: ‘If it is an open question whether an event-in-progress will complete, how can it be true to say that it is precisely a walk *to the river* that is underway? And for that matter, why say that the forest is burning *down*? Would it not be more accurate to leave the future out entirely and instead restrict ourselves to the atelic representations ‘The gazelle is walking in the direction of the river’ and ‘The forest is burning’? Of course, once the gazelle has walked to the river and the forest has burnt down — once the events are complete — we may form the corresponding perfective judgments. But until they are it is unclear why we should prefer (2a) and (3a) over, say, the thought that the gazelle is walking *to its death* or that the forest is burning *half-way down* (or, indeed, any number of other possible futures).¹⁷⁸

This sceptical doubt is not intended to be epistemological. The worry is rather that the possibility of incompleteness threatens to undermine the intelligibility of telic predication by making it indeterminate whether one or another process is underway. I do not mean to endorse the scepticism, of course. On the contrary, the sceptical doubt is surely absurd. For if events such as sunsets, walks to the river, and trips to the shop belong to the history of the world — as

¹⁷⁷ The practice of applying the telic/atelic contrast to particular event and process occurrences is sometimes objected to on the grounds that it entails that every telic occurrence is *also* atelic. Thus Krifka (1998, p. 207) says that “it is misleading to think that a particular event can be called ‘telic’ or ‘atelic’. For example, one and the same event of running can be described by running (i.e. by an atelic predicate, or by running a mile (i.e. atelic, or delimited, predicate). Hence the distinction between telicity and atelicity should not be one in the nature of the object described, but in the description applied to the object.” But the argument is weak. For Krifka does not explain why we cannot instead say that telic processes may be truly — though less than completely — described atelically.

¹⁷⁸ Cf. Small (2012, pp. 178-180) and Galton (1984, pp. 137-143)

they surely do — then there must have been times at which they were taking place. Indeed, the thought that we might help ourselves to the notion of such events having happened without the corresponding notion of their being in progress is borderline unintelligible.¹⁷⁹ The notions *complete event* and *event-in-progress* are co-relative, internally related concepts.¹⁸⁰ Still, the question remains: how *do* we justify the application of judgements such as (2a) and (3a)?

The way around the puzzle starts with the recognition that, although completion is not guaranteed, an event-in-progress is not simply indifferent to its own future. An indication of this lies in the intuitive idea that a forest which is burning down *will* burn down, *unless* the fire brigade (e.g.) intervenes. And similarly: the gazelle that *was* walking to the river presumably *would have* made it were it not attacked by lions. In both cases it would seem as if the unrealised state of completion is ‘coiled’, as Michael Thompson puts it, in the developing process itself.¹⁸¹

If we unpack the metaphor, then what these examples suggest, I think, is that the relationship between an event-in-progress and its future completion is such that an event-in-progress will complete *provided that nothing external interferes*. But we have already observed two other ways in which an activity may terminate prematurely: (i) by being *frustrated* (by prey that is too quick e.g.) or (ii) simply *given up* (as a result of a change of mind). It would be odd to think of activities that terminate in these ways as being interfered with by something external. But the important point to note is just that each possibility seems to presuppose that the event *was* heading towards completion *prior* to being interrupted, frustrated, or otherwise given up. This does not solve our problem, of course, but it allows us to put a finer point on it. For now the crucial question becomes: In virtue of what principles do events-in-progress tend toward completion? Call this the problem of the temporal unity of events.¹⁸²

4.2 Process and law

In the case of aetiological processes, I think that the answer to this question is relatively straightforward. For given that an aetiological process is one whose successive phases stand in direct causal relations to those that precede them, the principles which explain why processes of this kind tend towards completion are the relevant causal laws. As Anthony Galton remarks,

¹⁷⁹ It is true that instantaneous events (‘achievements’), such as winning or starting to act might be exceptions to the rule that what has happened was at one point happening (though cf. Z. Szabo (2004)). I discuss these below (§6.6). For now, however, we may restrict our attention to temporally extended events. For them the point stands.

¹⁸⁰ This point is emphasised by Small (2012, pp. 179-180).

¹⁸¹ See Thompson (2008, pp. 130-138)

¹⁸² Cf. Small (2012, p. 183)

any particular stage in the burning of a house is how it is because of the earlier stages. The extent and intensity of the conflagration at any time depend on the past history of the fire, and in turn help to determine its future history. For this reason, if at any time the event comes to a halt in mid-course (in this case, the house stops burning), then it cannot be resumed without a specific separate act of causation from outside.¹⁸³

In this case the laws are those of thermodynamics. But the tendency-towards-completion of a biological process, such as mitosis, can also be explained by the laws of the discipline, and so on for other types of aetiological processes. And if this is right, then we have an answer as to what justifies their progressive representation: it is the (implicitly assumed) fact that each stage follows from the next according to law. Small summarises the general idea as follows:

On this picture, [...] we distinguish between a case in which *S* is *actually* doing *B* from one in which *S* is only *apparently* doing *B* but is in fact doing *B'* by apprehending what is happening as exemplifying a law that is part of a system of laws (e.g. chemical laws) that together determine a generic kind of process (e.g. *chemical process*); the law that what *S* is doing exemplifies and through which it is understood specifies what doing *B* is and explains why *S* is going to do *B* [i.e. why *S* will do *B* if nothing interferes].¹⁸⁴

This much seems plausible. The idea is just that the presumption of causal law provides the basis for the predictive aspect of progressive representation. The moot question, however, is whether similar explanatory principles are available in the case of the activities characteristic of animal life.

To some it will seem as though there is no significant difference between the cases. It will readily be conceded, of course, that the *observable* parts of a lion's hunt (e.g.) do not stand in direct causal relations to each another — one leg movement is not the cause of the proceeding one, as everyone must admit. But this difference does not, it will be suggested, run deep. For if we knew everything there was to know about events going on within the animal's nervous system, its standing bodily states, reinforcement history, and so on, then we would see that the successive phases of its activity do follow one another in a law-like fashion: "the animal was structured thus, these processes were occurring in those parts of its central nervous system, and all this led by pure chemistry to the [series of] movement[s] in question"¹⁸⁵, as we might imagine a proponent of this sort of view saying.

¹⁸³ (1984, p. 140)

¹⁸⁴ (2012, p. 187)

¹⁸⁵ J. Bennett (1976, p. 73); Bennett is not espousing the view in question.

This is not the place to engage with this form of explanatory reductionism in any great detail. However, I do think that there is a very general reason to be sceptical of its prospects. For even if there were such laws, they would presumably relate the phases of the event under descriptions suited to enter into the supposed laws (be these physical, chemical, neurological, or whatever). Suppose, then, that we get ourselves in position to explain, in general, the mechanics of a lion's hunt; we can explain how feedback from its limb movements effects a gamut of neurological changes, which then co-ordinate with the 'information' processed by its perceptual systems, to affect further limb movements, and so on and so forth, in such a way that tends to result in the ingestion of protein. Have we explained the temporal unity of the lion's *hunt*? I suggest not. For what gets explained on the mechanical account is not the tendency of a hunting animal to catch and kill prey, but rather the tendency of various movements, neurological events, and so on, to result in a transfer of energy; in short, not the sort of thing that *animals* (as opposed to their parts) do and certainly not what we were trying to account for to begin with.¹⁸⁶

In light of this, we might instead look to the sort of teleological law developed by (amongst others) Jonathan Bennett.¹⁸⁷ In essence, Bennett's idea is that we can explain why animals act as they do by appeal to laws of the form:

$$(x)(f)((Rx \ \& \ x \text{ registers that } f/Gx) \rightarrow fx)$$

This says that whenever something is in a certain state, *R*, it will do whatever it 'registers' ($\approx$ perceives) to be necessary and sufficient for achieving some goal, *G*. For instance, if an animal is in a state of hunger and registers that the only way to achieve its goal of eating eggs is to climb, it will climb. Now the formula no doubt needs to be developed to account for the possibility of conflicting preferences, the concept of registration needs to be explained, and other such niceties need to be filled in before any laws of this form are going to so much as get off the ground (as Bennett is well aware). But leaving these details aside, I think that we can see that such laws cannot help us when it comes to explaining the temporal unity of an activity in progress.

The reason, in brief, is that the aspectual considerations that we have been stressing up to now are simply papered over by Bennett's formula. To see this, consider the case of the climbing

¹⁸⁶ This response is, I admit, dogmatic. In order to deal with the proposal adequately, I would need to argue that there is no hope of finding bridging principles that link the two forms of description. Here I can only signal my broad agreement with others that have done so. In this regard see especially R. McDonough (1989) and D. Davidson (1970). For the importance of whole-animal predicability, see T. Burge (2009) and Bennett and Hacker (2003).

¹⁸⁷ See Bennett (1976; chapter 2). Cf. C. Taylor (1964).

animal in a bit more detail. The example is Bennett's own and, like most of his others, he represents the means to achieving the goal by means of an *atelic* predicate: *climb*. This means that it doesn't much matter whether we conceive of the corresponding consequent of the law as the *process* of climbing (with imperfective aspect) or as a completed *event* of having climbed (with perfective aspect); after all, once something starts climbing, it has climbed, and, conversely, if it has climbed, then there was a time when it began (and was) climbing. But how can it be, one wants to ask, that *simply* climbing could be sufficient for achieving anything, let alone eating eggs? The animal won't be able to eat if it keeps slipping off the tree half-way up!

Insofar as the means are going to be genuinely sufficient, they need to be represented by a telic predicate that has taken a perfective determination; in this case, something like *has-climbed-to-where-the-eggs-are-located*. But if so, then what, exactly, is supposed to follow when the antecedent conditions of the law hold? Is it that the animal *starts* climbing to where the eggs are located? Or is the idea that it *will* climb to where they are located? Neither option is going to help us. For if the former, then the law implies nothing whatsoever about the tendency towards completion of the process in question. But if the latter, then we have no reason to believe that laws of this form are ever true: telic processes are always subject to interruption, as we have seen. To save it, we would have to add the clause: provided nothing interferes. And now the law simply *assumes* the relevant tendency without explaining it.¹⁸⁸

4.3 Anscombe on practical knowledge

In *Intention* Anscombe provides the resources for a different approach to this problem. Rather than locating the principles that explain the temporal unity of intentional actions in antecedent mental processes or teleological laws, her basic insight is that the instrumental structure of the action can explain its own tendency towards realisation. No doubt this will sound mysterious at first, but the idea actually turns on a feature of our everyday thought about action; indeed, it is one that we have already touched upon: the idea that what someone is doing can serve as the reason for which another is done.

We can see this thought at work in Anscombe's notorious discussion of the man pumping water into a cistern. The man is (A) moving his arms; he is (B) operating a pump; he is (C)

¹⁸⁸ It is true, of course, that similar saving clauses are also found amongst the causal laws governing aetiological processes (with the exception of the pure nomological laws of physics). But there is a crucial difference: because each part is the cause of the next in such cases, we have an account of the tendency that the notion of interference presupposes. Since agential processes differ in precisely this respect, we are left without the requisite account.

replenishing the water-supply, and he is (D) poisoning the inhabitants. What's more, he is doing each for the sake of the next and the whole lot for the ultimate end of realising the kingdom of heaven on earth. It is this A-D order of ends that is revealed by Anscombe's question 'Why?' Thus:

'Why are you moving your arm up and down?' — 'I'm pumping.' 'Why are you pumping?' — 'I'm pumping the water-supply for the house.' [...] 'Why are you pumping the water?' — 'Because it's needed up at the house' and (*sotto voce*) 'To polish that lot off.' 'Why are you poisoning these people?' — 'If we can get rid of them, the other lot will get in and...'¹⁸⁹

Of course, these questions might have received more sophisticated answers — answers that contain psychological verbs (e.g. 'I am operating the pump because I *want* to replenish the water supply') or 'final form' explanations that utilise the 'in order to' construction (e.g. 'I am replenishing the water supply in order to poison the inhabitants'). Indeed, it is characteristic of the Standard Approach to ignore the naïve pattern of explanation and to instead focus solely on the more sophisticated, principally psychological forms. But I am proposing here to follow Anscombe and her contemporary followers in taking naïve rationalisations seriously. The question must then be: What must an intentional action be like in order to explain its own tendency towards completion?

Anscombe's answer to this question, in brief, is this: an intentional action must be an object of *practical knowledge* — a kind of knowledge that, in contrast to its speculative counterpart, is productive of what is known. With a nod to Aquinas, Anscombe describes practical knowledge as 'the cause of what it understands'.¹⁹⁰

Anscombe says that this idea is 'difficult', and is so in part because "modern philosophy has an incorrigibly contemplative conception of knowledge." According to this conception, "The facts, reality, are prior, and dictate what is to be said, if it is knowledge."¹⁹¹ Her idea seems to be that in the case of practical knowledge the order of priority is reversed: the facts are *posterior* to the agent's knowledge of what they are doing and depend upon this knowledge for their own being. But can we make any sense of this idea? And if we can, can this understanding be applied to animals?

¹⁸⁹ (1963, p. 38)

¹⁹⁰ Anscombe (1963, p. 87). I should like to emphasise here that the topic of practical knowledge (in Anscombe's sense) has received a considerable amount of treatment recently and there is considerable disagreement over how Anscombe's notion is best understood. I cannot wade into exegetical questions here, however; instead, I shall focus on presenting the reading of Anscombe that I believe to be most philosophically interesting in what follows. My hope is that it is deserving of attention in its own right.

¹⁹¹ Anscombe (1963, p. 57)

Although this second question will no doubt weigh heavily on the reader's mind in what follows, I raise it here chiefly to set it to one side. In the remainder of this chapter my concern will be to bring Anscombe's theory into view as it might find application in the case of self-consciously rational human beings. In the subsequent final chapter, I take up the question of the theory's application to animals.

With that caveat in mind, I think that the best place to start in trying to understand the basic idea behind the notion of practical knowledge is with Anscombe's reference to Aquinas.¹⁹² For Aquinas, practical knowledge is used to characterise God's knowledge of creation. A student of nature must learn about the natural world by apprehending facts that exist independently of her. Her case is precisely one for which the facts are prior to her knowledge of them. If she is of a sound mind, her beliefs will *derive* from the facts, but they do not cause them. God's knowledge of the world is not like this. *He* knows that things are thus and so because his intention suffices for their *being* thus and so. According to Aquinas, His power to determine what is the case by willing it affords him practical knowledge of what is thereby determined.¹⁹³

This comparison with the divine is both instructive and problematic. It is instructive in so far as it suggests a way in which intention might serve to underwrite a distinctly practical form of knowledge — a form which, unlike its speculative analogue, is capable of bringing into being the very objects that it is said to understand. And if we do possess such knowledge, then it would seem to explain the tendency-towards-completion of the activities upon which this knowledge is directed. It is problematic, however, for the obvious reason that we are material beings whose practical powers are limited by the constraints of the natural world: for creatures like us, merely willing something to be the case is insufficient for making it so.

I do not think that admitting this point forces us to abandon the comparison entirely, however. On the contrary, it is just a useful reminder that executing our intentions means getting our hands dirty. Thus, although I cannot move the matchbox in front of me *merely* by willing it, I can move it nevertheless. Within certain limits I have the power to 'determine' where it ends up — not on the moon, perhaps, but in the cupboard, certainly. *Pace* Davidson, it is not 'up to nature' whether the matchbox gets put away;¹⁹⁴ it is up to me and my kind. Anscombe's

¹⁹² (1963, p. 87)

¹⁹³ This is a simplification of Aquinas's position, but not, I think, a falsification of it. For more detailed discussion of Aquinas's conception of practical knowledge and its relation to Anscombe's, see J. Schwenkler (2015).

¹⁹⁴ The reference is of course to Davidson (1971).

thought, I think, is that practical knowledge serves to explain this ability to alter the world in accordance with our will.¹⁹⁵

There are at least two aspects to this idea. The first (and perhaps the most recognisable from the text of *Intention*) is that ‘moving the matchbox’ can be used to describe something that I am doing *intentionally* — i.e., as a result of my will — only in so far as this happening is cognised by me in a distinctly practical way. My knowledge must be underwritten, in other words, not by a speculative *belief* (or perceptual *evidence*) that I am moving the matchbox, but rather by an intention to do that very thing.¹⁹⁶ According to Anscombe, this explains the difference between someone who notices that they are moving the matchbox as a result of something else they are doing (it is a side-effect of lifting the table, say) and someone who, upon noticing the matchbox sitting on the table, decides to put it away. “Without [practical knowledge]”, she says, “what happens does not come under the description — execution of intentions — whose characteristics we have been investigating.”¹⁹⁷

Understood in this way, Anscombe’s idea is that practical knowledge is the *formal* cause of what it understands: Just as the bronze out of which a statue is made lacks meaningful significance when considered independently from its structuring principle, so, too, do the material movements of a human being lack the unity of an intentional activity when considered independently from the agent’s practical knowledge of what they are doing.

To appreciate the motivation behind this idea, it is worth reminding ourselves of the commonplace that not everything that is done throughout an ongoing activity is done intentionally. My accidentally moving the matchbox illustrates this: I am lifting the table and moving the matchbox and casting a shadow and loosening the rug and disturbing the dog and... etc. Although any number of these things *might* be done intentionally in other contexts, in this case most of them are side-effects brought about as the result of my action proper. Their relation to my wider project is coincidental, so to speak; even if they are foreseen (or observed), they are not intended.¹⁹⁸

¹⁹⁵ In this context ‘explain’ does not mean ‘explain in terms of concepts that are somehow better understood or more fundamental’. The concepts of *practical knowledge* and *intentional action* are, for Anscombe, coeval. The sense of explanation at issue is therefore better understood as ‘connective’, in Strawson’s (1992, p. 19) sense.

¹⁹⁶ For discussion of the parallel between knowledge ‘in intention’ and knowledge ‘in belief’, see L. Campbell (2017).

¹⁹⁷ (1963, p. 88)

¹⁹⁸ This is not to say that they are not relevant to the moral evaluation or legality of the act. Just that the jurist’s notion of ‘oblique intention’ should not be confused with the real thing. On this point see, e.g., G. Williams (1987).

But amidst this chaotic welter of happenings, there is a rational order to be found. Quite generally, in fact, the things done intentionally over the course of an activity hang together in an instrumental structure of means and ends. This is one of the lessons brought out in the case of the pumping man, but it can also be seen in the example just mentioned: I am lifting the table *in order to* loosen the rug (let us suppose); the former is the means to the latter end. These cases are relatively simple. But there is no predetermined limit to the internal complexity of an activity: I am building a shed. At the present moment, I am holding a hinge in place and screwing it into the wall. I am doing this because I am hanging the door. In a few moments, I am going to lift the door into place and screw it onto the hinge. After that, I will fetch a ladder and set to work on the roof, etc. etc.

Notice that ‘Hanging a door’, ‘Fetching a ladder’, ‘Holding a hinge in place’ and ‘Hammering a nail’ (but *not* ‘Making a hammering noise’ or ‘Making a mess’) are descriptions of *how* I am building the shed. As Anscombe observes, the A-D order revealed in the case of the man pumping water can also be elicited in reverse by means of the question ‘How?’ And when looked at this way, we have not an order of ends, but an order of means — an order of practical reasoning.

It is here that the notion of practical knowledge becomes apposite. Anscombe’s idea is that this order of practical reasoning is held in place, so to speak, by the agent’s knowledge ‘in intention’ of what they are doing, why they are doing it, and how it is being done. Regardless of whether the agent has actually engaged in an explicit process of thinking through the premises of a practical syllogism prior to acting, it is this knowledge that, when elicited by means of the questions ‘Why?’ and ‘How?’, locates the action and its parts amidst the chaotic flux of cause and effect. As she comments in a crucial passage:

[...] if [practical syllogisms] were supposed to describe actual mental processes, it would in general be quite absurd. The interest of the account is that it describes an order which is there whenever actions are done with intentions; the same order as I arrived at in discussing what ‘the intentional action’ was, when the man was pumping water.

[...] [practical syllogisms] or my order of questions ‘Why?’ can be looked at as a device which reveals the order that there is in this chaos.¹⁹⁹

On my reading, the ‘order that there is in the chaos’ refers to the order of means and ends that *is* an intentional activity. This is a spatiotemporally extended material process. But it is also an

¹⁹⁹ (1963, p. 80) In both of these quotes Anscombe refers to Aristotle’s account of practical reasoning. I have replaced this with ‘practical syllogisms’ as I do not propose to address Aristotle’s account here.

essentially self-knowing one: its instrumental structure is constituted by the agent's practical knowledge of that very structure.

* * *

I think that this conception of intentional action as the formal object of practical knowledge is philosophically attractive in many ways. It reminds us that intentional activities are possessed of complex instrumental structure and it offers us an explanation of what constitutes that structure. Moreover, it does so by giving primacy to the agent's own conception of what they are doing: it is to her and her self-consciously practical apprehension of what is happening that the parts trace back. This makes the account especially well-suited for understanding the ethical and legal aspects of agency (a topic that thankfully sits beyond the scope of this thesis).

But it also raises several questions. First, and as I have already observed, the analogy with the divine is imperfect: we cannot bring about the change that we wish to see in the world simply by intending it. And even if one grants (as they should) that this doesn't mean that we are incapable of altering the world in accordance with our will *simpliciter*, one might nevertheless wonder whether this undermines the idea that intention could serve to underwrite *knowledge* of what happens when we act: since we can never be sure that we are doing what we intend, we cannot know it either.

The second concern is that the train of thought is beginning to lose touch with the problem of temporal unity that was supposed to motivate it. For what goes missing from the account of practical knowledge as a formal cause is, it might be said, its *material influence*.²⁰⁰ Thus even if it is granted to Anscombe that what happens in the material world would not come under the *description* 'execution of intentions' if not for the agent's practical knowledge, the knowledge itself does not seem to be playing any role in explaining the progress-towards-completion of the material process itself. Indeed, it seems to take it for granted: after all, in order to have practical knowledge that one is doing A, A must be happening (knowledge is factive). But practical knowledge was supposed to explain why it is precisely A-ing that *is* happening in the first place.

Of these issues, Anscombe engages most explicitly with the first. But she does make it clear, I think, that her conception of practical knowledge is richer than that of a mere formal cause. For although she does place emphasis on the idea that practical knowledge is the formal cause of what it understands, particularly in the passage where she says that "Without [practical

²⁰⁰ I borrow the term from J. Schwenkler (2015, p. 8).

knowledge] what happens does not come under the description — execution of intentions — whose characteristics we have been investigating”²⁰¹, she prefaces this with the comment that “[the idea that practical knowledge is the cause of what it understands] means *more* than that practical knowledge is observed to be a necessary condition of the production of various results[.]”²⁰² And this would seem to suggest that although practical knowledge involves something *besides* making a material difference (namely, being a formal cause), it makes a material difference nevertheless. Frustratingly, however, she says little about how we are to understand this second aspect of practical knowledge, and it is admittedly difficult to envisage precisely what sort of account that she has in mind.

In the next section, and with the help of some of Anscombe’s friends, I will take up the task of articulating an answer to this question on her behalf. In the subsequent section, I respond to the concern that the possibility of practical failure undermines very idea of practical knowledge.

4.4 The material influence of practical knowledge

In order to appreciate the way in which practical knowledge influences what takes place — and not merely how what takes place can be described — it will be useful to consider what happens when such knowledge is *lost*.

The effect that I have in mind is characterised by the experience, familiar to most of us I presume, of temporarily forgetting what we are doing: We get up and walk into the living room and then cannot remember why. Or we find ourselves driving but all of a sudden lose track of where we are going. An example much discussed in the literature on practical knowledge is that of a man who is walking up Fifth Avenue when he suddenly loses track of the larger project for which his walk is a means: he does not know whether he is meeting a friend, or going shopping, or is just stretching his legs.²⁰³

The experience is akin to the phenomenon of noticing an aspect: Just as grasping the duck-aspect of the duck-rabbit drawing makes one’s visual experience take on an entirely different dimension, forgetting what one is doing has a similarly paradoxical effect on one’s perception of the surroundings: it seems to change both everything and nothing at all. Somewhat unlike noticing a previously unrecognised aspect, however, losing one’s practical knowledge is disorienting. Commenting on the case of the man walking up Fifth Avenue, Richard Moran describes the shift like this:

²⁰¹ (1963, p. 88)

²⁰² (1963, pp. 87-88)

²⁰³ The example is taken from J.D. Velleman (1989, p. 15).

The objects and scenes of his environment no longer have a role assigned to them in his ongoing action (as goal, obstacle, distraction, background, etc.). He sees Fifth Avenue, and he sees it from a particular perspective which indicates what direction he was heading in, but he does not see his goal. He does not see how his surroundings orient themselves with respect to his goal. Instead, he sees *everything* around him, the traffic going by, the meaningless street-signs, the strangers' faces, and nothing in those details enables him to discern a destination, a point to his being right here facing in this direction.²⁰⁴

Generalising the idea, Anton Ford argues that from the perspective of an agent engaged in her activity,

[...] the objects of perception all have a practical valence — a positive or negative charge, as it were — presenting themselves as to be avoided, or to be pursued. [This mode of perception] is not only interested, but discriminating; for everything in the environment has a role assigned to it in action. What is seen in this mode are things like goals, opportunities, obstacles, and distractions, all of which are determined by relation to an end.²⁰⁵

Let me elaborate upon this. The point is not, I think, that perceiving practically salient objects produces feelings of attraction or repulsion which thereby lead us to act in certain ways. *That* suggests a distinctly non-rational view according to which the perception of the relevant object produces an impulse to act. This does sometimes happen, I suppose (when one perceives pastries in a shop window, say), but cases like this are not to the present point.²⁰⁶ The idea is rather that for someone with practical knowledge, certain objects solicit attention *by* appearing in a way that makes certain courses of action seem reasonable — and others less so. For a man walking to the supermarket, for example, the protest march blocking the route appears as a serious obstacle to achieving his end and thereby *gives him a reason* for taking an alternative route. Were he simply walking to the coffee-shop on the corner, the march might still grab his attention but, if it did, it would not be *because* of its relevance for where he is going. The idea, in other words, is that for an agent engaged in their activity the particulars stand out — or drop into the background — not because they are intrinsically eye catching or because they excite some

²⁰⁴ (2004, p. 57)

²⁰⁵ (2016, p. 16)

²⁰⁶ I mention this view because one way in which the notion of 'practical valence' tends to get cashed out in the embodied cognition literature is precisely in terms feelings of attraction and repulsion. Setting aside the dubious phenomenological implications of this idea, the central problem is that it fails to explain how the perception of such particulars could potentially *justify* a given course of action. For discussion of this point in a different context, see M. Boyle (2016).

primordial appetite, but rather because of the rational relationships that they bear to the agent's activity in progress.

Now, much more would clearly have to be said in order to develop these ideas into a workable theory of practical perception.²⁰⁷ But for our purposes this brief sketch is enough, I think, to highlight the rather dramatic effect that practical knowledge has on the tendency towards completion of an activity in progress. For without it, the agent is simply left at sea. The man that *was* walking to the supermarket will never make it there unless he can remember what he was up to, because the progress of his deed is directly dependent on his *using* his practical knowledge to orientate himself in relation to the particulars with which he is confronted: With it, he is able to locate objects that afford him an opportunity to take the means necessary to achieving his end; without it, there are simply no means to be taken.

We might put this point another way by saying that practical knowledge affords the agent an ability to keep track of the progress of her deed. If Sarah is baking a cake, for example, she needs to keep a running tally of what needs doing and what has already been done: having just beaten the eggs, she must now add the sugar; having done that, the next thing to do is to make the cake topping; if, to her surprise, the cake topping has already been made by the sous chef, then it is not as if she will continue on blindly: no, she will adjust to the (in this case fortuitous) circumstances. On the other hand, if it is discovered that the apples that were going to be used are rotten, she may need to adjust her plan accordingly, say, by using plums instead.

Adjusting one's activity as needed, correcting mistakes, and doing what needs to be done when the time comes, depends, of course, on using one's perceptual faculties: this sort of

²⁰⁷ When I first drafted this chapter, I thought that the Moran/Ford point could be usefully understood within the framework provided by Gibsonian theories of perception. Such theories emphasise our ability to perceive possibilities for action or 'affordances' and remain popular in both the philosophy of perception and (certain branches of) perceptual psychology. It soon became clear, however, that the notion of an affordance was too narrow for the task. In the example given in the text, part of the point is that the valence of each object is hugely different: they appear as goals, distractions, obstacles, and so on. To describe all of these different forms of significance in terms of the actions that they afford, however, is to blur important distinctions between them. As Matthew Ratcliffe usefully summarises, "Things do not simply 'afford' activities; they appear significant to us in all sorts of different ways. It is not helpful to say that a bull affords running away from, while a cream cake affords eating." [...] "[What we need are rather] distinctions between the many ways in which things appear significant to us." (2015, pp. 61, fn.24). So while I still think that Gibsonian theories are a good place to start, they stand in need of further development if they are going to accommodate the phenomenon that Moran and Ford highlight. (For a step in this direction, see R. Dings (2018)).

Such a theory will of course also need to be complemented with evidence that goes beyond the sort of anecdotal example given above. Here, however, I can at best gesture towards the enormous literature on selective attention, as well as the relation between perception and action, in perceptual psychology — a literature that is at the very least consistent with what I have claimed here. In this regard see especially A.D. Milner and M. A. Goodale (1995), as well as the essays collected in *Attention in Action* (2005). Cf. also G. Rizzolatti *et al.* (1985; 1994) and S. Hurley (2001).

knowledge is essential to keeping the progress of one's deed on track. As Patricio Fernandez puts it:

[...] perceptual and other cognitive inputs (from memory, imagination, etc.) about available options, features of the situation, and so on are incorporated *all the way down* as the agent advances from a general conception of her end to ever more specific means to its realization: *each* step of the transition is possible only through the 'coming together' of the representation of the desired end with a cognitive act making available to the agent some feature of the means to realize it.²⁰⁸

But, and this is the crucial Anscombean point, in order to make use of the information received via one's senses in the service of the goal, one must be able to orientate the relevant particulars in their practical relations *to* it. And this, I suggest, is precisely what their practical knowledge enables them to do.

If this is right, then the agent's self-conscious apprehension of what they are doing can indeed influence what happens. It is a 'cause' not merely in the formal sense, but also in the sense that, without it, events end up taking a different course. And although it is true that the particular course taken also depends on a number of other things — such as whether the agent chooses to take the means available to her — this is no reason to disparage the account. On the contrary, the fact that the Anscombean position leaves room for the agent to play a role in the development of the process itself is, from my perspective at any rate, a boon.

4.5 Solving the problem of temporal unity

Admittedly, this last point does highlight the fact that practical knowledge is only part of the story when it comes to explaining the temporal unity of intentional activities in progress. After all, no matter how well orientated the agent is in relation to her surroundings, actually *taking* the means in question is not something that she can do unless she possesses the requisite ability and skill. No one can climb up a cliff-face, for example, without a certain degree of strength and technical capability: Even if the agent's practical knowledge of what she is doing serves to orientate her with respect to practically salient particulars, this knowledge is in and of itself of little use if her attention is drawn to something that she is unable to utilise. If anything, the ultimate success of her activity is beholden to the exercise of good judgment in *resisting* the temptation to act on particulars that lie beyond her powers — a skill that is part and parcel of being a good rock-climber, I imagine.

²⁰⁸ P. Fernandez (2014, p. 171)

There is no need to multiply examples. For the connection between intentional action and know-how is in fact widely recognised.²⁰⁹ And while there are numerous points of contact, for our purposes the most salient is that doing something *intentionally* seems to presuppose knowing how to do it. For as the example of the rock climber suggests, without the relevant know-how the climber has no hope of making it to the top. And without that, there is surely no tendency towards completion that would allow us to describe her as climbing up the cliff-face. If we generalise this idea into a principle, we arrive at the following idea:

S is doing A intentionally only if S knows how to do A.

I think that this principle provides the key to completing our account of the temporal unity of intentional activities in progress. And it is certainly *prima facie* compelling to suppose that, in acting intentionally, we are exercising our knowledge how.²¹⁰ But the full generality of this principle has been challenged and it will be illuminating to consider the sort of case that allegedly poses a problem for it.

Kieran Setiya gives the following example:

I am trying to defuse a bomb, staring with confusion at an array of colored wires. Which one to cut? In desperation, not having a clue what the wires do, whether they will trigger the bomb or not, I disconnect the red wire — and the timer stops. Even though I did not know how to defuse the bomb, and managed to do so through dumb luck, I count as having defused the bomb intentionally.²¹¹

It is unclear to me just how robust the intuition that Setiya is relying on here really is.²¹² But I do not intend to rest much weight on that. Instead, I want to pick up on the fact that, so described, Setiya's own conclusion is not in obvious conflict with the principle that I have outlined.²¹³ For notice that his case is one in which he has supposedly *done* A intentionally. The principle relates, however, not to complete events, but to events in progress. Thus, in order for

²⁰⁹ Anscombe (1963, pp. 88-89) notes the connection, but does not elaborate upon it.

²¹⁰ It is worth emphasising that notable endorsements of this principle come from people with very different views on knowing how and, indeed, intention. For example, the arch-intellectualists J. Stanley and T. Williamson voice their acceptance of the principle (2001, p. 415) (though they hedge their acceptance, somewhat, by saying that the principle may only be true of a subset of intentional actions), as do recent defenders of Rylean conceptions of knowing-how, such as Hornsby (2016) and Small (2012).

²¹¹ (2008, p. 404). See also Setiya (2012, pp. 285-286).

²¹² For discussion of similar cases and their less than clear interpretations, see L. Campbell (2015).

²¹³ Note that this is also the version of the principle that Setiya is considering, save that I have 'doing A' where he has 'doing ϕ '.

the example to pose a genuine threat to the principle, it would have to be true that he was, at some point or other, defusing the bomb. And notice that Setiya's original description leaves this question open: he was *trying* to defuse the bomb, he says, but that does not imply that he was ever actually doing it.

Now it might be thought that the relevant progressive judgment is nevertheless implied by the fact that he did what he was trying to do: he disarmed the bomb. After all, I earlier emphasised that the very idea of something having happened presupposes that there was a time at which it was underway. But as I noted (in a footnote — see p.75), this is only true of temporally extended events. Instantaneous events — such as starting to act, or seeing an aspect, or discovering treasure — by definition take no time. Rather, they characterise the temporally extensionless *boundary* between two states. As Galton nicely summarises, “The hallmark of an instantaneous occurrence is that we cannot catch it, so to speak, in the act of happening. At any moment it either has already happened or is yet to come.”²¹⁴ And it is generally recognised that predicates that admit of aspectual distinctions can function to describe *both* sort of event in different contexts.²¹⁵ For example, ‘I finished my thesis last week’ might refer either to the event that occurred (a) over the temporal period during which I was finishing up the last bits of it or (b) at the moment of time during which my thesis changed state from incompleteness to one of completion.

How, then, are we to interpret Setiya's claim that he has disarmed the bomb? Is this a report of an event that was happening and, eventually, finished? Or does it refer, rather, to the event that occurs at the precise instant at which the bomb transitions from a state of being armed to a state of being unarmed?

Consider the second interpretation first. In this case, ‘I disarmed the bomb’ is comparable to ‘I discovered El Dorado’. That is, although both things required the doing of various *other* things prior to becoming true, in neither case were the respective progressive judgments ever true. Maybe I was traipsing about in the jungle trying to find El Dorado before stumbling upon it. Setiya might even say that I discovered it intentionally. But it does not seem mandatory to hold that there was a period of time during which a process describable as ‘discovering El Dorado’

²¹⁴ (1984, p. 61) It is of course true that there are various temporal events *associated* with each of these examples, and also true that we may be unable to specify empirically the precise moment at which someone ‘notices an aspect’ (say). But the distinction is a logical, not empirical. As S. Rothstein observes, “this imprecision, which leaves room for contextual determination of what the endpoint of an event actually is, does not take away from the fact that the description of the event entails that there is a point (in part contextually determined) at which the event is over.” (2004, p. 7)

²¹⁵ In this regard see A. Mourelatos (1978, p. 419).

was taking place. And if this is right, then we can grant Setiya's conclusion without abandoning the principle.

On the other hand, if we interpret the claim as reporting a temporally extended event, then there must indeed have been some time at which the progressive judgment was true. But now we can raise the question of temporal unity: In virtue of what principle is precisely *that* process being instantiated? As I have argued, the idea that a given telic process is underway presupposes that the temporal stages of its instantiation show a real tendency towards completion. But given that Setiya *lacks* the relevant know-how, it is difficult to see how there could be any such a tendency. If his limbs were (somehow) under the control of a benevolent puppeteer and *they* knew how to disarm the bomb, then some such tendency might be present. But they are not and without any *other* reason to believe that he will succeed by anything other than dumb luck, we have grounds for denying that 'disarming the bomb' describes something that he was ever doing intentionally.

To drive this point home, consider how things change when a bomb disposal expert is placed in Setiya's position. Suppose that she too disarms the bomb, but, unlike Setiya, does so as a result of her expertise: she inspects the bomb, tests the wires, identifies the one leading to the detonator, and then snips it. In this case it would be possible to suppose that her know-how provides the requisite tendency towards completion that would enable us to say that she was disarming the bomb *in* doing each of these things.

Some will object to this. All that she was *really* doing in inspecting the wires, it might be said, was *trying* to disarm the bomb. That's fine. Again, I do not want to do battle over intuitions. But I do think that it is important to notice that our pre-theoretic judgments can be pulled to-and-fro depending on the details. If, for example, we suppose that the bomb is of a well-known design, with recognisable features that our expert is familiar with, and that the chances of success are correspondingly high, then it will probably be allowed that she is disarming the bomb the whole time. ('Why aren't you running away?! That bomb is live!' — 'Don't worry, I'm disarming it. It will only take a minute.') On the other hand, if the bomb is unfamiliar and success is far from guaranteed, then it will probably be said that she was testing the wire (e.g.) *in order to* disarm the bomb, rather than doing so because she *was* disarming it. As Anscombe usefully observes, "the less normal it would be to take the achievement of the objective as a matter of course, the more the objective gets expressed *only* by 'in order to'."²¹⁶ Indeed, if the

²¹⁶ (1963, p. 40) Cf. M. Thompson (2008, p. 132): "[...] the type of explanation of action at stake in action theory [...] is uniformly a matter of locating the action explained in what might be called a developing

expert is unable to identify the correct wire and just happens to snip the right one fortuitously, then I think that we are well within our rights to treat her case in much the same way as Setiya's: she was never intentionally disarming the bomb.²¹⁷

What this brings out, I think, is that the question of whether someone is doing something, or merely trying to do it, or is doing some other thing with the intention of doing it, etc., is not all or nothing. As Anscombe rhetorically asks: "Is there much to choose between 'she is making tea' and 'she is putting on the kettle in order to make tea' — i.e. 'She is going to make tea'? Obviously not."²¹⁸ The distinction is rather fluid and flexible and tolerably vague — a result of the fact that it depends, as Anscombe says, on the extent to which a certain future outcome can be taken 'as a matter of course.' My further suggestion at this point is just that the extent to which a given objective *can* be taken as a matter of course runs in parallel with the question of whether — and to what extent — the agent knows how to do A.

This puts us in position, I think, to hazard the following explanation of the temporal unity of an intentional action-in-progress:

We begin with the observation that an agent's knowledge how to do A puts her in a position in which she can choose to do A. If she lacks the relevant know how, then all she can do is *attempt* to do it; perhaps she will get lucky and achieve her aim. But in such a case she was never actually doing A — or at any rate, not doing A intentionally. Supposing, however, that the requisite know-how *is* in place and that she starts doing A, then her practical knowledge that this is indeed what she is doing serves to orientate her in relation to her surroundings and allows her to keep track of the progress of her deed. It does this, I have claimed, by affording her a mode of perception in which the objects in the environment acquire various forms of practical significance: they appear as goals, obstacles, means, and distractions (as the case may be). But to take advantage of them, she must exercise her knowledge-how. Doing so successfully means that her activity will be brought one step closer to completion, and so on and so forth until she ends up having done A. And since we have supposed that she is doing something that she knows how to do, we have reason to believe that it will be no accident if she succeeds.

process; it is just that this progress, development or 'imperfection' must be understood to exhibit various types or grades." See also K. Falvey (2000, p. 26).

²¹⁷ As Small usefully remarks, "[...] if someone who aspires to do B brings it about that he has done B, we can ask two questions: was the present progressive judgment 'He is doing B' ever true? And, if it was, what explained it? If its truth is explained by the agent's sound practical thought, we have a case of intentional action; otherwise, we have something less than that[.]" (2012, pp. 192-193)

²¹⁸ (1963, p. 40)

If this is right, then the two forms of practical knowledge — knowledge in intention and know how — together serve to explain the tendency towards completion that justifies the progressive judgment that she *is doing A*. Practical knowledge is indeed the cause of what it understands.²¹⁹

* * *

Before drawing this chapter to a close, let me briefly return to the objection that, due to our limited powers as material beings, the possibility of having our aims thwarted vitiates the idea that intention could serve to underwrite knowledge of a material process taking place ‘outside’ the mind. Although I cannot hope to respond to this objection in full here — to do so would be to vindicate the central and remarkable thesis of *Intention* — I want to at least attempt to blunt it.

The objection that I am imagining might be put like this: ‘According to Anscombe, an intentional activity is (i) an instrumentally structured process whose (ii) structuring principle is the agent’s knowledge in intention of what they are doing. But unlike God’s will, our intentions can be frustrated. We do not always do what we intend since the world does not always co-operate. A man may think that he is pumping water into a cistern when he is in fact not (since there is a hole in the pipe), for example. Since we cannot rule out that this will happen in any particular case, our intentions cannot provide us with knowledge of what is done intentionally — Anscombe’s theory is a non-starter.’

As it stands, the objection invites a cursory response: the mere possibility of failure in the case of action no more vitiates an intention’s status as a knowledge-bearing state than does the lack of absolute certainty undermine the status of a true belief as knowledge. But at this point the objection will no doubt be transformed into one about justification: ‘Given that our intentions can in general be frustrated, we need some reason to think that they are *not* being frustrated in this particular case if the intention is going to amount to knowledge. But that justification must derive from observing what is happening in the material world; thus, the knowledge that it underwrites is speculative, not practical knowledge.’

The suspicion that I want to raise is that this objection hinges on a failure to recognise the asymmetrical relationship that practical and speculative knowledge bear to their respective objects. For notice that the objection is predicated on the idea that one needs independent

²¹⁹ See also Small (2012, pp. 205-206) and Falvey (2000). Useful hints in this direction can also be found in Annette Baier’s work. See her (1970) and (1977).

justification for believing that a certain fact is true. This is wholly appropriate when the fact in question is itself independent of my own attitude towards it. For we need some reason to believe that it is not a happy accident that the two align. But we have seen that this is precisely *not* how we should understand practical knowledge and the intention that underwrites it: If I know how to do A and decide to do it there *already is* reason to think that it will be no accident if my intention aligns with the facts. If this is right, then the burden of proof is reversed. That is, one does not need independent reason for believing that one's intention is being executed; rather, it is special proof of *failure* that is needed to vitiate one's claim to practical knowledge.

In saying this, we need not deny the truism that much observational knowledge is needed, both to get into a position to act and to sustain the deed as it progresses. Indeed, it is precisely this role for perception that I have emphasised above. The point is just that the deliverances of perception (testimony, etc.) play a role that is subsidiary to the agent's practical knowledge. This is Anscombe's point in claiming that observational knowledge is akin to the role of the eyes in knowing what I am writing:

[...] isn't the role of all our observational knowledge in knowing what we are doing like the role of the eyes in producing successful writing? Once given that we have knowledge or opinion about the matter in which we perform intentional actions, our observation is merely an aid, as the eyes are an aid in writing. Someone without eyes [...] may not realise he is going over the edge of the paper on the table or overwriting lines already written; here is where the eyes are useful; but the essential thing he does, namely to write such and such, is done without the eyes.²²⁰

One of the lessons of the forgoing discussion is that this description is only half-right; it is not simply that the use of one's senses ensures that what gets done gets done *well* (as the analogy with legible writing would suggest). No, it is essential to ensuring that it gets done at all. But the analogy is useful since it highlights the fact that the role played by one's senses over the course of an unfolding activity is not that of providing us with reasons for believing that a certain proposition is true, but, rather, by providing us with reasons for acting: reasons for avoiding the crowd, or taking a left-hand turn, or for correcting our observed mistakes. Crucially, however, these reasons are not independent of the agent's practical knowledge: as the case of the man walking up Fifth Avenue shows, the fact that turning left affords a direct path to the supermarket is only a reason for one who is already going there.²²¹

²²⁰ (1963, p. 53)

²²¹ Cf. Falvey (2000, p. 36): "The ability to recognize, through observation, that one has done something unintentionally plays, as already noted, an important role in intentional action. But again I think this is an

4.6 Summary

I began this chapter by arguing that activities of animals differ from the processes triggered by inanimate objects and the vital operations of plants insofar as they admit of decomposition into parts that are coordinated by the animal. My claim was that this truism underlies the naïve thought that animals are agents in a way that inanimate objects are not: since effecting each part requires a separate contribution on the part of the animal, the animal is ‘triangulated’ as a constant source of motion of change.

As I conceded, this basic insight needs to be clarified. As a means to doing so, I began with some basic features of the way in which events and processes are represented generally. This led us to consider the question of what sort of principle unifies the temporal stages of a given processes and thereby justifies the predictive aspect of telic representation. In the case of aetiological processes, recall, my claim was that the various stages were unified by the relevant causal laws. But I also argued that similar principles are unavailable when we turn our attention to the activities characteristic of animal life.

It was here that I pivoted, somewhat obliquely, to Anscombe’s notion of practical knowledge and her corresponding conception of intentional action. I suggested that conceiving of practical knowledge as a formal cause helps to explain the intensionality of intentional action, but leaves us without an account of its material influence. To remedy this, I drew on work by some of Anscombe’s recent defenders to argue that we should see practical knowledge as affording us a mode of perception through which particulars take on various practical roles. This is essential, I suggested, to keeping track of the progress of our deeds.

In the remainder of the chapter I argued that knowing how to do A is a necessary condition for doing A intentionally. This put me in a position to offer an account of the temporal unity of intentional action and to respond to the objection that intention could ever be a knowledge producing state.

It goes without saying that I do not take myself to have vindicated Anscombe’s theory of intentional action in these few pages. But I hope to have said enough to highlight the theory’s attractive features, as well as to remove some of its most obvious hurdles. The question that I have not attempted to answer, however, is whether the theory can be applied to animals. It is to this question that we now turn.

enabling role; the agent’s ability to recognize a mistake as such, and his knowledge of how to correct it, are part of his ability to perform actions of a certain type.”

Animal Agency

“...what is so commonly said, that reason and cause are everywhere sharply distinct notions, is not true.”²²²

5.0 Preliminaries

In the previous chapter I made a case for the idea that our ability to see a phenomenon as a manifestation of agency is internally related to the instrumental structure of the material process itself: it is the complex means-end organisation of the action-in-progress, I have claimed, that allows us to triangulate a centre of motion and change — an *agent* properly so called.

I have also tried to bring the Anscombean understanding of what constitutes this structure, at least in the case of mature human beings, into view: it is the agent’s practical knowledge of what they are doing, why they are doing it, and how it is being done that unifies what happens into a rational order of means and ends. This knowledge also goes some way, I have suggested, towards explaining the genesis of the material process itself. By orienting the subject in relation to the practically significant particulars of her environment, her practical knowledge allows her to identify the means available to her. Together with her knowledge-how, practical knowledge thus serves to explain the temporal unity of the action-in-progress.

Can this account also be used to explain the temporal unity of the activities characteristic of animal life? Anscombe, for her part, maintains that animals do act intentionally and hence (one would suppose) that her account can be applied in this context. But she offers little by way of explanation of this assessment,²²³ and there are a number of reasons to be sceptical of it. Think again of the passage from Thompson with which I began this thesis:

²²² Anscombe (1963, p. 24)

²²³ As far as I know, Anscombe’s only explicit engagement with the question of whether animals act intentionally occurs in a paragraph-length discussion towards the end of *Intention* (1963, p. 86) and an equally brief series of remarks in ‘Under a description’ (1979, p. 221). Both passages are terse, however, and do not deal with the questions of practical knowledge that will concern me in what follows. (I discuss the passage

An agent, we may say, is something that operates on the strength of practical reasons or thoughts or considerations. It is something that can be viewed as doing one thing for the sake of another, according to concepts, or equivalently, I think, simply as a realizer of concepts. I see a process as a phenomenon of agency, in this sense, when I see the concept through which I describe or represent the process as itself at work in the genesis of the process I describe or represent.²²⁴

This view of intentional agency is deeply Anscombean. An intentional agent does not merely fall under the concept of something do-able; her knowledge that she is doing that very thing is part of what explains the progress of the deed in question. But remember the generality constraint: if we credit an individual with the compositionally structured thought that *a* is *F*, then we are committed to holding that they will similarly be able to think that *b* is *F*, for any *b* with which they are sufficiently familiar. The same point applies in the case of action. In particular, she cannot think of *herself* as doing *A* unless she can also entertain the thought that any number of other individuals might also do the same thing.

If this is right, however, then the Anscombean conception of intentional action seems to be even more restrictive than that of the Standard Approach. Not only does it presuppose the conceptual abilities associated with propositional content, it requires that the animal have conceptual representations of themselves and others.

Although there are ways in which this conclusion might be challenged, I propose to accept it: Anscombe's account *cannot* be straightforwardly applied to animal activities; animal activities and intentional actions are, I think, two species of a common genus — that of action.²²⁵ The question that I want to pursue in this final chapter, then, is whether we can develop a recognisably Anscombean account that is comparable to, but nevertheless distinct from, the account that applies to the intentional actions of human beings.

My strategy for doing so is based on two ideas. The first is that practical knowledge, like knowledge more generally, is internally related to the capacity to use a fact as a reason: an agent that is intentionally doing *A* can use the fact that she is doing *A* as a reason for (*inter alia*) doing *B* (where *B* is a means to doing *A*). The second idea is that the ability to use a fact as a reason

from *Intention* in §5.5 below.) In 'Practical Truth' (1993, p. 156) Anscombe allows that there is a kind of action that is unavailable to animals (and young children) on the grounds that animal action is not aimed at 'truth in agreement with right desire'. The connection with practical knowledge is here muted but left undeveloped. Her point in this paper seems to be that animals are not capable of deliberating about what it would be right to do. This seems plausibly true, but the concerns that I mean to develop are independent of this claim.

²²⁴ Thompson (2004, p. 352)

²²⁵ There are multiple ways of understanding the relationship between different species of a common genus. (For a useful overview in the context of action see Ford (2011). Although I shall not insist a specific understanding of it here, see Fernandez (2014) for a way of approaching the relation between intentional action and animal activity with which I am sympathetic.)

can be seen to lie on a spectrum: at one end there is the sophisticated, two-way ability to take a fact ‘into consideration or account’; in the middle there is the capacity to be ‘guided’ by a fact; and at the other end is the involuntary susceptibility to a fact’s imperious ‘control’.

Developing the conceptual tools needed to spell these metaphors out in a satisfactory way is the task of §5.2 and §5.3. In §5.4 I show how the tools that I develop might be applied to several examples of animal activity and, in §5.5, I bring these points, along with those made over the course of the last chapter, into a more general theory of animal activity. But first I need to spell out (i) the connection between knowledge and the ability to use a fact as a reason and (ii) why it is plausible to think that this ability lies on a spectrum.

5.1 Facts, reasons, and painful experiences

According to John Hyman, one way in which propositional knowledge manifests itself is by enabling the subject to use the fact known as a ‘reason’ — for “[m]aking a decision or an inference, believing, doubting, performing a calculation in one’s head, doubting or believing or hoping something, conceiving a desire or forming an intention[.]”²²⁶

To get the flavour of the idea, consider someone who is (intentionally) milking a cow. As we have already seen, such an agent knows what she is doing. Moreover, the fact that she is doing it can be used to explain why she is taking the means available to her. For example, were I to ask her why she is tugging on a cow’s udder, I could receive the answer ‘Because I’m milking it’. In this case she uses the fact that she is milking a cow as a reason for *acting*. But her knowledge enables her to use it as a reason for other things as well — as a reason for believing that she would be happier moving to the city, for instance. By contrast, if she does not know that she is milking a cow (perhaps because she has misidentified a bull and is tugging on something else), then the reason that explains her action is obviously not the *fact* that she is milking a cow (explanation is *factive*), but, rather, a different item of knowledge; the fact that she *intends* to be milking a cow, perhaps. Hyman’s idea is that this connection between knowledge and the ability to use a fact as a reason is no accident: *what it is* to know that *p* is to be able to use the fact that *p* as a reason.

Whether Hyman is right in thinking that knowledge can be defined in this way is not something that I am going to concern myself with here. But I do propose to assume that, at least

²²⁶ (1999, p. 439)

in paradigmatic cases, one knows that p if and only if one can use the fact that p as a reason.²²⁷ The question that I want to consider is how we might understand what this ability amounts to.

In the first place, and as Hyman emphasises, it is not an ability to do anything specific. We cannot explain the difference between knowing that p and knowing that q in terms of the particular actions that they are abilities to perform. It is an ‘adverbial ability’, he says, one that relates not to the *manner* in which something is done, but rather to the fact that in doing it one “takes the fact into consideration or account.”²²⁸

This is not especially helpful, of course, until we are told what is involved in taking a fact into consideration or account. Hyman says frustratingly little on this score and instead relies on the metaphor of being ‘guided’ by a fact to get the idea across (an idea I shall return to later). But at a minimum, I think that his talk of ‘consideration’ suggests that the ability is a two-way, voluntary power. That is, it is an ability to do or refrain from doing something and is only exercised when there are available alternatives. So, although someone who instinctively cries out in pain might be said to do so ‘because’ they are in pain, they are not using the fact that they are in pain as a reason (they do not *choose* to writhe). By contrast, someone who decides to take an aspirin because she has a headache does do so.

This last example is interesting because it is on this basis that Hyman argues (*contra* Wittgenstein) that it is possible to know that one is in pain. Moreover, he suggests that the ability to use the fact that one is in pain as a reason is what makes room for the distinction between *being* in pain and *knowing* that one is in pain. He writes:

The difference between ‘Katy knows she is in pain’ and ‘Katy is in pain’ is that ‘Katy knows she is in pain’ implies that Katy can be guided by the fact that she is in pain, whereas ‘Katy is in pain’ does not. A fish or a new-born baby cannot be aware of the fact that it is in pain — although it can of course be in pain — because a fish or a baby cannot be guided by reasons in the way it acts. Hence, if Katy is a new-born baby, ‘Katy is in pain’ may be true, but ‘Katy knows that she is in pain’ will certainly be false.²²⁹

This seems plausible. Naïvely, at least, it is difficult to take the suggestion that a baby — let alone a fish — is able to ‘consider’ the fact that she is in pain.²³⁰ But the deeper point is that

²²⁷ Note that this formula plausibly holds whether we are dealing with reasons for thinking that something is true (theoretical reasons) or reasons for thinking that it would be good or useful to do something (practical reasons).

²²⁸ (2015, p. 169)

²²⁹ (2015, p. 186)

²³⁰ Whether fish feel pain is a contested empirical issue. (For a somewhat sceptical overview, see Rose *et al.* (2014)). I take it as uncontroversial that many other species of animal do feel pain. (Although it is not an

someone with propositional knowledge presumably has a belief (or is at least capable of thinking a thought) with content that mirrors the structure of the fact known;²³¹ if so, then knowing that one is in pain presupposes the abilities that I have previously associated with propositional content. In particular, someone who knows that they are in pain will *ipso facto* be able to ascribe pain to others.²³²

I think that this point needs to be handled with care, however. For it is easy to describe the distinction between an infant (or animal) that experiences pain and a mature human being that knows that she is in pain in such a way that the relation between the two becomes mysterious. Thus, consider the following passage from Richard Rorty:

[...] the way in which the pre-linguistic infant [or animal] knows that it has a pain is the way in which the record-changer knows the spindle is empty, the plant the direction of the sun, and the amoeba the temperature of the water. But this way has no connection with what a language-user knows when he knows what pain is. [...] The mistake which Wittgenstein exposed was to assume that we learn what a pain is in the second sense by casting linguistic garb over our knowledge of what pain is in the first sense — by clothing our direct acquaintance with special felt, incommunicable qualities in words.²³³

On the most charitable interpretation of this passage that I can sustain, the idea is similar to Hyman's in one sense: while animals and infants might *feel* pain, they do not know that they are in it — not really. But unlike Hyman, Rorty makes the further claim that there is *no connection* between this knowledge and the experience characteristic of infants and animals.

Now the reason Rorty makes this further claim, I take it, is that we cannot imagine a pre-conceptual form of experience *justifying* knowledge proper: knowledge requires independent evidence, and this evidence must exist in a form that is fit to be taken up as the content of judgment. Thinking otherwise is to fall into the myth of the given (something to be avoided at all costs, I gather).²³⁴ Accordingly, Rorty thinks that he must say that the pre-conceptualised

article in a scientific journal, it would be remiss to fail to mention David Foster Wallace's sublime 'Consider the Lobster' (2004) in this context.)

²³¹ This is generally accepted as a point of common ground between classical justified-true-belief and ability accounts of at least theoretical knowledge. On this point, see Campbell (2017).

²³² Cf. Evans (1982, p. 103)

²³³ (1981, pp. 110-111)

²³⁴ The myth of the given is the idea that knowledge could be justified by experiences that are not already conceptually structured. I am following McDowell (1989, p. 287ff) in reading this passage of Rorty's as a response to this sort of worry. As McDowell puts it, "The fundamental point is the distinction between foundations and (mere) causal antecedents: non-conceptual pain (in pre-linguistic infants) is a causal antecedent of the ability to have conceptual pain episodes, *not a continuing ingredient in them which grounds the conceptual structures involved.*" (1989, p. 288)

experience of the un-enlanguaged bears ‘no connection’ to the knowledge that mature human beings have of their pain.

But this is a mistake. It is fostered by the idea that a language-user’s knowledge of her own pains is justified by her painful experience. But no one — animals, infants, and mature human beings included — knows that they are in pain *on the basis of* their pain. The reason is familiar: since being in pain and experiencing the pain are one and the same thing, the fact that one feels pain cannot serve as an evidential basis for judging that one is in pain.²³⁵ Rorty misses this point because, although he claims to be espousing lessons learnt from Wittgenstein, his position is actually diametrically opposed to the one that Wittgenstein recommends. In a notorious passage, Wittgenstein writes:

[H]ow does a human being learn the meaning of the names of sensations? — of the word ‘pain’ for example. Here is one possibility: words are connected with the primitive, the natural, expressions of the sensation and used in their place. A child has hurt himself and he cries; and then adults talk to him and teach him exclamations and, later, sentences. They teach the child new pain-behaviour.²³⁶

The basic idea is that the linguistic report is a natural extension of the primitive reactions to pain characteristic of infants and animals. They are *both* ways of expressing one’s pain. And while it is true that only the linguistic report can properly be said to be evaluable for truth or falsity (rather than sincerity or insincerity), it does not follow that the imposition of linguistic (subject-predicate) structure somehow transforms *what* gets expressed into something propositional. Rorty’s claim that there is ‘no connection’ here is the result of a failure to grasp this point, and as a result, he fails to understand why it is wrong to ask for *justification* for one’s claim to know one is in pain. An adult no more needs justification for that claim than a child — or animal — needs justification for writhing about.²³⁷

But can this Wittgensteinian point be accommodated alongside the plausible idea that propositional knowledge enables the subject to use the fact as a reason?

Indeed it can. To accommodate the Wittgensteinian point, we simply need to recognise that the fact that one is in pain can enter into the determination of what one does in various ways. Here is a rough sketch, to be expanded upon in due course. At one end of the spectrum the fact

²³⁵ Wittgenstein makes this point succinctly in *The Blue Book*: “ ‘How do you know that you have pains?’ — ‘Because I *feel* them’. But ‘I feel them’ means the same as ‘I have them’. Therefore, this was no explanation at all.” (1969, p. 68)

²³⁶ (2009, §244) I discuss this passage in greater depth in Cash (2017).

²³⁷ On these matters, I am in broad agreement with Finkelstein (2003; esp chapters 3 and 6).

that the organism is in pain *controls* what it does by causing the various instinctive responses which are familiar to us all. This is the way to describe the case of a dog that yelps because someone stepped on its foot — or, indeed, the case in which *I* do so. Moving gradually forward, there are cases in which the fact that one is in pain exerts a less imperious influence. The fact *guides* (as we might try saying) what the organism does causally (in the broad sense that without it, the organism would act otherwise), but in such a way that what it incentivises doing can be resisted. This is the way to describe a case in which you avoid putting pressure on a wounded foot, even though, in a pinch, you might decide to do it anyways. Finally, we reach the two-way *ability* to take the fact into consideration; to use it as a reason for inferring that it would be a good idea to see a doctor, or for complaining about it, or for any number of other things.

Again, this sketch is rough-and-ready in the extreme. It relies on the standard metaphors and these demand explanation. But what matters for me at this stage is just that the reader recognise the schematic — and surely plausible — idea that the fact that one is in pain can carry different degrees of cognitive significance for an organism, and that these degrees of significance belong on a spectrum. My suggestion is that, *contra* Rorty, the acquisition of conceptual abilities does not transform the experience of pain into one that is conceptually articulate; one who can speak does not experience *that* they are in pain. Rather, an organism with conceptual abilities can use those abilities to separate the experience from the subject undergoing it. By reflecting on the experience, the subject moulds the experience into a propositional shape suited to entering into acts of inference.

Of course, this proposal commits me to the idea that someone could only use the fact that they are in pain as a *reason* after their experience has been conceptualised in this way. She must, to use Rorty's metaphor, dress up a pre-conceptual experience in conceptual garb. And this will no doubt prompt the objection that the position that I am recommending is an un-tenable one: to suppose that something non-conceptual could be taken up into one's thought and thereby be made to stand in rational relations to others is to succumb to the myth of the given — or so one might say.

Spelling out this worry in enough detail to make it tractable would no doubt take us too far afield. But let me close this section by making two remarks that will hopefully alleviate the concern. (Readers who are unsympathetic to this kind of consideration in general can skip ahead to the next section).

In the first place, I think that the question of whether the painful experiences of mature human beings are in some sense conceptually 'shaped' is ultimately orthogonal to my project. I could allow that conceptual capacities transform the pain experiences of rational animals into

something conceptual — something whose shape is such as to be taken up and made an object of judgment, as McDowell likes to put it.²³⁸ I don't know what this means in the case of an experience like pain, but all that really matters for me is that we can make sense of the idea that a fact can enter into the determination of what an animal does (i) in a way that does not presuppose the conceptual abilities associated with propositional content, but (ii) nevertheless resembles the ability furnished by propositional knowledge.

Secondly, though, it is worth noting that one of the ways in which the transformative position that I am imagining gains plausibility in the case of *desire* depends on drawing a distinction between 'impulses' to act, such as the impulse to vomit associated with nausea, and states which present their objects *as* desirable, as typical human desires purportedly do.²³⁹ Only the latter are states with conceptual content, on this kind of view. But notice that it is presumably possible to know that one is feeling nauseous in exactly the same sense in which one knows whether they are in pain. Indeed, in drawing the distinction between impulses and typical human desires, Matthew Boyle points out that one might use the fact that one is nauseous as a reason for taking some medication.²⁴⁰ If that's right, however, and the ability to use this fact as a reason implies knowing it, then then someone that wants to defend a distinction between impulses and desires proper is already committed to the idea that the fact that one is in a state with non-conceptual content can be *used* as a reason (even if it doesn't present a certain course of action as reasonable). Accordingly, it would be possible to endorse the view that I am recommending in the case of certain states (the appetites, as we might call them) while maintaining a transformative conception with respect to others.

5.2 Emotions, incentives, and spatially significant auditory information

I have discussed the example of pain at some length because I want to make space for the idea that a fact might come to influence an animal's behaviour in ways that are recognisably similar to, if nevertheless distinct from, the sophisticated capacity to use a fact as a reason. Indeed, the case of pain is especially interesting since it relates to a fact about the subject in question. Mightn't the fact that an animal is hunting (say) influence what further she does in a way that is analogous to the way in which the fact that she is in pain does? If so, then we will have isolated a non-conceptual analogue of practical knowledge (or so I will suggest).

²³⁸ See, e.g., McDowell (1989, p. 285).

²³⁹ I am thinking especially of Boyle (2016, p. 540).

²⁴⁰ (2016, p. 540)

I expect this idea to sound far-fetched. But as a step towards making it plausible, consider the case of the primordial emotions — the feelings of hunger, thirst, tiredness, nausea, and sexual arousal, as well as the sensations that typically precede urination and defecation (to name a few). According to Derek Denton and his colleagues,

[...] a primordial emotion has two components — an imperious specific sensation and the compelling intention. The intention, like the primordial emotion, is specific. The genetic organization determines that severe thirst is generative of a powerful intention to drink — not to eat or copulate. Thus, for example, with suffocation, there is a ‘hunger for air’ sensation [...] and a compelling intention to fight for breath. Neuroimaging shows the concurrent air hunger and hypercapnia to evoke powerful activation in phylogenetically ancient areas of the brain[.]²⁴¹

There are several points worth remarking upon here. In the first place, the primordial emotions are similar to pain insofar as they have a strong experiential component. Moreover, and again like pain, the strength of the experience can range from subtle to severe. But these emotions also have a *directive* aspect to them: they produce a ‘compelling intention’, as the authors put it, to act in various ways.

I think that there are grounds for resisting this last characterisation, however. Chief amongst them is the fact that an intention (as this notion is typically understood) is subject to rational evaluation. If someone intends to operate heavy machinery, for example, they might change their mind after considering reasons that speak against doing so (‘I’ve been drinking and I suppose one shouldn’t mix alcohol and bulldozing...’). An intention is in this respect like a motive: you can argue against it.²⁴² Yet it would surely be absurd to suggest that the inclination to fight for breath is something that could be given up as a result of rational considerations (even if, in certain circumstances, such considerations might allow the agent to *resist* acting upon it).

Nevertheless, I think that something is right about the idea that primordial emotions direct organisms to act in ways that extend beyond the instinctive forms of response characteristic of pain. Thirst, hunger, and sexual arousal can obviously lead an organism to act in pursuit of distant objectives, as the animal’s instincts demand (I return to the topic of instinct below). But to capture this idea we need a notion that does not carry the cognitive associations of intention. To this end let me borrow a term from Christine Korsgaard and say that the primordial emotions produce *incentives* to act.

Korsgaard explains the idea as follows:

²⁴¹ (2009, p. 501)

²⁴² See Anscombe (1963, pp. 31-32).

An incentive is a motivationally loaded representation of an object. I am using the term ‘object’ broadly here to include not only substances but also states of affairs and activities. The object may be actually perceived, or conceived as a possible item in the environment, a way that things might be. You are subject to an incentive when you are aware of the features of some object that make the object attractive or appealing to you. Perhaps the object satisfies one of your needs; [...] It interests you, it arouses the exercise of your faculties, it excites your natural curiosity, or it provides some sort of emotional comfort or satisfaction.²⁴³

I think that Korsgaard’s description can be usefully understood in terms of the primordial emotions and *vice versa*. The basic idea, as I am proposing to understand it, is that an emotion can *incentivise* a certain course of action by producing a representation of it which the animal can then utilise in deciding what to do (as I shall argue later (§5.4), where there is no room for choice, there is no room for incentivisation either). Like an intention or desire, an incentive’s intentional object is not a proposition, but something do-able: for example, to hunt, eat, or copulate. But it differs from them, I suggest, insofar as it does not necessarily present its object as something that there is a *reason* to do.

This will no doubt sound cryptic, but the basic contrast is easy enough to grasp. The idea is just that, paradigmatically, a desire (or intention) presents its object in such a way that the subject knows *why* they desire it: e.g. ‘I want to go out for dinner with my mom because... I haven’t seen her for a while and that’s the sort thing that good sons do/ she’ll pay and I can save some money/ I want to, that’s all/ etc.’ As the last answer indicates, one doesn’t *need* a further reason for wanting something. But it seems built into the very notion of desire, at least as it appears in the case of rational animals and is distinguished from the appetites, that desires present their objects in such a way that the question *has application* (as Anscombe might have put it).²⁴⁴

My minimal claim, then, is that an incentive is not like this: an incentive to do A does not presuppose — though it certainly does not exclude — apprehension of what it is about doing A that incentivises one to do it. And this is surely plausible: nausea can produce an incentive to vomit regardless of whether the animal is capable of understanding why they are subject to it. (As I shall suggest in §5.5, the capacity of an animal to integrate its incentives with its more distant goals and to make decisions on this basis marks a significant cognitive achievement.)

As Korsgaard’s description indicates, incentives can also be produced by perceptual experiences and it will be useful to consider how the idea might find application here. To this

²⁴³ (2009, pp. 104-105)

²⁴⁴ For discussion of these points see especially Boyle (2016) and Boyle and Lavin (2010).

end, consider the experience of hearing a sound as coming from a certain direction in space. You hear it as coming from your right or left, from above or below you, etc., in such a way that you can immediately orientate yourself in relation to it — by turning in its direction, or by pointing towards it, or, perhaps, by running away from it. In discussion of this case, Evans remarks that

[w]hen we hear a sound as coming from a certain direction, we do not have to think or calculate which way to turn our heads (say) in order to look for the source of the sound. If we did have to do so, then it ought to be possible for two people to hear a sound as coming from the same direction and yet to be disposed to do quite different things in reacting to the sound, because of differences in their calculations. Since this does not appear to make sense, we must say that having spatially significant perceptual information consists at least partly in being disposed to do various things.²⁴⁵

Evans's negative claim in this passage might be challenged: is it absurd to suppose that two people might hear the sound as coming from a certain direction and yet be disposed to do different things? After all, and as he admits, the things that an individual who hears the sound as coming from such and such a direction will be disposed to do are incredibly various. As such, there doesn't seem anything *obviously* absurd in supposing that two people might be 'disposed' to react to the directionality of the sound in different ways. Perhaps person A is the curious type and is disposed to investigate the noises that she hears whereas person B tends to ignore them. This is presumably not the sense of 'disposition' that Evans has in mind; but if that notion is going to cut any cloth, it needs to be sharpened.

In any case, I think that it is the positive claim that is more interesting. On my (perhaps overly charitable) interpretation, Evans is claiming that there is an internal relationship between hearing the sound as coming from a certain direction and knowing-how to orientate oneself in relation to it. And this does seem plausible. For although it is presumably possible to misperceive the actual direction from which the sound came (one hears it as coming from the right when it in fact originated from the left), or perhaps even to hear it *as* coming from a certain direction without knowing how to orientate oneself to it in special cases,²⁴⁶ our ability to interpret one another as experiencing the directionality of the sound in egocentric terms only makes sense against a background of shared responsiveness to the sound's direction of origin: our instinctive

²⁴⁵ (1982, pp. 154-155)

²⁴⁶ For example, it seems to be an intelligible possibility that the receptive and active components *normally* involved in hearing the direction of the sound can become dissociated as the result of, say, a stroke. Thanks to Michael Potter for drawing my attention to this issue.

responsiveness is part of what enables us to make sense of the basic egocentric notions *right, left, up, down, front, back*, and, hence, to understand one another's experiences in these terms.²⁴⁷

Now it is true, of course, that mature human beings do not always 'react' to a sound in the same way: the forms of responsiveness to hearing a sound as coming from the right (say) are incredibly various and include doing nothing at all. But to repeat the Wittgensteinian point: the two-way power to choose whether or not to react to the sound (to use the fact that the sound came from such and such a location as a *reason* for, say, shutting a particular window) is a natural outgrowth of our primitive reactions to its directionality — reactions that lie on a range from the startled head-movements of a new-born, to the curious searching for the source of the sound characteristic of infants, to the ability to point in its direction possessed by young children, and onwards towards the mature ability to describe its direction of origin in egocentric terms.

It is in partial explanation of this fact that the notion of an incentive is useful, I think. As applied to this sort of case the idea is that, at one end of the spectrum, the experience and the primitive reaction are rolled into one, so to speak. In this sort of case the experience is the subjective aspect of the instinctive response; it occurs more or less simultaneously (or even after) the response (as William James said of instinct and emotional excitements that accompany them, the two "shade imperceptively into one another."²⁴⁸) Moving forward, however, the experience and the instinctive response can come apart. Hearing a commotion as coming from outside one's window might prick one's curiosity, for example, but the urge to look can be resisted (especially so, if one is doing something else — like trying to finish a thesis). My idea is that in cases like this the response is 'coiled', so to speak, into a directive representation produced by the experience — an incentive. Indeed, even when we reach, at the other end of the spectrum, the sort of experience that merely indicates the sound's direction of origin (i.e., its motivational force is nil), I suggest that the link between the experience and the capacity to orientate oneself in relation to the sound is maintained through a primitive representation of how to turn in response to it — a representation that says, roughly, *to turn in the sound's direction, turn to the right*. At any rate, that is how I am proposing to understand Evans's claim that "the spatial information embodied in auditory perception is specifiable only in a vocabulary whose terms derive their meaning partly from being linked with bodily actions."²⁴⁹

²⁴⁷ Just a part, mind. The links between perception and action that I am envisioning doubtlessly extend to other sensory modalities.

²⁴⁸ (1890, p. 442)

²⁴⁹ (1982, p. 157)

There are two points that I want to emphasise at this stage. The first is that the notion of an incentive helps us to understand how a fact can enter into the determination of what one does in ways that are different from the ability to use a fact as a reason. The fact that a sound came from one's right, or that one is hungry can *guide* — to adopt Hyman's metaphor — what one does by incentivising certain courses of action. But the form of guidance that incentives provide is different from that of, say, desire. As I indicated a moment ago, the fact that I desire such-and-such directs the subject to perform a certain action in such a way that its instrumental connection to a further end (if such there be) is known to the agent: desires enable the subject to answer the question 'Why?' But there is a further difference in relation to the question 'How?' For desires are, so to speak, pure: they present their object as something that it would be good to do, but they leave the question of how to achieve goodness open. Otherwise put, desires must combine with (true) beliefs about the instrumental relation between the available courses of action and the desired end in order to guide an activity to its successful completion. As Millikan puts it,

[...] representations that tell what to do have no utility unless they can combine with representations of facts. It follows that a capacity to make mediate inferences, at least practical mediate inferences, must already be in place if an animal is to use [...] purely directive representations.²⁵⁰

One of the lessons of the forgoing, however, is that the incentives produced both through directional perception and emotion are *directly* relevant to action in such a way that we do not need to think about *how* to achieve the courses of action that they recommend. And the reason for this, I want to suggest, is that what the incentive is an incentive *to do* — i.e., what its content is — is in part determined by the agent's *instincts* and *skills*. It is to an elaboration of these notions that I now turn.

5.3 Instinct and skill

According to Merriam Webster, an instinct is “a largely inheritable and unalterable tendency of an organism to make a complex and specific response to environmental stimuli without involving reason.” If we broaden this definition to include internal stimuli, it serves to capture the basic notion that I have in mind. As I shall put it, an instinct is an unlearned predisposition to act in accordance with a rule that links environmental, physiological, and psychological

²⁵⁰ (1995, p. 192)

inputs to specific forms of action and reaction — to fight for breath when suffocating, for example.²⁵¹

Like dispositional notions generally, the term ‘instinct’ can *also* be used as high level description of the trait(s) that serves to partly explain the tendency: just as a fragile vase (i) is one that has a disposition to break when dropped and (ii) breaks when dropped *because* of its fragility, so, too, a person’s tendency to fight for breath when suffocating can be explained by their ‘instincts’, where this is a cover term for whatever traits go into the explanation of that tendency. I will mark this distinction with the subscripts ‘instinct_[t]’ (when referring to the tendency) and ‘instinct_[e]’ (for the explanatory trait); when both are intended, I leave the subscript off.

Now in some organisms the links between the sensory input, the resulting physiological state, and the ensuing reaction may be unmediated by any kind of experience. This is the way in which to conceive of the tropistic behaviour of magnetosomes, clams, and fly-traps, I assume: sensory stimulation leads directly to behavioural output without any need for intermediary perceptual or emotive states. There may also be organisms (frogs perhaps — see below §5.4) that are capable of perception and emotion but for whom the links to action are invariant and unmediated by incentives. As before, the animal’s instincts_[e] play an imperious role in the determination of what happens: they completely control what the organism does in response to what without any co-ordination on the part of the animal.

But for many species, I think, the response *is* mediated by experiences that produce incentives to act: dogs, for example, can learn to resist the impulse to urinate when in the house. And my suggestion is that in such a case we need to think of the instinctive tendency as being coiled into an incentive: this explains how the reaction and the emotional experience can come apart whilst maintaining their internal relationship to one another.

In this last sort of case, the incentive produced by the emotion is of the same general type no matter what organism is in question: the need to pee produces an incentive to urinate (if it does) regardless of whether it is me, or a dog that is in question. But in other cases, the instincts_[e] characteristic of the species explain what the relevant incentive is an incentive to do.²⁵² For

²⁵¹ The appeal to rules, tendencies, etc. raises a host of philosophical problems that I cannot enter into here. I must rely on the naïve idea that we can isolate functionally significant links between various perceptual/emotive cues and the activities which follow them — where ‘functionally significant’ relates both to a broad notion of fitness, as well as more specific ‘goods’ which characterise the species of animal in question. On these topics see and cf. Thompson (2008, pp. 63-82), Burge (2009, pp. 267-278), Boyle (2012, pp. 403-406), and Foot (2001; chapter 3).

²⁵² This is not to deny that in the former case what it is an incentive to do is determined by the animal’s instincts_[e] — the point is just that they are shared across species.

example, although a leopard and a vulture might both experience hunger, their differing instincts_[e] mean that the incentive produced will in each case be different — the leopard's hunger directs her to hunt, whereas the vulture's hunger motivates scavenging.

I have been considering cases in which the animal's instincts produce incentives in response to emotional experiences — such as hunger — which then direct the animal to *start* acting in certain ways. However, it seems plausible that the converse is also true — i.e., that certain instincts only become operative once an animal is *already engaged* in the relevant activity. For example, even if they are hungry, wolves will often ignore prey that stumbles upon them: rather, their instinctive tendency to chase, capture, and kill prey only becomes operative after they have already started their approach.²⁵³

Of course, similar tendencies can be learnt, developed, and perfected through experience. These are what I call *skills*, though the name is admittedly less than ideal.²⁵⁴ For me, skills are just the acquired analogue of instincts. Accordingly, they can similarly be divided into the tendency to act in accordance with a rule linking a complex form of response with various forms of inputs, on the one hand, and the acquired traits which serve to partly explain the tendency, on the other. (As before, I shall mark the distinction with subscripts.)

The distinction between instincts and skills is not, I think, a sharp one. Much learning in the wild involves the *refinement* of instinctive tendencies (as when predators learn to avoid certain prey in preference of others), and seemingly novel behaviours, like termite and ant fishing in chimpanzees, bear some connection to instinct. Thus, captive chimps seem naturally inclined to put long thin things through holes even when there is no prospect of reward.²⁵⁵ Still, it is certainly true that in the wild the chimpanzees refine this instinctive tendency in impressive ways: they learn craft sticks one way for fishing out ants and another way for termites.²⁵⁶ There is even some evidence for different methods between communities.²⁵⁷ Whether we call these refined instincts or skills is a terminological choice; I will opt for the latter.

Needless to say, the topic of skill and of the various forms of learning available to animals deserves a richer treatment than I can provide here. I invoke the idea of skill here simply because it is obvious that many animals do acquire new — and refine old — practical tendencies. But

²⁵³ See Mech *et al.* (2015, p. 8). I discuss wolf-hunting strategies below (§5.4).

²⁵⁴ It is less than ideal because the topic of skill has received treatments of late that emphasise its connection with *propositional* knowledge (e.g., Stanley and Williamson (2017)) and its internal relation to the capacity to instruct others (Small (2014)). Both of these connections are interesting and plausible, but I do not want my notion to carry the same associations.

²⁵⁵ Gould and Gould (1998, p. 55)

²⁵⁶ Byrne (1995, pp. 96-97)

²⁵⁷ Luncz and Boesch (2015, pp. 72-73)

what is ultimately important for me is not how animals acquire their skills, or how to distinguish between learnt and instinctive tendencies. What I want to focus on is rather how to understand the way in which skills_[e] and instincts_[e] interact with incentives, when they do, in the context of a complex activity. To this end it will be useful to see how these ideas might apply in practice.

5.4 Some applications

(1) Start with a simple case — a fly passes in front of a frog and is snapped up. As I pointed out in §2.2, frogs are doing more than merely responding to surface irritations. They perceive the object's real size and are thus able to distinguish their quarry from obstacles and threats (though not, of course, from anything else that is small-dark-and-moving).²⁵⁸ But if philosophical legend is to be believed, frogs are unable to inhibit their instinctive response to any of these perceived objects: upon detecting something that resembles a passing fly entering snapping range, the frog reacts invariably; deviation occurs only when the instinctive response gets overridden by detection of a threat.²⁵⁹

Suppose, for the moment, that this is accurate and that, in general, the frog's entire repertoire of activity can be modelled according to a few simple rules relating perceptions to action ('if something small dark and moving comes within range, snap'; 'if it is out of range, approach and then snap'; 'if a threat is detected, hop away'; and so on). We can think of these decision rules as describing the frog's instincts_[i]: they describe perception-action cycles which the frog's nature has predisposed it to act in accordance with.²⁶⁰

Froggy instincts_[i] are rigid, as far as I am aware: frogs cannot learn to inhibit them or adapt them in light of previous experience. And this means, I think, that there is no role for the notion of an incentive to play. Frogs do not perceive their prey and *then* feel incentivised to act. To hold otherwise would be to overcomplicate matters. For remember that the point of bringing in the notion of an incentive is to explain how an individual might control their response whilst retaining the ability to act in accordance with the instinctive tendency. When there is no question of control, there is no need for incentivisation either.

What would happen if we imagined that frogs did not snap invariably, but only when hungry? Here I think that it is tempting to say that the frog's hunger produces an incentive to find food, which its instincts_[e] further determine into an incentive to snap at anything small, dark, and moving that passes in front of its eyes. However, it is not clear that this is necessary.

²⁵⁸ Ingle (1998)

²⁵⁹ Neander (2006, p. 171)

²⁶⁰ Cf. Millikan (2004, p. 166ff)

For if the onset of a certain physiological state leads the frog to invariably adopt a new pattern of activity which is then carried out rigidly, it seems an easy task to simply tweak our description of the frog's instincts^[1] so as now to include the conjunctive condition 'when hungry' — e.g., 'if something small dark and moving comes within range *and* you are hungry, snap'. As before, the introduction of incentivisation would appear to merely add on an idle wheel to a process that is more perspicuously explained without it.

I do not know whether this more minimal account is indeed the right way to think about frogs. Certainly, the mere fact that a pattern of behaviour can be *modelled* according to rules like these does not entail that the animal is *not* being motivated to act as the result of an incentive.²⁶¹ What I do want to suggest, however, is that forms of activity that are unmediated by incentivisation are limiting cases of animal agency. This is not to deny that the frog's behaviour has an instrumental structure to be discerned and accounted for: a frog that hops towards its prey and then snaps it up does two things, and one for the sake of the other. But this instrumental structure is constituted behind its back, as it were.²⁶² That is, natural selection has provisioned it with a tendency to procure prey by hopping and then snapping — a tendency that, when the conditions are right, results in the procurement of nourishment. And while the coordination between the environment and the behaviour requires the frog to utilise its perceptual faculties to distinguish between features of its environment, what it does in response to these distinctions seems to lie entirely beyond its control: it is at the mercy of its instincts.

(2) In order to see how the notion of incentivisation might enter into the explanation of an activity I want to consider another much-discussed example: the prey-retrieval behaviour of the great golden digger wasp (*Sphex ichneumoneus* — henceforth simply 'wasps'). I'll begin with a brief description of the background of the activity before going into more detail about the particular stage of the process that I want to discuss.

²⁶¹ For discussion of this point in a related context see Buckner (2011). Still, there does seem to be evidence that the purely instinctive form of explanation is on the right track. As Papineau explains, "the sensory information which guides the frog's prey-catching behaviour is not available to other systems of behavioural control. One channel of sensory information guides its fly-catching behaviour, another guides its obstacle-avoiding behaviour, and yet another its ability to jump away from looming threats." (2003, p. 117) As such, the information picked up influences behaviour quite directly, seemingly without the need for mediation by any form of incentivisation. Accordingly, Milner and Goodale suggest that "[...] one might characterize the organization of the visual midbrain [of frogs] not as a visual system, nor even as a set of visual subsystems, but rather as a set of visuomotor subsystems." (1995, p. 11)

²⁶² I have borrowed the metaphor from Fernandez (2014) and Ford (2008, p. 156).

Wasps construct intricate burrows in the soil where they lay their eggs. The burrows consist of a long central tunnel with various side chambers where the larvae and prey are deposited.²⁶³ At some point after the construction is complete, the wasp leaves the nest in search of food. Having captured and paralysed her quarry, she returns to the nest and initiates a complex process of getting the prey into the brood-chamber.²⁶⁴

It is this last pattern of activity that I want to focus on. To begin, (A) the wasp deposits her catch in the ‘normal position’: four to five centimetres away from the edge of the burrow, orientated so that the prey’s antennae face the entrance (the wasp typically uses these to grip the prey, though she can also employ other means if they are destroyed). If there is a disturbance around the nest — if, for example, there is an ant infestation, or she has been pursued by a parasitic fly — she will back into the tunnel and remain alert, sometimes exiting to bite at the ants or lunge at the fly, shooing it away. (B) Once all is clear, she then enters the tunnel (headfirst) and proceeds to remove dirt from the side tunnel leading to the brood chamber. At this point (C) she enters the brood chamber, turns around (she cannot do this in the main tunnel) and then returns to the surface, before (D) pulling her prey inside.

Like the frog’s snapping response, this activity is unquestionably the product of instinct_[e]: the wasp has an unlearned tendency to initiate the process of getting the wasp into the nest just described after returning to the burrow. Indeed, we can think of the transition between each stage as a further instinct_[i]; that is, the wasp has a tendency to do B once A is completed, to do C once B is completed, and then to do D once C is completed.

This case is usually trotted out as the example *par excellence* of an activity which looks complex on its surface, but is in fact controlled entirely behind the animal’s back. The wasp’s instincts, it is supposed, fully determine what happens after each stage.

This idea, however, is founded on a fable, popularised by Daniel Dennett and Douglas Hofstadter,²⁶⁵ according to which the wasp can be caught out by a simple trick. The trick is played like this: once the wasp has entered the burrow, an experimenter moves the prey outside of the normal area. The wasp reappears, drags the prey back to the edge of the burrow, and then re-enters only to emerge and find that the prey has again been moved. So she drags it back, re-enters the burrow, and so on, indefinitely. To some it will thus seem as if the activity is

²⁶³ For a nice description, see Brockman (1980, pp. 428-431).

²⁶⁴ What follows is a much condensed version of the description provided by Brockman (1985, pp. 632-642).

²⁶⁵ See, e.g., Dennett (1973, p. 169ff) and Hofstadter (1985, p. 528ff).

“nothing more than an automatic reflex, genetically hardcoded behavior. It’s just like a calculator that adds 10 plus 10.”²⁶⁶

If the wasp’s activity is truly no more than automatic reflex, then it would seem that there is again no need to bring the notion of incentivisation into the picture: the structure is fully explained through a series of simple decision rules — rules which describe the wasp’s instincts^[1] and which are implemented by means of genetically inherited mechanisms. But let us consider some of the data. Contrary to the story as told by Dennett and Hofstadter, George and Elizabeth Peckham report the following observation:

[...] when she had brought a grasshopper to the entrance of the nest, and while she was below, we moved it back five or six inches. When she came out, she carried it to the same spot and went down as before. We removed it again, with the same result, and the performance was repeated a third and a fourth time, but the fifth time that she had found her prey where we had placed it she seized it by the head, and going backward dragged it down into the nest without pausing.²⁶⁷

They conducted the same experiment on successive days with similar results: after initially repeating the process, the wasp eventually circumvented the trick (on the fourth and seventh trials, respectively).

This observation was made over 100 years ago, but the most recent study (that I could find, at any rate) provides support for it. Jane Brockman reports the following experiment, in which she moved the prey to one of fifteen non-normal locations (on successive trials) before leaving it in the normal position on the sixteenth. Once the wasp reappeared at the entrance, she gave it two minutes in which to retrieve the prey before ending the trial. She describes the results as follows:

Thirty-one wasps were tested with the following results [...]. (1) Twelve of the wasps pulled the prey into their nests only when it was in the normal position and thus these wasps completed all 16 trials of the experiment. (2) Ten wasps took the prey into the nest from some location other than the normal position and hence ended the experiment early. (3) Five wasps gave up searching for the prey before the two minutes elapsed (one each on trials 5, 14, 15 and two on trial 6) returned to the burrow and then left the area. (4) Two wasps picked up the prey and left with it (both on trial 9) (5) one wasp set the prey down so close to the nest entrance (on trial 12) that it

²⁶⁶ Fogel (2000, p. 284)

²⁶⁷ Peckham and Peckham (1905, pp. 304-305)

fell in spontaneously, and (6) one pulled the prey into a neighboring burrow (trial 9) which was very close to her own nest.²⁶⁸

These results are interesting in a number of respects. In the first place, some of the wasps do indeed appear to be susceptible to the trick: twelve of them went the entire time without dragging the wasp inside. But others circumvented it: after initially falling for the trick, ten wasps eventually pulled the prey inside the borrow without first re-entering. And yet others acted differently: some gave up, others took the prey and flew off elsewhere

How are we to explain this variety if we are left only with the decision rules that (let us assume) explain the prey-catching activities of frogs? The first thing to notice is that the difference cannot be put down to the presence or absence of a simple conjunctive condition such as whether or not the wasps are hungry. The prey is being retrieved for their young, and, in any case, each wasp is *already engaged* in the relevant activity when the variation occurs. Still, many will no doubt be tempted to think that all that is needed is a more complex algorithm. All that either experiment shows, it might be said, is that the activity cannot be accounted for purely in terms of instincts_[e] that relate *present-tense* stimuli to behavioural outputs; the inputs to the wasp's 'programming' must therefore include past failures.²⁶⁹

On its own, of course, this does not account for the data. In the first experiment, the prey was pulled in after different numbers of repetitions, while in the second there is obviously diversity between the individual wasps as to whether they bring it in at all. Admittedly, the experiment did not control for things like weather conditions, nest locations, and other features that might be thought to influence the wasp's activity; so it remains a dialectical possibility that an algorithmic form of explanation *could* be found. But it seems clear that things are going to get complex very quickly and, I suspect, somewhat *ad hoc*. At any rate, to demonstrate otherwise is the challenge that I want to pose to anyone who thinks that there *must* be a purely instinctual form of explanation in this case.

Here is an alternative possibility: the activity *is* structured by the wasp's instincts_[e], but the instincts_[e] operate by producing incentives to act as the activity unfolds: B is incentivised after A is completed (in normal conditions), C is incentivised after B is completed (in normal conditions), and D is incentivised after C is completed (in normal conditions). In the experiment, the conditions are abnormal after C — the prey is not on the threshold. This produces a new incentive: to look for it. Having found it, the wasp's instincts_[e] incentivise starting over with A,

²⁶⁸ Brockman (1985, pp. 140-141)

²⁶⁹ Thanks to Michael Potter for pressing me to get clearer on the dialectic here.

and then proceeding to B and C. But after several repetitions, doing B and C after A can become *de-incentivised*, to a greater or lesser extent. For compare: after hearing the same sound coming from a certain direction over and over again, the incentive to look in its direction becomes less and less powerful. If this analogy is on the right track, then it begins to make sense why certain wasps, upon finding the prey, will do something other than A. For the incentivising effect no longer carries the same kind of imperious influence; the wasp achieves some distance from the instinctive reaction.

But what of the wasps that do not figure out the trick? How are we to explain the *diversity* exemplified by the different forms of response of the individual wasps? The first thing to note about those that never figure it out is that there is actually a functional explanation for re-entering the burrow after bringing the cricket to the entrance: it provides the wasp the opportunity to turn around in the brood chamber. Thus, one of the interesting features of the wasps that drag the cricket into the den from a non-normal location is that they need to adopt a *new* pattern of behaviour: they need to turn around outside the burrow and to then back in with the prey.²⁷⁰ Each of these are things that the wasps are able to do in other contexts (recall that, if there is a disturbance at the nest when the wasp originally arrives she will turn and back in), but not together. The wasp must therefore de-couple and recombine means and ends into a new pattern of behaviour. No small feat, I suspect, especially when the wasp has an incentive to re-enter the burrow headfirst.

There is not enough information here to form any determinate judgments about how the wasps that circumvent the trick do it. I would certainly not want to make any definite claims about insight or creative uses of the imagination (though neither would I rule either out on the presumed simplicity of the organism in question). But at this stage the important point that I want to emphasise is just that *part* of what makes this sort of re-combination possible is surely the fact that its instincts are not imperious. Rather they operate via the production of intermediary incentives that the wasp can, ultimately, resist. And if this is right, then although her instincts *contribute* to the explanation of what the wasp does in response to what, they are not the whole story: unlike the frog, she is not entirely at their mercy.

(3) Perhaps the previous example will occasion some doubt due to the presumed simplicity of the organism in question. As a final case study, then, let me show how these ideas might be

²⁷⁰ On this point, see Keijzer (2013, p. 513)

applied to what is a presumably ‘higher’ form of animal life. As before, I’ll offer a brief description of the activity before commenting on its theoretical interpretation.

Wolf-hunts are initiated when hunger prompts the wolves to leave their resting area and to begin searching for prey, the direction of travel being determined in part by their past experiences of success and failure.²⁷¹ After locating their quarry (via vision, scent, or tracking),²⁷² the wolves begin to approach, sometimes by stalking but at other times by moving towards their prey in plain sight.²⁷³

The next stage of the hunt occurs when the wolves stop their approach and survey the prey. This need not happen, but if it does, it can last for anywhere between minutes and hours. At this point, the wolves might leave the prey and continue searching, or they may approach the prey further, or they may approach and then proceed to the initial attack phase.²⁷⁴ This is the stage during which the wolves either pursue a fleeing group of animals or instead start to lunge and harass a group that stands its ground.

In order to bring the hunt to a successful conclusion, the wolves need to isolate an individual. Chasing and harassing the group as a whole is useful in this respect because it makes vulnerable individuals easier to detect and to separate from the heard (though wolves do sometimes target individuals directly after one of the preceding stages). If no single individual is isolated, the hunt will typically transition to one of its earlier phases, particularly watching or searching. But if an individual does become isolated, it will be chased (if it flees) or harassed (if it stands its ground) until the wolves give up or are able to capture and kill it.²⁷⁵

This pattern of activity is in many respects unremarkable. We have not needed to mention any explicit forms of cooperation, coordination, or planning. And although anecdotal evidence of cooperative strategies has been documented,²⁷⁶ there is a large body of literature suggesting that, as Madden *et al.* put it, “group behavior is [...] not a well structured set of strategies but rather generalized ‘rules of thumb’ that are used to react to the prey’s escape behavior in order to minimize the risk of injury to themselves.”²⁷⁷ Some examples of these ‘rules of thumb’ (what I earlier called ‘decision rules’) might include: ‘if one individual behaves differently from the rest of the group, focus on it’; once an individual has been separated, move directly towards it until

²⁷¹ Madden *et al.* (2010, p. 1044)

²⁷² Mech *et al.* (2015, p. 9)

²⁷³ Mech (2015, p. 8)

²⁷⁴ MacNulty *et al.* (2007, p. 600)

²⁷⁵ The previous two paragraphs rely on Madden *et al.* (2010, pp. 1044-1045)

²⁷⁶ See, e.g., Mech (2007, pp. 147-149)

²⁷⁷ Madden *et al.* (2010, p. 1043)

a safe distance is reached’; ‘once a safe distance is reached, move away from the other wolves’; ‘if it is not running, harass it’; ‘if it runs, chase it’; ‘if the chase covers more than two kilometres, stop chasing’; and so on.²⁷⁸

As before, I suggest that we conceive of these rules as abstract representations of the wolves’ instincts_[i] and skills_[i]. That is, they are principles to which the wolves have a tendency (either innate or acquired) to correspond to, but *are not* — I am going to assume — explicit objects of wolf-thought: the wolves do not represent these rules to themselves.²⁷⁹

How might incentivisation enter into this picture? The sceptical doubt will be that it doesn’t, or at least not in any significant way. Rather, upon receiving certain perceptual cues (e.g., the ones that indicate moving prey) the wolves implement a certain action plan (e.g., chasing), whereas upon receiving different perceptual cues (e.g., prey standing still) they implement a different one (e.g. harassing). At best, incentivisation comes in as an initial launching point: hunger produces an incentive to hunt, but after that everything would seem to take place according to a pattern that can be described purely in terms of instincts and skills.

I think that this is too quick, however. For there is in fact an incredible range of diversity with respect to the internal complexity of different wolf-hunts. Some of this was mentioned above, but Table 1 provides a more detailed record of observed transitions between stages over many years (these data pertain to wolves hunting elk; the percentages differ depending on the prey in question).²⁸⁰

Table 1

Following State						
Preceding state	Search	Approach	Watch	Attack Group	Attack individual	Capture
Search	0%	68%	0%	31%	1%	0%
Approach	9%	0%	12%	69%	9%	1%
Watch	32%	35%	0%	27%	6%	0%
Attack Group	24%	9%	3%	13%	51%	0%
Attack Individual	16%	6%	2%	16%	8%	52%

²⁷⁸ These are adapted from Packard (2012, p. 25) and Muro *et al.* (2011, p. 193). Computer simulations that deploy rules not unlike these have been shown to model simplified hunts to a high degree of accuracy. See Madden *et al.* (2010) and Muro *et al.* (2011).

²⁷⁹ Again, this is an assumption that I am making for the sake of the example. If some of the rules are represented explicitly, then that is quite interesting.

²⁸⁰ These data are taken from MacNulty *et al.* (2007, p. 600).

As these data illustrate, various stages of the hunt can be interchanged, depending on the circumstances: Wolves do not always stop to watch their prey, and when they do, they do not necessarily proceed to attack. Similarly, even after isolating an individual and attacking it, the wolves might return to attacking the group or instead back down and continue their search. The simple decision rules just described do not account for this complexity: they record tendencies, but these tendencies are not rigid.

Given this flexibility, what determines which stage of the hunt an individual wolf will transition to in any particular case? Commenting upon this question, Madden *et al* suggest that:

In nature, what decides which transition is chosen is a combination of situational factors such as the number of wolves in the pack, the number of prey individuals in the herd, terrain features, as well as the wolf's individual attributes such as age, weight, and personality ([e.g.] aggressive individuals are more likely to move more quickly toward capture).²⁸¹

One way to read this passage (I shall consider another below) is to interpret it as offering an alternative answer to the naïve idea that it is the *wolf* that decides which stage of the hunt to enter into next. That is, we can replace the wolf in the explanation by positing something *other* than the organism, which makes her 'decision' for her, *viz.*, a vast combination of situational factors and personal attributes: if only we knew all of the variables we could predict what would happen in every instance without any recourse to the idea that the wolf chooses whether or not to pursue its prey.

As indicated in the previous chapter, I do not have a satisfactory answer to this form of explanatory reductionism, especially when it is bought with promissory notes. My thought is just that, even if it were true that, at some level of description, the wolf's movements are determined by factors such as this, the laws would have to relate descriptions cast in terms of muscle contractions, neurological events, and so on and so forth. My scepticism relates to the possibility that there is any way to translate explanations cast in this language into that which we naïvely use to describe the activity: descriptions that represent wolves as, say, attacking musk-oxen.²⁸²

Perhaps, though, the best way to resist this kind of reductionism is to investigate the idea that tends to motivate it, namely, that there is something problematic about the idea that an animal could decide which course of action to pursue.

²⁸¹ (2010, p. 1045)

²⁸² See again section §4.2. Cf. also Bermúdez (1998, p. 86).

I think that this might be conceived of as problematic for one of two reasons. The first is that the notion of deciding is sometimes taken to imply the Cartesian idea that, in so deciding, we produce a volition that in turn causes the movement.²⁸³ This is of course problematic for a host of familiar reasons that need not concern us here.²⁸⁴ But this is not the Anscombean view. For remember that according to Anscombe there is *no* ‘additional feature’ that prompts the action and thereby constitutes it as intentional. If anything, intentional and, more generally, voluntary movement is ‘marked’ by the *absence* of certain features — the absence of surprise being perhaps the most obvious.²⁸⁵ As Wittgenstein put it with respect to human action,

There is a particular interplay of movements, words, expressions of face, as of manifestations of reluctance or readiness, which are characteristic of the voluntary movements of a normal human being. If one calls a child, he does not come automatically: there is e.g. the gesture ‘I don’t want to!’ Or coming cheerfully, the decision to come, running away with signs of fear, the effects of being addressed, all the reactions of the game, the signs and the effects of consideration.²⁸⁶

Of course, the case of an animal is different in many respects; but the general point is just that the concept of voluntary movement has its place in a certain wider context that is marked by features that, I believe, we can also see in the case of many animals: the wolves stop and observe their prey, they show signs of hesitation in response to larger animals, and at times give up their pursuit of one animal and refocus their attention on others. The idea that they are making decisions is at home in this context and ought not be taken to imply anything metaphysically suspect.

This brings us, however, to the second reason why someone might find the idea that wolves make decisions problematic. For it might be said that the capacity to reason practically — to think about what to do — *belongs* to the ‘normal surroundings’ of voluntary movement and choice that Wittgenstein describes. Since it seems unlikely that the wolves utilise practical reasoning in the context of their hunts, then the appeal to the Wittgensteinian point cannot be used to *support* the idea in the context of the wolves; on the contrary, it seems to undermine it.²⁸⁷

Before I respond to this objection, I want to briefly consider another way in which we might read the passage from Madden *et al*, quoted above, in which they suggest that ‘what decides

²⁸³ As Descartes (1985, p. 343) put it, “simply by willing something [the soul] brings it about that the little gland to which it is closely joined moves in the manner required to produce the effect corresponding to this volition.”

²⁸⁴ But see Hyman (2015; esp chapters 1 and 2).

²⁸⁵ As Wittgenstein (2009; §628) famously observed.

²⁸⁶ (1987, §594)

²⁸⁷ This would appear to be the attitude of Bermúdez (2002, p. 252ff), for example.

which transition is chosen' depends on a vast range of situational factors. On this interpretation, the idea is simply that the transition from one stage of the process to another becomes more or less probable depending on a vast range of contextual and personal factors, including, for example, how aggressive the wolf in question happens to be. And when understood in this way, the passage seems to me to be undeniably correct: what the wolves do in response to which situation is not random, as these data show, and I have no doubt that the factors mentioned by Madden *et al.* bear on the wolf's adopted course of action. But this raises the question of *how* this vast amount of information influences what the wolf does.

Unsurprisingly, I think that this is where the notion of an incentive can help us. By way of analogy, consider a tennis master whose skill allows her to pick up, synthesise, and use information pertaining to the particularities of the concrete situation — information which correlates with the shot that gets played. These particularities might include background knowledge of the nature of her opponent (whether she is a baseline player, or likes to come to the net), but also subtle cues about the opponent's balance (as well as her own), speed, and so on. Clearly, the way in which these cues influence the shot that gets played is not by being consciously thought about: the master does take these nuances 'into consideration' in the sense that she first observes them and then reasons about what sort of shot would be most likely to win the point. *That* is the sort of thing that a novice might do (to adapt a quote from Ryle, the novice engages in an "internal process of avowing to himself propositions about what is to be done [...]; only then can he execute his performance in accordance with those dictates."²⁸⁸)

Still, the question remains: how does this information influence what happens (as the evidence from correlation would suggest)? My positive proposal is that it does so by incentivising a narrow range of possible shots: the opponent's position on the court incentivises a cross court forehand, to a certain extent, and a nearside drop-shot, to another; the possibility of a lob or a forehand slice, by contrast, does not cross her mind. And while it is ultimately up to the athlete to choose which shot she hits, we need not conceive of this decision as a choice made after a period of reflection. Indeed, there need be no gap between making a decision and acting on it — one can make the decision *by* so acting.²⁸⁹

²⁸⁸ Ryle (1949, pp. 29-30). Ryle is not describing the performance of a novice in this quote, but is rather mocking the view that all knowledge is propositional (which he takes to be committed to this description of expert performance). But whether or not he is right about that, his remarks make for an especially apt description of what it is like to be a novice tennis player — as I can reliably report.

²⁸⁹ This is not to deny that some shots involve no choice and are better conceived of as pure reactions. But again, my thought is that these things sit on a spectrum.

We can conceive of the wolf's relation to the particularities of her situation in an analogous way. Just as the tennis master's skills_[e] transform the subtle cues of the situation into a narrow range of incentivised options, so, too, do the wolf's instincts_[e] and skills_[e] transform the wealth of information about its surroundings into a limited range of options which it can select between. As Korsgaard remarks,

The fact that an animal has certain instincts explains why he is subject to the associated incentives. In this sense the animal's instincts play a double role in the account of his actions. They both explain why the animal is subject to certain incentives in the first place, and what he does in response to those incentives once they are present.²⁹⁰

If this is right, then it suggests a way of responding to the objection that the language of decision only has place in a context in which the subject reflects upon what to do. For this language is also at home in the context of skilled performance where reflection is noticeably *out* of place. And if the analogy between the wolf and the tennis master is on point in this respect, then there ought to be no barrier to conceiving of the wolf as deciding between available alternatives.

5.5 A theory of animal activity

I hope that these examples have given a better sense of the theoretical notions that I am proposing to adopt, and have served to illustrate the way in which they might find application in specific cases. Now, though, I want to bring the points made over the course of the last two chapters together into something resembling a more general, Anscombean theory of animal activity.

Let me begin with a brief review. I began chapter 4 with the observation that animal activities are generally complex, agential processes: animals do one thing, and then another, (etc.) because they are doing some further thing. The problem that this raised, recall, was that of accounting for the tendency towards completion of the process as a whole, given that there are no causal laws which govern the transition from one stage of the activity (so described) to the next.

In the case of the intentional actions of mature human beings, I have defended the Anscombean view that it is the agent's practical knowledge of what they are doing, why they are doing it, and how it is being done that both (i) unifies what happens into a rational order of means and ends and (ii) serves to partially explain the tendency towards completion demanded

²⁹⁰ (2009, pp. 106-107)

by telic representation. But I have also claimed that this explanation cannot be applied unproblematically to animals, and thus proposed that we classify intentional actions and animal activities as two species of a common genus.

The question that I have been pursuing in this chapter is whether a recognisably Anscombean account of the latter can be developed. My strategy for doing so has been based on two ideas. The first is that practical knowledge, like knowledge more generally, is internally related to the capacity to use a fact as a reason: an agent that is intentionally doing A can use the fact that she is doing A as a reason for (*inter alia*) doing B (where B is a means to doing A). The second idea is that the ability to use a fact as a reason can be seen to lie on a spectrum: at one end there is the sophisticated, two-way ability to take a fact ‘into consideration or account’; in the middle there is the capacity to be ‘guided’ by a fact; and at the other end is the involuntary susceptibility to a fact’s imperious ‘control’. The fact that one is in pain, I suggested, can enter into the determination of what an individual does in each of these various ways.

It was in an effort to spell out these metaphors that I introduced the notion of an incentive, along with the corresponding concepts of instinct and skill. My basic idea was that (i) incentives are intermediary states that are produced by perceptual and/or emotive experiences which in turn (ii) serve to motivate certain courses of action that are (iii) further determined by the animal’s instincts and skills and, hence, (iv) do not stand in need of support from practical reason. This explains how an animal — or person — is able to inhibit the instinctive forms of response associated with various perceptual and emotive experiences, whilst nevertheless retaining their ability to act in accordance with their instinctive tendency.

So much by way of review. The important question now is whether these notions can account for the temporal unity of animal activities in a recognisably Anscombean way. To this end, let me start by quoting Anscombe’s own defence of the idea that ‘intention concepts’ can reasonably be applied to animals:

[...] we certainly ascribe intention to animals. The reason is precisely that we describe what they do in a manner perfectly characteristic of the use of intention concepts: we describe what *further* they are doing *in* doing something (the latter description being *more* immediate, nearer to the merely physical): the cat is stalking a bird *in* crouching and slinking along with its eye fixed on the bird and its whiskers twitching. The enlarged description of what the cat is doing is not all that characterises it as an intention (for enlarged descriptions are possible of any event that has describable effects), but to this is added the cat’s perception of the bird, and what it does if it

catches it. The two features, knowledge and enlarged description, are quite characteristic of the description of intention in acting.²⁹¹

Insofar as I understand it, the central idea is that our naïve judgment that the cat is stalking the bird by crouching (and, I take it, our judgment that she is crouching like that *because* she is stalking a bird) is warranted because (i) the cat is a perceiver (a ‘knower’) and (ii) we have a rough idea of what cats tend to do when they catch birds.

I think this brief justification that Anscombe provides here is neither compelling nor illuminating. Why does a future in which the cat catches and kills a bird constitute the realisation of the tendency inherent in what the cat is doing here and now, whereas a future in which the bird escapes does not? Why think that the cat is stalking the bird in order to capture and kill it rather than, say, scare it away? If the cat were an intentional agent, we could say that she is using the fact that she is hunting the bird (and not that she is trying to scare it) as a *reason* for stalking it. But on the assumption that the cat lacks the practical knowledge which would serve to underwrite that ability, Anscombe does not leave us with a clear answer to these questions.

It is here, I think, that the tools I have marshalled can be of use to contemporary Anscombeans. For although the cat lacks the practical knowledge that would enable her to use the fact that she is hunting the bird as a reason, on my account, that fact can nevertheless enter into the determination of what happens in a similar way. For given that the cat is an animal with a suite of instincts and skills that (i) become operative once she starts hunting and (ii) determine what incentives she is, in turn, subject to, then it seems reasonable to suppose that, (iii) upon perceiving a bird, she will be incentivised to stalk, capture, and kill it. After all, cats are not in general possessed of instincts to flush birds up into the air, in the way in which some hunting dogs are; if she could be trained to do so, then perhaps that *would* be what she was up to. But assuming that she has not been so-trained, and that her instincts operate through the production of incentives, then *what* she is incentivised to do must be understood in terms of these instincts. And if I am also right that her instincts resolve the question of *how* to act in accordance with the incentive, then we can cite the fact that she is stalking the bird to explain *why* she is slinking along, inching ever closer. That is, the naïve form of rationalisation that Anscombe appeals to above (i.e., that the cat is stalking the bird in crouching and, conversely, that she is crouching along because she is stalking the bird) is genuinely explanatory precisely because the fact that she is stalking a bird entails that the instincts which explain her activity’s tendency towards completion are operative.

²⁹¹ (1963, p. 86)

If this sort of account is on the right track, then the explanation of why it will be no accident if she captures and kills the bird — why it is precisely *that* future towards which her present movements are tending — is similar, in certain respects, to the account provided in the case of intentional agents. For, just as an intentional agent's practical knowledge serves to orientate him in relation to the particularities of the here and now by determining the practical significance of his surroundings in relation to his end, an animal's instincts_[e] and skills_[e] provide her with a similar ability to take advantage of her circumstances, in this case by producing further incentives as the activity unfolds. And just as the progress of an intentional agent's deed is beholden to his ability to take the means that are available to him — something he does, I have suggested, by exercising his knowledge-how — the progression of an animal activity towards completion is similarly dependent on *its* ability to act in accordance with the dictates of its instincts_[e] and skills_[e].

On the other hand, though, there is also a crucial difference. For in the case of an agent with practical knowledge, I have argued that the perception of practically charged particulars is such that certain courses of action seem reasonable (and others unreasonable) *in light of* what he is up to. Returning to the case of the man walking up Fifth Avenue, the protest march presents itself as an obstacle to turning left on Broadway, and thereby makes the prospect of turning left unreasonable *as a means to walking to the supermarket*. That is to say, practical knowledge enables the agent to understand *why* a certain course of action is compelling in this situation.

But I have also suggested that this is not the way in which incentives present their objects. Rather than making a certain course of action seem reasonable as a means to some further end, an incentive simply directs the animal to do something: to turn to the right, or to vomit, or to chase the rabbit, and so on; *why* such courses of action are incentivised is a further question — a question that is not provided for simply by being subject to one.

The way to put the distinction is, I think, like this: in the intentional case, the fact that S is doing A influences what is done in pursuit of that end in a way that is, so to speak, transparent to the subject. Her practical knowledge allows her to articulate *why* the protest march (say) is an obstacle for her wider project. When an animal is incentivised to do B, by contrast, the animal is 'guided', if you like, by the fact that she is doing A, but the form of guidance is opaque: the cat is subject to the incentive produced by her perception of the bird partly because she is hunting, but she does not need to apprehend the instrumental relationship between stalking the bird and catching and killing it in order to *be* incentivised to stalk it.

This will perhaps become clearer if we change the example to one in which the animal's 'goal' is not within view. Think again about the wasp. She is incentivised to bring the prey to

the edge of the burrow; then she is incentivised to enter and turn around; and then she is incentivised to drag the prey in. She need not have any overarching representation of the activity's end-point: it is possible that it is guided from one situation to the next, by a chain of linked incentives. As Millikan puts it in a related context, such an animal “does not, as it were, *project* its goals. Its behaviours are controlled completely from behind by emerging environmental contingencies. It does not represent its goals as purposed future occurrences or states to which the actual accomplishments will be compared.”²⁹²

I think that the basic idea that Millikan is driving at is correct: for many animals, the representation of a distant goal does not figure in the explanation of what happens here and now. But Millikan's description of these animals as being ‘controlled completely from behind’ is offkey. That is the sort of language that is appropriate in the case of animals like frogs (at least as I have imagined them) whose instincts fully determine what they do in response to their perceptions. But as I have been wont to point out, the central idea behind the notion of an incentive is that it affords the *animal* some control over what it does in response to what. And, indeed, even in the case of animals that are not capable of ‘projecting their goals’, there is still room for choice. As the example of the wasp illustrates, she may give up and fly away, or drag the prey in from another location, or, indeed, keep repeating the process.

Indeed, this point is even clearer when the animal must select between competing incentives. For example, the wolves discussed previously have instinctive tendencies to stop chasing prey after (roughly) two kilometres, as well as instinctive tendencies to capture and kill prey that runs. And while the former often wins out — that's why it's a tendency — there is also evidence of wolves pursuing prey more than twenty kilometres.²⁹³ Part of the point of bringing in the notion of incentivisation is that it can help us make sense of this variability: it allows the wolves to resist their instinctive tendencies, if they so choose, and to act otherwise than we might expect.

Now my suggestion is emphatically *not* that animal activities are *in general* structured by instincts that only produce incentives in response to ‘emerging environmental contingencies’. In more advanced animals, memory, imagination, and perhaps even conclusions drawn as the result of practical inferences can serve as inputs to their instincts and skills. Such animals can become incentivised to do something *by* doing some further thing and *in order to* do something else: they can project their goals and compare their accomplishments in light of them. As indicated in the case of the wolves, for example, the direction of their search is influenced by

²⁹² (2004, pp. 169-170)

²⁹³ Mech *et al.* (2015, pp. 3-4)

past experiences of success and failure. And it seems within the realms of possibility to imagine that this memory is integrated with their current state of hunger to produce an incentive to hunt *by* searching these past locations.

The point is speculative with respect to the wolves, but future planning by apes and corvids has been persuasively documented. For example, a chimpanzee that learns to make use of a straw to extract the juice from a fruit smoothie will select that tool when given a choice between various objects— including a tempting grape — even if the reward is not made available for over an hour. And, notably, they will *not* select the tool if they already have one — in that case, they invariably choose the grape. Though there is room for discussion about the correct interpretation of these results, at the very least it suggests that some animals are capable of representing distant goals and then coordinating what further they do around this basis: they select the tool for the sake of drinking the smoothie.²⁹⁴

If I had more time, I would like to be able to offer a more detailed description of these cognitive powers and to show how they can be integrated into the theory that I have provided. But in lieu of that, let me make two points. The first is that there is nothing in what I have said that would block their integration; indeed, my hope is to have laid the groundwork out of which these more advanced cognitive powers can be seen as having arisen. The second, and to my mind more important point, however, is that possession of these further powers is not essential to our being able to see a phenomenon as a manifestation of animal agency; *that* is secured by the animal's ability to coordinate its perceptions and actions in the context of a wider pattern of activity — an ability that is underwritten, I have argued, by its instincts and skills.

5.6 Summary

In this chapter I have made a case for an Anscombean approach to animal agency. To do so, I have relied on the idea that a fact about what an agent is doing can enter into the determination of what further it does in ways that are recognisably similar to, but nevertheless distinct from, the way in which someone with practical knowledge uses the fact that she is doing something as a reason for taking the means available to her. My suggestion has been that the instincts which become operative as a result of the fact that an animal is doing something can guide (or, in simpler animals, control) what further the animal does, thereby furthering the progress of the activity towards completion.

²⁹⁴ This example is drawn from Osvath and Osvath (2008); Cf. Suddendorf *et al.* (2009) and Osvath (2010). For similar findings with respect to corvids, see Kabaday *et al.* (2017).

As I went on to point out, this account makes the influence of these facts indirect: they are not taken up directly as objects of knowledge and then *used* in such a way that the agent understands *why* they are being influenced as they are. Nevertheless, I have suggested that even here there is much room for selection between various courses of action and, as such, highly flexible patterns of activity (as exemplified, I think, by the wolves).

I have also admitted, however, that many animals *do* seem capable of projecting their goals in ways that my account cannot explain in detail: I have not offered a description of animal memory, or imagination, and the account of practical reasoning offered at the end of chapter 2 was speculative. Still, there is no reason to think that these further powers cannot be integrated into the account that I have provided.

References

- Alvarez, M., & Hyman, J. (1998). Agents and Their Actions. *Philosophy*, 73(284), 219-245.
- Andrews, K. (2015). *The Animal Mind*. London: Routledge.
- Annas, J. (1976). Davidson and Anscombe on The Same Action. *Mind*, 85(338), 251-257.
- Anscombe, G. E. M. (1963). *Intention* (2nd ed.). Cambridge, Mass: Harvard University Press.
- Anscombe, G. E. M. (1975). The First Person *Metaphysics and the Philosophy of Mind* (pp. 21-36). Oxford: Basil Blackwell.
- Anscombe, G. E. M. (1979). Under a Description. *Nous*, 13(2), 219-233.
- Anscombe, G. E. M. (1983). The Causation of Action. In M. Geach & L. Gormalley (Eds.), *Human Life, Action and Ethics: Essays by G.E.M. Anscombe* (pp. 89-108). Exeter: Imprint Academic.
- Anscombe, G. E. M. (1989). Practical Inference. In M. Geach & L. Gormally (Eds.), *Human Life, Action, and Ethics* (pp. 109-147). Exeter: Imprint Academic.
- Anscombe, G. E. M. (1993). Practical Truth. In M. Geach & L. Gormally (Eds.), *Human Life, Action and Ethics: Essays by G.E.M. Anscombe*. Exeter: Imprint Academic.
- Aristotle. (1984a). *De Anima* (J. A. Smith, Trans. Vol. One). Princeton: Princeton University Press.
- Aristotle. (1984b). Movement of Animals (A. S. L. Farquharson, Trans.). In J. Barnes (Ed.), *The Complete Works of Aristotle* (Vol. 1, pp. 1087-1096). Princeton: Princeton University Press.
- Aristotle. (2016). *De Anima* (C. Shields, Trans.). In L. Judson (Ed.), *Clarendon Aristotle Series*. Oxford: Oxford University Press.
- Baier, A. (1970). Act and Intent. *The Journal of Philosophy*, 67(19), 648-658.
- Baier, A. (1977). The Intentionality of Intentions. *The Review of Metaphysics*, 30(3), 389-414.
- Barner, D., Wood, J., Hauser, M., & Carey, S. (2008). Evidence for a non-linguistic distinction between singular and plural sets in rhesus monkeys *Cognition*, 107, 603-622.
- Bennett, J. (1976). *Linguistic Behaviour*. Cambridge: Cambridge University Press.
- Bennett, M. R., & Hacker, P. M. S. (2003). *The Philosophical Foundations of Neuroscience*. Oxford: Blackwell.
- Bermúdez, J. L. (1998). *The Paradox of Self-Consciousness*. Cambridge, MA: The MIT Press.
- Bermúdez, J. L. (2002). Rationality and Psychological Explanation Without Language. In J. L. Bermúdez & A. Millar (Eds.), *Reason and Nature* (pp. 233-265). Oxford: Clarendon.
- Bermúdez, J. L. (2003). *Thinking Without Words*. Oxford: Oxford University Press.
- Boyle, M. (2012). Essentially Rational Animals. In A. Günter & C. James (Eds.), *Rethinking Epistemology* (Vol. 2, pp. 395-428). Berlin: De Gruyter.
- Boyle, M. (2016). Additive Theories of Rationality: A Critique. *European Journal of Philosophy*, 24(3), 527-555.
- Boyle, M., & Lavin, D. (2010). Goodness and Desire. In S. Tanenbaum (Ed.), *Desire, Practical Reason, and the Good* (pp. 161-201). Oxford: Oxford University Press.
- Boysen, S. T., & Bernston, G. G. (1995). Responses to Quantity: Perceptual and Cognitive Mechanisms in Chimpanzees. *Journal of Experimental Psychology and Animal Behavioural Processes*, 21, 82-86.
- Boysen, S. T., Mukobi, K. L., & Bernston, G. G. (1999). Overcoming response bias using representations of number by chimpanzees. *Animal Learning and Behaviour*, 27(2).
- Boysen, S. T. M., K.L.Bernston, G.G. (1999). Overcoming response bias using representations of number by chimpanzees. *Animal Learning and Behaviour*, 27(2).
- Brandom, R. (1994). *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Bridges, J. (2006). Davidson's Transcendental Externalism. *Philosophy and Phenomenological Research*, LXXIII(2), 290-315.
- Brockman, J. (1980). The Control of Nest Depth in a Digger Wasp. *Animal Behaviour*, 28, 426-445.
- Brockman, J. (1985). Provisioning Behavior of the Great Golden Digger Wasp, *Sphex ichneumoneus* (L.) (Sphecidae) *Journal of the Kansas Entomological Society* 58(4), 631-655.
- Buckner, C. (2011). Two Approaches to the Distinction Between Cognition and 'Mere Association'. *International Journal of Comparative Psychology*, 24, 314-348.
- Burge, T. (2009). Primitive Agency and Natural Norms. *Philosophy and Phenomenological Research*, LXXIX, 251-278.

- Burge, T. (2010a). *Origins of objectivity*. Oxford: Oxford University Press.
- Burge, T. (2010b). Steps Towards Origins of Propositional Thought. *Disputatio*, IV(29), 39-67.
- Burge, T. (2011). Border crossings: Perceptual and post-perceptual object representation. *Behavioral and Brain Sciences*, 34, 125.
- Byrne, R. (1995). *The Thinking Ape*. Oxford: Oxford University Press.
- Böhm, J., Scherzer, S., & all, e. (2016). The Venus Flytrap *Dionaea muscipula* Counts Prey-Induced Action Potentials to Induce Sodium Uptake. *Current Biology*, 26, 286-295.
- Campbell, L. (2015). *Action, Intention, and Knowledge*. (PhD), University of Cambridge, Cambridge.
- Campbell, L. (2017). An epistemology for practical knowledge. *Canadian Journal of Philosophy*, 1-19. doi: 10.1080/00455091.2017.1341073
- Campbell, L. (2018). Propositionalism About Intention: Shifting the Burden of Proof. *Canadian Journal of Philosophy*, 1-33. doi: <https://doi.org/10.1080/00455091.2018.1512819>
- Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.
- Carey, S. (2011). Précis of *The Origin of Concepts*. *Behavioral and Brain Sciences*, 34, 113-167.
- Carruthers, P. (2004). On Being Simple Minded. *American Philosophical Quarterly*, 41(3), 205-220.
- Cash, L. (2017). Wittgenstein and the Animal Origins of Linguistic Communication. *Philosophical Investigations*, 40(4), 303-328.
- Clayton, N., Emery, N., & Dickinson, A. (2006). The rationality of animal memory: Complex caching strategies of western scrub jays In S. Hurley & M. Nudds (Eds.), *Rational Animals?* (pp. 197-216). Oxford: Oxford University Press.
- Comrie, B. (1976). *Aspect*. Cambridge: Cambridge University Press.
- Crane, T. (2002). *The Mechanical Mind* (2nd ed.). Abingdon: Routledge.
- Davidson. (1999). Interpretation: Hard in Theory, Easy in Practice. In M. d. Caro (Ed.), *Interpretations and Causes: New Perspectives on Donald Davidson's Philosophy* (pp. 31-44). London: Springer.
- Davidson, D. (1963). Actions, Reasons and Causes *Essays on Actions and Events* (pp. 3-19). Oxford: Clarendon.
- Davidson, D. (1967). Causal Relations. *Journal of Philosophy*, 64(21), 691-703.
- Davidson, D. (1970). Mental Events. In L. Foster & J. W. Swanson (Eds.), *Essays on Action and Events* (pp. 207-224). Oxford: Clarendon Press.
- Davidson, D. (1971). Agency *Essays on Actions and Events* (pp. 43-61). Oxford: Clarendon.
- Davidson, D. (1973). Freedom to Act *Essays on Action and Events* (pp. 63-81). Oxford: Clarendon.
- Davidson, D. (1978). Intending *Essays on Action and Events* (pp. 83-102). Oxford: Clarendon.
- Davidson, D. (1982). Rational Animals. *Dialectica*, 36(4), 317-327.
- Davidson, D. (1987a). Knowing One's Own Mind. *Proceedings and Addresses of the American Philosophical Association*(3), 441-458.
- Davidson, D. (1987b). Problems in the Explanation of Action *Problems of Rationality* (pp. 101-116). Oxford: Oxford University Press.
- Davidson, D. (1991). Epistemology Externalized. *Dialectica*, 45(2-3), 191-202.
- Davidson, D. (1992). The Second Person. *Midwest Studies in Philosophy*, XVII, 255-267.
- Davidson, D. (1994). Reply to Stoecker. In R. Stoecker (Ed.), *Reflecting Davidson* (pp. 286-289). Berlin: Walter de Gruyter.
- Davidson, D. (1995). Donald Davidson. In S. Guttenplan (Ed.), *A Companion to the Philosophy of Mind* (pp. 232-244). London: Wiley-Blackwell.
- Davidson, D. (1997). Seeing Through Language. In J. M. Preston (Ed.), *Thought and Language* (pp. 15-27). Cambridge: Cambridge University Press.
- Davidson, D. (1999a). Interpretation: Hard in Theory, Easy in Practice. In M. d. Caro (Ed.), *Interpretations and Causes* (pp. 31-44). New York: Springer.
- Davidson, D. (1999b). The Emergence of Thought. *Erkenntnis*, 55(1), 7-17.
- Davidson, D. (2001). What Thought Requires *Problems of Rationality* (pp. 135-150). Oxford: Oxford University Press.
- Davidson, D. (2003). Quine's Externalism. *Grazer Philosophische Studien*, 66, 281-297.
- Dennett, D. (1968). Features of Intentional Action. *Philosophy and Phenomenological Research*, 29(2), 232-244.
- Dennett, D. (1973). Mechanism and Responsibility. In T. Honderich (Ed.), *Essays on Freedom of Action* (pp. 157-184). London: Routledge and Kegan Paul.

- Dennett, D. (2009). Intentional Systems Theory. In A. Beckermann, B. McLaughlin & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind* (pp. 339-350). Oxford: Oxford University Press.
- Dennett, D. (2013). The Evolution of Reasons. In B. Bashour & H. Muller (Eds.), *Contemporary Philosophical Naturalism and its Implications* (pp. 47-81). London: Routledge.
- Denton, D. A., McKinley, M. J., & Farrell, M. E., G.F. (2009). The role of primordial emotions in the evolutionary origin of consciousness. *Consciousness and Cognition*, 18, 500-514.
- Descartes, R. (1985). *The Philosophical Writings of Descartes* (J. Cottingham, R. Stoothoff & D. Murdoch. Eds. Vol. 1 & 2). Cambridge: Cambridge University Press.
- Descartes, R. (1991). *The Philosophical Writings of Descartes* (J. Cottingham, R. Stoothoff, D. Murdoch & A. Kenny Eds. Vol. III). Cambridge: Cambridge University Press.
- Dickinson, A. (1994). Instrumental Conditioning. In N. J. Mackintosh (Ed.), *Animal learning and cognition* (pp. 45-79). San Diego: Academic Press.
- Dings, R. (2018). Understanding phenomenological differences in how affordances solicit action. *Phenomenological Cognitive Science*, 17, 681-699.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.
- Falvey, K. (2000). Knowledge in Intention. *Philosophical Studies*, 99(1), 21-44.
- Fernandez, P. (2014). Reasoning and the Unity of Aristotle's Account of Animal Motion. In B. Inwood (Ed.), *Oxford Studies in Ancient Philosophy* (Vol. 47, pp. 151-203). Oxford: Oxford University Press.
- Finkelstein, D. (2003). *Expression and the Inner*. Cambridge, Mass: Harvard University Press.
- Finkelstein, D. (2007). Holism and Animal Minds. In A. Cray (Ed.), *Wittgenstein and the Moral Life: Essays in Honor of Cora Diamond* (pp. 251-278). Cambridge, Ma: MIT University Press.
- Fogel, D. B. (2000). Imagining machines with imagination. *Proceedings of the IEEE*, 88, 284-288. *Proceedings of the IEEE*, 88, 284-288.
- Foot, P. (2001). *Natural goodness*. Oxford: Oxford University Press.
- Ford, A. (2008). *Action and Generality*. (PhD), Pittsburgh.
- Ford, A. (2011). Action and Generality. In A. Ford, J. Hornsby & F. Stoutland (Eds.), *Essays on Anscombe's Intention* (pp. 76-104). Cambridge, Ma: Harvard University Press.
- Ford, A. (2015). The Arithmetic of Intention. *American Philosophical Quarterly*, 52(2), 129-143.
- Ford, A. (2016). On What is in Front of Your Nose. *Philosophical Topics*, 44(1).
- Frankfurt, H. (1978). The Problem of Action. *American Philosophical Quarterly*, 15(2), 157-162.
- Frege, G. (1953). *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number* (J. L. Austin, Trans.). Oxford: Blackwell.
- Frege, G. (1956). The thought: a logical inquiry. *Mind*, 65(259), 289-311.
- Frege, G. (1980). Letter from Frege to Peano 29 Sep. 1896. In G. Gabriel (Ed.), *Philosophical and Mathematical Correspondence* (pp. 115). Oxford: Blackwell.
- Galton, A. (1984). *The Logic of Aspect: An Axiomatic Approach*. Oxford: Clarendon Press.
- Galton, A. (2012). *States, Processes and Events, and the Ontology of Causal Relations*. Paper presented at the Seventh International Conference on Formal Ontology in Information Systems, Graz, Austria.
- Geach, P. T. (2000). Intention, Freedom, and Predictability. In R. Teichmann (Ed.), *Logic, Cause, and Action* (pp. 73-81). Cambridge: Cambridge University Press.
- Glock, H. J. (2000). Animals, Thoughts and Concepts. *Synthese*, 123, 35-64.
- Glock, H. J. (2012). Non-Human Knowledge and Non-Human Agency. In S. Tolksdorf (Ed.), *Conceptions of Knowledge* (pp. 557-587). New York: de Gruyter.
- Glock, H. J. (2013). Animal Minds: a Non-Representationalist Approach. *American Philosophical Quarterly*, 20(3), 213-232.
- Gould, J. L., & Gould, C. J. (1998). Reasoning in animals. *Scientific American*, 9, 52-59.
- Hanser, M. (1998). Intention and Teleology. *Mind*, 107(426), 381-401.
- Heck, R. (2007). Are there different kinds of content? In B. P. McLaughlin & J. D. Cohen (Eds.), *Contemporary Debates in Philosophy of Mind* (pp. 117-138). Oxford: Blackwell.
- Heyes, C., & Dickinson, A. (1990). The Intentionality of Animal Action. *Mind and Language*, 5(1), 87-104.
- Hofstadter, D. (1985). On the seeming paradox of mechanizing creativity *Metamagical Themas* (pp. 526-546). London: Penguin.
- Holland, P. C. (1979). Differential Effects of Omission Contingencies on Various Components of Pavlovian Appetitive Conditioning. *Journal of Experimental Psychology: Animal Behavior Processes* 5, 178-193.

- Hornsby, J. (1999). Anamalousness in Action. In L. Hahn (Ed.), *The Philosophy of Donald Davidson* (pp. 623-636). Chicago: Open Court.
- Hornsby, J. (2016). Intending, knowing-how, infinitives. *Canadian Journal of Philosophy*, 46(1), 1-17.
- Humphreys, G. W., & Riddoch, M. J. (Eds.). (2005). *Attention in Action*. New York: Psychology Press.
- Hurley, S. (2001). Perception and Action: Alternative Views. *Synthese*, 129(1), 3-40.
- Hurley, S. (2003). Animal Action in the Space of Reasons. *Mind & Language*, 18(3), 231-256.
- Hursthouse, R. (2000). Intention. *Royal Institute of Philosophy Supplement*, 46, 83-105.
- Hyman, J. (1999). How Knowledge Works. *The Philosophical Quarterly*, 49(197), 433-451.
- Hyman, J. (2015). *Action, Knowledge, and Will*. Oxford: Oxford University Press.
- Ingle, D. (1983). Prey selection by frogs and toads. In E. Satinoff & P. Teitelbaum (Eds.), *Motivation* (pp. 235-261).
- Ingle, D. (1998). Perceptual constancies in lower vertebrates. In V. Walsh & J. Kulikowski (Eds.), *Perceptual Constancy* (pp. 173-191).
- James, W. (1890). *The Principles of Psychology*. New York: Henry Holt and Company.
- Jamieson, D. (2009). What do animals think? In R. Lurz (Ed.), *The Philosophy of Animal Minds* (pp. 15-34). Cambridge: Cambridge University Press.
- Kabaday, C. T., L.A., von Bayern, A. M. P., & Osvath, M. (2017). Ravens, New Caledonian crows and jackdaws parallel great apes in motor self-regulation despite smaller brains *Royal Society Open Science*, 3(160104), 1-7.
- Keijzer, F. (2013). The SpheX story: How the cognitive sciences kept repeating an old and questionable anecdote *Philosophical Psychology*, 26(4), 502-519.
- Keil, G. (2007). What Do Deviant Causal Chains Deviate From? In C. Lumer & S. Nannini (Eds.), *Intentionality, Deliberation and Autonomy* (pp. 69-90). London: Routledge.
- Korsgaard, C. (2009). *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.
- Krifka, M. (1998). The Origins of Telicity. In S. Rotstein (Ed.), *Events and Grammar* (pp. 197-235). Dordrecht: Kluwer.
- Lavin, D. (2013). Must There Be Basic Action? *Noûs*, 47(2), 273-301.
- Lavin, D. (2015). Action as a form of temporal unity: on Anscombe's Intention. *Canadian Journal of Philosophy*, 45(5-6), 609-629.
- Luncz, L. V., & Boesch, C. (2015). The Extent of Cultural Variation Between Adjacent Chimpanzee (Pan troglodytes verus) Communities; A Microecological Approach. *American Journal of Physical Anthropology*, 156, 67-75.
- MacNulty, D. R., Mech, L. D., & Smith, D. W. (2007). A Proposed Ethogram of Large-Carnivore Predatory Behavior, Exemplified by the Wolf. *Journal of Mammalogy*, 88(3), 595-605.
- Madden, J. D., Arkin, R. C., & MacNulty, D. R. (2010). Multi-robot system based on model of wolf hunting behavior to emulate wolf and elk interactions *Proceedings of the 2010 IEEE International Conference on Robotics and Biomimetics* 1043-1050.
- Marcus, E. (2012). *Rational Causation*. Cambridge, MA: Harvard University Press.
- McDonough, R. (1989). Towards a non-mechanistic theory of meaning. *Mind*, xcvi(389).
- McDowell, J. (1989). One Strand in the Private Language Argument. *Grazer Philosophische Studien* 33.
- McDowell, J. (2011). *Some Remarks on Intention in Action*. Paper presented at the The Amherst Lecture in Philosophy.
- Mech, L. D., Smith, D., & Macnulty, D. (2015). *Wolves on the Hunt*. Chicago: The University of Chicago Press.
- Millikan, R. G. (1995). Pushmi-Pullyu Representations. *Philosophical Perspectives*, 9, 185-200.
- Millikan, R. G. (2004). *Varieties of Meaning*. Cambridge, MA: The MIT Press.
- Milner, A. D., & Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.
- Mody, S., & Carey, S. (2016). The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition*, 154, 40-48.
- Moran, R. (2004). Anscombe on 'Practical Knowledge'. In H. Steward & J. Hyman (Eds.), *Agency and Action* (pp. 43-68). Cambridge: Cambridge University Press.
- Moran, R., & Stone, M. (2011). Anscombe on Expression of Intention: An Exegesis. In F. Stoutland, J. Hornsby & A. Ford (Eds.), *Essays on Anscombe's Intention* (pp. 33-75). Cambridge, MA: Harvard University Press.
- Mourelatos, A. (1978). Events, Processes, and States. *Linguistics and Philosophy*, 2, 415-434.

- Muro, C., Escobedo, R., Spector, L., & Coppinger, R. P. (2011). Wolf-pack (*Canis lupus*) hunting strategies emerge from simple rules in computational simulations. *Behavioural Processes*, 88, 192-197.
- Neander, K. (2006). Content for Cognitive Science. In G. Macdonald & D. Papineau (Eds.), *Teleosemantics* (pp. 195-207). Oxford: Oxford University Press.
- Osvath, M. (2010). Great ape foresight is looking great. *Animal Cognition*, 13, 777-781. s
- Osvath, M., & Osvath, H. (2008). Chimpanzee (*Pan troglodytes*) and orangutan (*Pongo abelii*) forethought: self-control and pre-experience in the face of future tool use. *Animal Cognition*, 11, 661-674.
- Packard, J. (2012). Wolf Social Intelligence. In A. P. Maia & H. F. Crussi (Eds.), *Wolves: Biology, Behavior, and Conservation* (pp. 1-47). Hauppauge: Nova Science Publishers.
- Papineau, D. (2003). Is Representation Rife? *Ratio*, 16(2), 107-123.
- Peckham, G. W., & Peckham, E. G. (1905). *Wasps, Social and Solitary*. Westminster, UK: Constable.
- Prior, A. N. (1949). Determinables, Determinates, and Determinants. *Mind, New Series*, 58(229), 1-20.
- Putnam, H. (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 131-193.
- Ratcliffe, M. (2015). *Experiences of Depression*. Oxford: Oxford University Press.
- Rescorla, R. A., & Skucy, J. C. (1969). Effect of response-independent reinforcers during extinction. *Journal of Comparative and Physiological Psychology*, 67, 381-389.
- Rizzolatti, G., Gentilucci, M., & Matelli, M. (1985). Selective spatial attention: one centre, one circuit, or many circuits? In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 251-265). Hillsdale, NJ: Erlbaum.
- Rizzolatti, G., Riggio, L., & Sheila, B. M. (1994). Space and selective attention. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV. Conscious and nonconscious information processing* (pp. 231-265). Cambridge, MA: The MIT Press.
- Rorty, R. (1981). *Philosophy and the Mirror of Nature*. Oxford: Blackwell.
- Rose, J. D., Arlinghaus, R., S J Cooke, S. J., Diggles, B. K., Sawynok, W., Stevens, E. D., & Wynne, C. D. L. (2014). Can fish really feel pain? *Fish and Fisheries*, 15, 97-133.
- Rothstein, S. (2004). *Structuring Events: A Study in the Semantics of Lexical Aspect*. Oxford: Wiley-Blackwell.
- Ryle, G. (1949). *The Concept of Mind*. New York: Routledge.
- Rödl, S. (2002). Practice and the Unity of Action. In G. Meggle (Ed.), *Social Facts and Collective Intentionality*. Frankfurt: Frankfurt: Hänsel-Hohenhausen
- Saidel, E. (2009). Attributing Mental Representations to Animals. In R. Lurz (Ed.), *The Philosophy of Animal Minds* (pp. 35-51). Cambridge: Cambridge University Press.
- Schwenkler, J. (2015). Understanding Practical Knowledge. *Philosopher's Imprint*, 15(15), 1-32.
- Setiya, K. (2003). Explaining Action. *The Philosophical Review*, 112(3), 339-393.
- Setiya, K. (2008). Practical Knowledge. *Ethics*, 118(3), 388-409.
- Setiya, K. (2012). Knowing How. *Proceedings of the Aristotelian Society, CXII*, 285-307.
- Sinhababu, N. (2015). Advantages of Propositionalism. *Pacific Philosophical Quarterly*, 96, 165-180.
- Small, W. (2012). Practical Knowledge and the Structure of Action. In G. Abel & J. Conant (Eds.), *Rethinking Epistemology* (Vol. 2, pp. 133-228). Berlin: De Gruyter.
- Small, W. (2014). The Transmission of Skill. *Philosophical Topics*, 42(1), 85-111.
- Smith, M. (2004). The Structure of Orthonomy. *Royal Institute of Philosophy Supplement*, 55, 165-193.
- Stanley, J., & Williamson, T. (2001). Knowing How. *Journal of Philosophy*, 98(8), 411-444.
- Stanley, J., & Williamson, T. (2017). Skill. *Noûs*, 51(4), 713-726.
- Steward, H. (2012). *A Metaphysics for Freedom*. Oxford: Oxford University Press.
- Stoutland, F. (2011). Interpreting Davidson on Intentional Action. In J. Malpas (Ed.), *Dialogues with Davidson: Acting, Interpreting, Understanding*. Cambridge, MA: The MIT Press.
- Strawson, P. (1992). *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.
- Suddendorf, T. C., M.C.Collier-Baker, E. (2009). How great is great ape foresight? *Animal Cognition*, 12, 751-754.
- Szabó, Z. G. (2004). On the Progressive and the Perfective. *Noûs*, 38(1), 29-59.
- Taylor, C. (1964). *The Explanation of Behaviour*. London: Routledge.
- Thompson, M. (2004). What is it to wrong someone? A puzzle about justice. In P. P. R. Jay Wallace, S. Scheffler, M. Smith (Ed.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (pp. 333-384). Oxford: Clarendon.

- Thompson, M. (2008). *Life and Action: Elementary Structures of Practice and Practical Thought*. Cambridge, Ma: Harvard University Press.
- Tinbergen, N. (1969). *The Study of Instinct*. Oxford: Oxford University Press.
- Travis, C. (2004). The Silence of the Senses. *Mind*, 113(449), 57-94.
- Velleman, J. D. (1989). *Practical Reflection*. Princeton: Princeton University Press.
- Velleman, J. D. (1992). What Happens When Someone Acts? *Mind*, 101(403), 461-481.
- Velleman, J. D. (2004). Précis of The Possibility of Practical Reason. *Philosophical Studies*, 121(3), 225-238.
- Vendler, Z. (1957). Verbs and Times *Linguistics in Philosophy* (pp. 97-121). Ithaca, NY: Cornell University Press.
- Vogler, C. (2002). *Reasonably Vicious*. Cambridge, Mass: Harvard University Press.
- Vogler, C. (2016). Nothing Added: *Intention* §§19 and 20. *American Catholic Philosophical Quarterly*, 90(2), 229-247.
- Wallace, D. F. (2004). Consider the Lobster. *Gourmet*, 50-64.
- Williams, G. (1987). Oblique Intention. *The Cambridge Law Journal*, 46(3), 417-438.
- Wittgenstein, L. (1969). *The Blue and Brown books*. Oxford: Basil Blackwell.
- Wittgenstein, L. (1987). *Zettel* (G. E. M. Anscombe, Trans. G. E. M. Anscombe & G. H. v. Wright Eds.). Oxford: Basil Blackwell.
- Wittgenstein, L. (2009). *Philosophical investigations* (G. E. M. Anscombe, P. M. S. Hacker & J. Schulte, Trans. Rev. 4th ed.). Oxford: Wiley-Blackwell.
- Woodfield, A. (1982). Preface. In A. Woodfield (Ed.), *Thought and Object*. Oxford: Clarendon Press.