

Cultivating Ethical Expertise

By

Marcus Lee

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

August 2019

CONTENTS

Page:

Acknowledgements	v
Abstract	vii

Chapters:

1. Introduction	
1. The ancient skill model of ethical expertise	1
2. The contemporary skill model of ethical expertise	4
3. A constraint and a desideratum	5
4. Summary of my argument	8
2. Annas' account of practical expertise	
1. Introduction	11
2. The right kind of skills	12
3. Intellectualism about practical expertise	16
4. Acting-in-flow	19
5. Conclusion	22
3. The articulation requirement	
1. Introduction	25
2. The Dreyfus model of practical expertise	29
3. A four-stage model of practical expertise	33
4. Expertise in recreational activities	39
5. Expertise in vocational activities	41
6. Conclusion	46
4. The special requirement	
1. Introduction	49
2. Kinds of expertise	54
3. Theoretical understanding	58

4.	Practical understanding	68
5.	The potential value of understanding	76
6.	Conclusion	80
5.	The phenomenology of expert action	
1.	Introduction	83
2.	The standard view of action	84
3.	Merleau-Ponty's view of action	89
4.	Ways of being aware	93
5.	Attentive engagement	98
6.	Awareness, know-how and practical understanding	100
7.	Conclusion	104
6.	The skill model of virtue	
1.	Introduction	107
2.	Hills' arguments for intellectualism	111
3.	A tension in Annas' account	119
4.	A commitment to goodness	125
5.	Conclusion	127
7.	The skill model of <i>wu-wei</i>	
1.	Introduction	131
2.	<i>Wu-wei</i> as acting without intentions	132
3.	A familiar analogy	135
4.	A commitment to <i>Dao</i>	139
5.	Conclusion	142
8.	Conclusion	145
	References	147

Acknowledgements

I've had lots of help during this project. I am very grateful to Midlands4Cities DTP and the Arts and Humanities Research Council for generously funding this PhD. I especially want to thank my supervisors, Neil Sinclair, for his patient guidance throughout all my postgraduate studies, and Chris Woodard, for his advice and encouragement, and taking the time to provide detailed comments on previous drafts of this thesis. I also want to thank Komarine Romdenh-Romluc for introducing me to Merleau-Ponty's ideas during my MA, and thereby laying the foundation for this investigation.

My sincere thanks also go to Jules Holroyd, Uri Leibowitz, Steve Barker, Ian James Kidd, Simon Thunder, Joseph Adams, Tom Crawley, Caterina Moruzzi, James Chamberlain, Katrina O'Keefe, Kipros Lofitis, Alice Monypenny, Ozun Cetinkaya, Matthew Duncombe, Craig French, Jon Robson, Aness Webster, Zach Hoskins, Andy Fisher, Lina Jansson, Jonathan Tallant, Jonathan Birch, and many others, for their advice, encouragement, and many thought-provoking conversations.

I also want to thank my mum, dad, sister and brother, and Si, Dave, and Pete for their love and support. Above all, I want to thank Claudia, whose love and understanding sustain me always.

Abstract

The skill model of ethical expertise holds that becoming an ethical expert is (like) becoming an expert in a practical skill. This model of ethical education was advocated by philosophers in both ancient Greece and ancient China. In this dissertation, I critique a prominent contemporary account of the skill model of virtue by attending to the phenomenology of the learning process involved in acquiring a practical skill. Situating this critique within a Merleau-Pontyan framework, I then develop a two-tiered account of ‘awareness’ which I use to explicate novel views of the epistemology of both virtue and the ancient Chinese ethical ideal *wu-wei*.

1. Introduction

1. The ancient skill model of ethical expertise

The skill model of ethical expertise holds that becoming an ideal ethical agent is (like) becoming an expert in a practical skill. This idea goes back to ancient times. For instance, Aristotle writes that:

‘[T]he way we learn the things we should do, knowing how to do them, is by doing them. For example people become builders by building, and cithara-players by playing the cithara; so too, then, we become just by doing just things, moderate by doing moderate things, and courageous by doing courageous things.’¹

According to this passage, the way we cultivate moral virtues like justice, moderation and courage is the same way we cultivate practical skills like building and playing a musical instrument, namely, by doing them. Thus, for Aristotle virtue is like practical expertise in that both are acquired by practising the very thing we are trying to achieve.

The early Daoist² thinker Zhuangzi also appears to have thought that the ideal agent is similar to a practical expert.³ This is best seen in a story about a nobleman who witnesses expert butchery skill (Z3). In the story, Zhuangzi describes Lord Wenhui witnessing Cook Ding deftly carve an ox as if he were ‘dancing or conducting an orchestra in rhythm’. Cook Ding tells Lord Wenhui that it took him three years of practise before he came to see the ox as more than just meat, and that because he now ‘go[es] along with the natural makeup’ of the ox and ‘follow[s] things as they are’, his knife never touches bone, so he hasn’t had to sharpen it in

¹ *Aristotle: Nicomachean Ethics* (Broadie & Rowe 2002), (hereafter *NE*) 1103a30 – 1103b5. A cithara is a musical instrument.

² By ‘early Daoism’ I am referring to two texts: *The Complete Works of Zhuangzi* (Watson 2013) and *The Daodejing of Laozi* (Ivanhoe 2002). I will refer to portions of each text by the initial letter of the eponymous author’s name (Zhuangzi and Laozi, respectively) and the number of the chapter the text appears in, e.g., Chapter 1 of the *Daodejing* would be (L1).

³ Eno (1996), Yearley (1996), Slingerland (2003), Barrett (2011).

over a decade of use. To this Lord Wenhui replies that Cook Ding has taught him 'how to care for life'. From this and other stories of practical expertise (Z19), it seems appropriate to interpret Zhuangzi as suggesting that the cultivation and exercise of an expert skill (like butchery) is analogous in some way to the cultivation and exercise of the ideal *wu-wei*, often translated as 'non-doing' or 'no-action'.⁴ In early Daoism, *wu-wei* is presented as the way sages comport themselves that manifests in natural spontaneity and effortless efficacy.

There is reason to think, then, that both Aristotle and Zhuangzi advocated some version of the skill model of ethical expertise. I think this is a remarkable fact that warrants further investigation. However, although parallels have frequently been drawn between Aristotle and Kongzi (Confucius), there has been very little investigation into any points of similarity between Zhuangzi's ethical perspective and Aristotle's ethics.⁵ If Zhuangzi and Aristotle did both advocate a skill model of ethical expertise, then gaining a better understanding of the cultivation and exercise of practical skills will shed valuable light on how these thinkers conceived of *wu-wei* and virtue, respectively, as well as the similarities and differences between their conceptions of ideal agency. Of course, even if Zhuangzi and Aristotle did both advocate a version of the skill model of ethical expertise, they may have been wrong. It may be that neither *wu-wei* nor virtue is best modelled along the lines of expertise in activities like playing a musical instrument, building a house, or butchering an ox. But if the analogy with practical skill is the best way of understanding these ideals, then investigating the nature of practical expertise would also be valuable for informing our educational practices today.

Despite the fact that Aristotle and Zhuangzi both appear to have advocated a skill model of ethical expertise, the ethical ideals they endorsed are importantly different. For Aristotle, the ideal agent is someone who has cultivated a fully virtuous character. In *NE*, this requires cultivating moral virtues like justice and courage, mentioned in the passage above, but it also requires cultivating practical

⁴ I follow Slingerland (2003) and Barrett (2011) in using *wu-wei* in a broad way to refer to an ideal that incorporates a cluster of notions or abilities including not least: *wang* ('forgetting'), *ziran* ('self-so'), *cong* ('following'), *you* ('wandering'), *shun* ('flowing along with'), *an* ('at ease') and *yi* ('fitting').

⁵ For an exception, see Machek (2011).

wisdom (*phronesis*).⁶ For Aristotle, practical wisdom is an intellectual excellence that is marked by the ability to deliberate well about what to do in terms of what is good for human beings qua human beings.⁷ Thus, in addition to having cultivated the moral virtues (like courage, moderation and justice), the fully virtuous agent also has the ability to decide between the various needs of the situation, in terms of what is most conducive to the overall human good in general. Zhuangzi, by contrast, expresses scepticism about the value of theoretical knowledge and understanding, i.e., knowledge and understanding achieved by testimony and drawing conceptual distinctions (Z2, Z7). For Zhuangzi, *wu-wei* requires one to ‘fast the mind’ (*xinzhai*) of such distinctions, so that one’s actions can be in harmony with one’s natural environment (Z19). In this way the Daoist sage is presented as being able to naturally adapt to any situation in part because they do not deliberate about what to do. So, despite their both apparently endorsing an important analogy with practical expertise, the ethical ideals that Aristotle and Zhuangzi advocated are crucially different. Aristotle’s writings suggest an intellectual picture of ethical expertise that involves learning to think carefully about what to do, whereas Zhuangzi’s writings suggest an anti-intellectual picture of ethical expertise that involves learning to not think about what to do.

The question I am interested in, however, is not about which ethical ideal, if either, we should endorse. I am happy to be pluralistic in this regard. Rather, the question I am most interested in is: What is the nature of practical expertise? And more particularly: Is practical expertise intellectual or anti-intellectual in nature? Chapters 2 to 5 are taken up with answering these questions. My main thesis is that practical expertise essentially requires an inarticulate, embodied, practical kind of understanding of what one is doing. The last two substantive chapters are taken up with discussing the accounts of virtue and *wu-wei* that emerge on the skill model if this is true.

⁶ NE 1138b35-1139b5. Cf. Broadie & Rowe (2002: 17, 46).

⁷ NE 1140a25, 1140a30. Cf. Broadie & Rowe (2002: 49).

2. The contemporary skill model of ethical expertise

The most prominent contemporary advocate of a skill model of ethical expertise is Julia Annas. In her earlier work, Annas favoured a ‘Socratic’ conception of virtue based largely on Plato’s appeals to practical skill,⁸ and she once wrote that Aristotle is ‘a lone voice’ in that he ‘rejects the idea that virtue is a skill’ when most Greek thinkers ‘followed Plato and the Stoics in holding that virtue is a skill... virtue is, as the Stoics put it, the skill of living’.⁹ In her more recent book, *Intelligent Virtue*, however, Annas defends the skill model, and provides a compelling account of virtue, by drawing heavily from Aristotle, as well as Plato and the Stoics.¹⁰ A core claim of both Annas’ early and more recent defences of the skill model that has not changed is the claim that practical expertise requires coming to understand the activity, and that in virtue of having such understanding experts in practical skills can explain why certain actions are appropriate in some situations and not others, given the goals of the activity.¹¹ Thus, Annas has always defended an intellectual account of practical expertise that has served to descriptively ground an intellectual conception of virtue.

A highly salient feature of Annas’ recent account of practical expertise is her appeal to studies in empirical psychology about the phenomenology of skilful action. By so doing, Annas argues that when acting in character the fully virtuous person’s behaviour results in a ‘flow experience’.¹² The term ‘flow’ was introduced by Mihaly Csikszentmihalyi to refer to the phenomenology that people reported experiencing when totally absorbed in doing something that takes a lot of skill, e.g., playing a musical instrument.¹³ This kind of phenomenology is also sometimes referred to as ‘being in the zone’. In their research, Csikszentmihalyi

⁸ Annas (1993, 1995, 2003).

⁹ Annas (2003: 16).

¹⁰ Annas (2011a). All page references to Annas’ views in the main text are to this work unless otherwise stated. What I am calling Annas’ ‘recent defence’ of the skill model of virtue is part of a bigger project that she has in *Intelligent Virtue* of providing mutually reinforcing accounts of both virtue and ‘happiness’ (*eudaimonia*). However, my focus here is on her accounts of practical skill and virtue and so I won’t be discussing her account of happiness.

¹¹ Annas (1993: 67-68, 1995: 231-233, 2011a: 17-19, 2011b).

¹² Annas (2008, 2011a: Ch. 5, 2011b).

¹³ Csikszentmihalyi (1990, 2014).

and colleagues found that flow experiences are robustly correlated with skilful action across differences in age, gender, and culture.¹⁴ Annas' appeal to the phenomenology of flow in giving her account of virtue is highly salient because Zhuangzi scholars have often compared the ideal of *wu-wei* to the ability to 'act-in-flow'.¹⁵ In her appeal to flow experiences, Annas thus creates a space to build a bridge between the contemporary literatures on virtue and *wu-wei*.

In *Intelligent Virtue*, Annas makes clear that she is not proposing a 'theory' of virtue or practical skill based on 'intuitions' (of some technical kind) that 'can be refuted by a single counterexample' (pp. 3-4). Rather, her aim is to give a 'holistic' account of virtue and practical skill based on 'our everyday views' about virtue and expertise. She says, 'I talk simply about what we think and believe about virtue, and also about skill and expertise' (p.3). The 'we' here is, Annas tells us, used in an 'inclusive, invitational' sense (*Ibid.*).¹⁶ In the concluding chapter, Annas tells the reader:

'If you have found yourself disagreeing with the claims about 'our' everyday views, this is relevant evidence, and best put to use by showing how it does or doesn't suggest modifications to the account.' (p. 170).

In this dissertation, I take up Annas' invitation. In the following chapters, I offer a critique and a modification of Annas' account of practical expertise. In the final two substantive chapters, I apply this modified account of practical expertise to virtue and *wu-wei*.

3. A constraint and a desideratum

It seems clear that both Zhuangzi and Aristotle intended their ethical writings to be practically applicable to real people. And Annas too aims to provide an account of the cultivation of virtue applicable to real people, by appeal to the way real people typically cultivate practical expertise (p.172). Given the goal of practical

¹⁴ See Nakamura & Csikszentmihalyi (2009) for a recent review of some of this research.

¹⁵ E.g., Barrett (2011) and De Prycker (2011), Csikszentmihalyi (1990: 150-151) and Velleman (2008).

¹⁶ Following Annas, my use of 'we' is intended in the same inclusive, invitational fashion.

applicability, the plausibility of an account of how we become ideal ethical agents is, to a significant degree, an empirical matter. So, an important constraint on the skill model of ethical expertise is that the account of practical expertise invoked is empirically adequate. Call this the *empirical constraint* (EC).

EC: A plausible account of practical expertise in real people must be empirically adequate.

EC does not say that the plausibility of an account of practical expertise must be conclusively supported by the empirical evidence, for this would set the bar too high. As Annas says, an empirically adequate account can tolerate some counterexamples (pp.3-4). This does raise the question, however, of how many 'some' is. In giving an account of practical expertise suitable to ground a human ethical ideal, we are making empirical generalisations about the way practical experts actually cultivate and exercise their skills. Intuitively, an empirical generalisation of this sort must hold for the majority of cases to have a chance of being true. If so, then, at minimum, for a claim to qualify as a plausible candidate for a true empirical generalisation of this sort, it must hold in at least 50% of actual cases (and probably a lot more).¹⁷ For instance, if it were the case that only 49% of birds could fly, then the claim 'typically birds can fly' would look highly dubious. In similar fashion, if it were the case that only 49% of gyms had cafes, then the claim 'generally gyms have cafes' would seem false. The claims about cultivating practical expertise in real people at stake on the skill model look to be empirical generalisations in the same way as 'typically birds can fly' and 'generally gyms have cafes' are empirical generalisations. By parity of reasoning then, if a claim about what cultivating practical expertise in real people requires does not hold in at least 50% of actual cases, this seems to disqualify it as a plausible candidate for being a true empirical generalisation. Insofar as an account of practical expertise endorses such a claim, it thereby fails to be empirically adequate, and thus is implausible in the sense intended by EC. Of course, that an account of practical expertise complies with EC is by no means, and is not intended to be, a guarantee for the

¹⁷ This sort of empirical generalisation contrasts with 'by nature' empirical generalisations, e.g., the claim 'by nature birds can fly', which we might still think is true even if only 49% of actual birds could fly.

truth of the account. Rather, EC is intended to guarantee only the bare minimum: that the account is plausible in the sense that it is suitable to play the ground in a skill model of ethical expertise applicable to real people.

As well as the empirical constraint, there is also an important theoretical desideratum on an account of practical expertise, aside from any role it might play in the skill model of ethical expertise. It seems an important part of the concept of expertise that there is something special about the person who has acquired it. Becoming an expert at something is usually something to be celebrated, and something that attracts admiration (and envy) from others. It seems obvious that we value expertise.¹⁸ I think we value expertise because of what the expert can do for us, which makes expertise appear instrumentally valuable. But I also think we value and admire experts for their skill in itself, as a quality of the expert themselves, even if it is of no benefit to us. And this makes expertise look non-instrumentally valuable. If this is right, then an account of practical expertise that has the resources to explain one or both of these ways expertise appears valuable would be more plausible for it. So, a desideratum on an account of practical expertise is that it has the resources to account for what is valuable about practical expertise. Call this the *special requirement* (SR).

SR: A plausible account of practical expertise has the resources to account for what is valuable about practical expertise.

SR is formulated to allow that an account of practical expertise can be plausible without having a specific view about the value of expertise, so long as it contains the materials needed to build one. The following critique of Annas' account of practical expertise is guided by EC and SR.

¹⁸ This might seem to be more controversial than I'm making out. For instance, the 'anti-expert' narrative portrayed in the media might make it seem as if many people don't value expertise all that much. However, I am not convinced of this narrative myself and I assume that any readers who are sympathetic to this anti-expert narrative would still agree that it is obvious that we value practical experts like skilled mechanics, butchers, hairdressers, builders, firefighters, nurses, musicians, and so on.

4. Summary of my argument

In Chapter 2, I reconstruct Annas' account of practical expertise in a way that lays the foundations for critique in subsequent chapters. In Chapter 3, I argue that Annas' account of the cultivation process of practical expertise is not empirically adequate in the way required by EC. The key implication drawn is that practical expertise does not require coming to understand the activity in the way Annas claims. However, since for Annas the (supposed) fact that practical expertise requires understanding is the most plausible candidate for fulfilling SR, the implication that practical expertise does not require this kind of understanding leaves an explanatory gap in her account. In Chapter 4, I try to fill this gap. I suggest that there are two kinds of understanding one might have in regard to a practical activity—*theoretical understanding* and *practical understanding*—and that only theoretical understanding is ruled out by EC. Further, I suggest that the features that made theoretical understanding a plausible candidate to fulfil SR are shared by practical understanding, and that this makes practical understanding apt to fulfil SR.

In Chapter 5, I examine the phenomenology of skilled action, and argue that the 'flow experiences' reported by practical experts are best accounted for by appealing to a Merleau-Pontyan framework of action and perception. Extending this framework, I propose that the phenomenology of flow manifests when the agent is aware of what they are doing in a very specific way, and that this provides the agent with a specific kind of epistemic access to the world. I conclude that the ability to exercise one's skill 'in-flow' is the defining feature of practical expertise, and that doing so manifests the expert's embodied, practical understanding of the activity.

In Chapter 6, I apply the account of practical expertise developed in previous chapters to virtue. I argue that even though the skill model does not descriptively ground an intellectual account of virtue, there is independent reason to hold that full virtue requires having a theoretical understanding of one's actions. The skill model of virtue that emerges holds that the ideally virtuous person would have

both theoretical *and* practical understanding of what they are doing, *plus* a commitment to goodness. In Chapter 7, I apply the modified account of practical expertise to *wu-wei*. In the contemporary literature, it has been common for theorists to interpret *wu-wei* as either identical to, or requiring, ‘acting-in-flow’. I address a recent objection to this view and argue that the extended Merleau-Pontyan framework provides the resources to defuse it. I conclude that on the skill model, *wu-wei* requires an embodied, practical understanding of what one is doing, *plus* a commitment to *Dao*.

2. Annas' account of practical expertise

1. Introduction

In her recent book *Intelligent Virtue*, Julia Annas defends what I am calling a skill model of ethical expertise.¹⁹ Annas gives a detailed account of the process by which she thinks practical skills are cultivated and claims that an analogous process applies to the cultivation of a virtuous character. In doing so, she defends two theses, which I formulate:

Process – cultivating virtue is (like) cultivating practical expertise.

State – the fully virtuous person is (like) the ideal practical expert.

For Annas, the way we develop virtues like courage and generosity is the same way we learn practical skills like playing the piano and tennis (p.16). Further, Annas thinks the crucial analogy between virtue and these kinds of practical skills is the way these skills are 'taught and learnt' (p.21). Thus, her primary focus is on *Process*. Let's call the quality of being *maximally* skilled at some practical activity—like playing a musical instrument, building, butchery, or tennis—having *full* practical expertise at that activity. An ideal practical expert, on this view, is someone that has full practical expertise in the sense that when exercising their skill they never make a mistake, and always perform an appropriate action. In this way, someone with *full* practical expertise—viz. the ideal practical expert—is like the *ideally* virtuous agent who always does the right thing. It is important to note, however, that this notion of practical expertise is gradable in the sense that one can be more or less skilled. After all, we clearly think that there are highly skilled (expert) builders, tennis-player, pianists, butchers, etc., but plausibly none (or very few) of them has *full* practical expertise of the relevant activity in the sense of being

¹⁹ Annas (2011a). Unless otherwise stated, all references to Annas' views will be to this text.

maximally skilled at that activity. (I'm sure many would admit there is still room for improvement). If so, then to qualify as a practical expert it is not necessary that one have *full* practical expertise in this maximal sense, but only that one have a (large) measure of practical expertise in the sense of being (more or less) highly skilled at the activity. Two people might both be expert tennis players, and yet one be more skilled than the other. Given this, we can say that a learner can acquire a practical skill to a certain degree, without thereby qualifying as an expert, i.e., someone who has acquired the skill to a high degree. And that someone who has cultivated a practical skill to a high degree thereby has some measure of practical expertise, and thus qualifies as an expert. To be a practical expert is, on this view, to possess a practical skill to a sufficiently high degree.

It is worth emphasising the distinction between *State* and *Process*. In my terms, practical expertise qua state is something the practical expert has, or is in, so to speak. Cultivating practical expertise qua process, on the other hand, is the way in which practical expertise qua state is achieved. Bearing this distinction in mind should hopefully help in what follows, since trying to discuss one without mentioning or alluding to the other is extremely difficult, and likely counterproductive. Indeed, even though in the early stages of this dissertation my focus is predominantly on *Process*, I think the best way of proceeding is to frame Annas' account in terms of claims related to *State*. In this chapter, I reconstruct Annas' account of practical expertise in these terms. In the following three chapters, I challenge various aspects of that account, and in doing so propose an alternative account of practical expertise. In Chapters 6 & 7, I discuss what kind of account of ethical expertise the skill model generates given my alternative proposal by applying it to virtue and *wu-wei*, respectively.

2. The right kind of skills

On Annas' account, practical skills are dispositions to behave in certain ways that are acquired by habituation, i.e., repeated performance. But for Annas practical skills (like tennis and playing the piano) are different from 'routine habits' (like

tying one's shoelaces), which we also learn by habituation. Unlike habituating a routine habit, the habituation process required to cultivate the kinds of practical skills that are analogous to virtue, Annas tells us, 'needs constant monitoring' because rather than being 'static conditions', the dispositions that underlie practical skills are 'always developing, being sustained or weakened' (p.14). The idea here is that once someone has acquired (say) the ability to tie their shoelaces, they don't need to continue to practise it in order to maintain it, where this is not the case for (say) playing the piano. If a pianist wants to maintain (or increase) their skill, they need to practise. Moreover, the way they practise will affect the quality of their skill. If they don't concentrate when practising, their skill will degrade, whereas if they focus on what they are doing, they will increase their skill. Given this, we can say that a practical skill, as opposed to routine habit, is 'trainable' in the sense that the dispositions underlying it can be weakened or sustained depending on the attention we pay to the way we practise.²⁰ I will often refer to such trainable dispositions as 'abilities'.²¹ In these terms, practical skills are constituted by abilities, which are trainable dispositions acquired by habituation.

Given this, Annas thinks that not every ability that we are prepared to call a skill is the right kind of ability to ground a conception of virtue (pp.12-14, 16, 18 n. 3, 20). For instance, we sometimes refer to the ability to tie one's shoelaces as a skill, but the ability to tie one's shoelaces is *not*, Annas says, a practical skill of the right kind to ground a conception of virtue.²² Annas gives the following examples of practical skills that she thinks *are* analogous to virtue: playing the piano, pottery, rock-climbing, dancing, ice-skating, playing tennis, playing golf, playing football, throwing the javelin, running, athletics, farming, building, plumbing, fixing a computer, surgery, translating Homer, speaking a second language and driving. Importantly, the idea is not that virtue is analogous to any and all of these skills whenever and however they may manifest. Rather, the idea is that the way that skill in these activities is cultivated and *expertly* exercised is analogous to the way virtuous character traits are cultivated and exercised. For instance, Annas gives the

²⁰ Cf. Ryle (1949: 42ff).

²¹ Cf. Swartwood (2013).

²² Annas (2011b).

following example of the cultivation and exercise of her ability to drive to work, which she says is *not* analogous to virtue, but is rather an example of ‘habit becoming routine’.

‘I drive to my university job every day, following the same route to the parking garage. At first I have to think consciously about the best way to do this, avoiding traffic without going too far from the most direct way, modifying the route at different times of day and so on. Gradually I become used to driving on this route, and it becomes habit with me. I no longer have to think about which way to turn at every corner, where to slow down and the like. My driving has become routine. This does not make it mindless: I am still at some level aware of where I am going, since I stop at red lights, drive at the right speed, and behave cautiously around dangerous drivers. But driving has become detached from my conscious thinking, and my conscious and deliberate thoughts may fail to be properly integrated with it. I may find myself at the garage when I started out intending to go somewhere else on route, or find myself at the usual entrance even when I know it is closed for construction. A decision to act differently from usual has not penetrated the patterns of routine, which carry on unaffected. A change to routine has to be conscious, explicit, and sometimes repeated if it is to have the appropriate impact’ (p.13).

Although Annas’ ability to drive her normal route to work described in this passage was acquired by habituation, it is not, she says, a practical skill of the right kind to draw an analogy with virtue. However, it is worth noting that when discussing the kind of enjoyment that characterises both practical expertise and virtue, Annas says the following:

‘There are people who enjoy driving, but for them the experience is an exercise of skill rather than mere habituated routine’ (p.71, n. 11).

While the first passage quoted above suggests that in some cases driving is not the right kind of skill to ground a conception of virtue, the second quotation suggests that in other cases driving skill can manifest in a way analogous to the exercise of virtue.²³ Thus, it appears that Annas’ thought is not that the ability to drive *simpliciter* is the wrong kind of skill to ground a conception of virtue, but rather that the ability to drive which has plateaued into routine habit is the wrong kind (pp.13-15). Since the ability to drive is ubiquitous, it serves as a relatable example of practical skill that many people have experience of. For this reason, I return to

²³ Despite this, I am unsure whether Annas does ultimately take driving expertise to be analogous to virtue. However, I cannot see any particular reason why she would exclude it. So, in what follows I will assume that expert driving skill (e.g. as manifested in Formula 1) is a practical skill of the right kind to ground a conception of virtue on Annas’ view.

examples of driving skill in Chapter 3 and again in Chapter 5. The point being highlighted for now is just that for Annas there appears to be a distinction between exercising a skill in a routine habitual way and exercising *the same skill* in an expert way, i.e., a way analogous to the exercise of virtue.²⁴ If this is right, then the difference between routine habits and practical skills lies, partly, in the way we exercise them.

According to Annas, part of what qualifies an ability as a practical skill, rather than a mere routine habit, is something about the agent, i.e., that they have the ‘need to learn’ and the ‘drive to aspire’ (p.16). Annas says little about what the need to learn amounts to, but I take it she intends to at least exclude innate capacities like our sense of sight from being practical skills.²⁵ Although our sense of sight develops (and degrades) over time, this development does not itself count as a learning process. On the other hand, Annas says quite a few things about the drive to aspire. In Annas’ usage ‘drive to aspire’ is a technical term, which we can define as being a property of the learner of the skill in question, consisting in a cluster of motivational drives: the motivation to continual improvement, to self-direction in performance, and to understanding what one is doing and why (pp.17-19). The most salient of these drives, for my purposes, is the motivation to understand what one is doing and why. This is the requirement that one is motivated to understand the activity, its goals and the principles that underlie it. On Annas’ account, a learner comes to understand an activity by being given explanations by an expert:

‘Something has [to be] conveyed from the expert to the learner which cannot be reduced to showing the learner something to repeat. [...] This may be a simple aim, easily understood, or it may be something more complex, involving grasp of a principle or a set of principles. [...] With skills of any complexity, what is conveyed from the expert to the learner will require the giving of reasons. [...] Reasons enter in here as a medium of *explanation*: the teacher can, by giving reasons for what she does, explain to the learner why she must [act] in such and such a way, as opposed to just showing her that you do it this way. [...] The ability both to teach and to learn a skill thus depends on the ability to convey an explanation by giving and receiving reasons. It thus requires some degree of articulacy’ (pp.18-19, emphasis in original).

²⁴ In support of the idea that this is ‘the same skill’ consider the case of a Formula 1 driver dropping their children off at school.

²⁵ Cf. NE 1103a25 – 1103a30.

In order to cultivate a practical skill the learner needs role models from whom they can learn. But, according to Annas, this is not just a matter of ‘rote copying’, but rather requires that the role model provide the learner with the ‘right information’ (p.17), which Annas thinks requires giving the learner reasons that explain why a particular course of action is the right one in the circumstances. This leads onto another important point. For Annas, part of what makes for the difference between a routine habit and a practical skill is that a routine habit is ‘detached’ from conscious thought, whereas a practical skill of the right kind is responsive to the agent’s conscious thoughts and reasoning. In the case of practical skills, she tells us:

‘[T]he ability, though a habituated one, is constantly informed by the way the person is thinking [...] practical mastery is at the service of conscious thought, not at odds with it’ (pp.13-14).

Practical experts on Annas’ account think about what they are doing when acting expertly, which is not the case during the exercise of a routine habit. Indeed, this is a key reason Annas’ advocates a skill model of virtue—she opens her book with the claim that:

‘[E]xercising a virtue involves practical reasoning of a kind that can illuminatingly be compared to the kind of reasoning we find in someone exercising a practical skill’ (p.1).

3. Intellectualism about practical expertise

Given the above comments, we can attribute to Annas the following claims:

- 1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.²⁶
- 2) S has practical expertise of ϕ -ing only if S has come to understand ϕ -ing.

²⁶ In some cases, it may not be possible to give a reason at the time the action is taken. For example, when playing the trumpet. In such cases, according to 1) the expert will still be able to give a reason after the fact or in relation to hypothetical situations.

Here, ‘ ϕ ’ refers to some practical activity. Call the combination of 1) and 2) ‘intellectualism’ about practical expertise.²⁷ Thus, Annas advocates intellectualism about practical expertise. For ease of exposition, I will sometimes refer to 1) as the *articulation requirement* (AR) on an account of practical expertise, and to the ability that AR attributes to the practical expert as the ability to give ‘reasons-explanations’.²⁸ Reasons-explanations in this sense are utterances that express considerations that help the learner make sense of the prescribed action. For Annas, it is by giving reasons-explanations that the expert facilitates the learner’s coming to understand the skill. For example, suppose Albert, the expert plumber, is teaching his apprentice Glenda how to install a sink. According to intellectualism, if Glenda is to come to understand the task, Albert will need to explain things to her, e.g., that they need to put in a u-bend to stop the sink smelling foul. If this is right, then in order to cultivate practical expertise we need a teacher who can explain why we should act in one way or another in different situations. Thus, on Annas’ account, the nature of the learning process involved in cultivating a practical skill entails that the expert qua teacher can give learners reasons-explanations (pp.17-19). In the next chapter, I examine this argument in detail.

A second, related reason to think AR is true is because it might seem that 1) is entailed by 2), i.e., the claim that the practical expert has come to understand the activity at which they are expert. As we saw above, for Annas, the kind of understanding required for practical expertise involves knowing what the goal of the activity is, which, for complex activities, requires a unified grasp of the principles that underlie the activity. The thought then is that, generally, having such understanding enables the expert to explain their skilful actions.

This raises the question of why Annas thinks that cultivating practical expertise requires coming to understand the activity in the first place. Annas gives two

²⁷ Cf. Hursthouse (1999: 124) and Hills (2015).

²⁸ In later chapters, I also refer to the analogous claim about the virtuous person as the articulation requirement (AR). In most cases which version of AR is being discussed will be clear from the context. When this is not the case, I will be explicit about which of these two versions of the articulation requirement is being considered.

reasons why she thinks 2) is true. The first reason is that coming to understand the activity is what makes for the difference between the practical expert and the ‘clone-like impersonator’ (p.17). Annas gives the following example:

‘Imagine a pianist whose goal is to play like Alfred Brendel, but mistakenly thinks that she will achieve this by copying all his mannerisms and niceties of style, playing only the pieces he plays and playing them just as he plays them. The result would be an impersonation of Brendel, not the achievement of his skill’ (p.17).

The thought here is that if someone acquires a skill just by copying what an expert does, they will end up habituating aspects of the expert’s performance that are inessential to the skill. In this case, what’s missing is that the learner has not acquired the skill *‘for herself’* which requires having the ‘drive to aspire’ and thus, amongst other things, the motivation to understand what one is doing and why (pp.17-18, emphasis in original). Without having an understanding of the activity, just being able to perfectly imitate the practical expert is, Annas thinks, a failure to cultivate the skill in question (*Ibid.*).

The second reason Annas thinks that cultivating practical expertise requires coming to understand the activity is that this is what allows the expert to ‘go ahead in different situations and contexts, rather than simply repeat the exact same thing that was done’ (p.19). The idea here seems to be that by coming to understand the activity the practical expert can reason their way to the correct action, by applying the principles that underlie the activity to the particulars of each new situation in light of the goals of the activity, whereas the learner who does not understand what they are doing and why will not be able to do this. Thus, for Annas, that the practical expert has come to understand the activity explains how they can adapt their performance to different circumstances.

One implication of this is that it is the fact that practical experts have come to understand the activity that makes them special as compared to the non-expert who can perfectly imitate them. This is because having understanding of the activity is what enables the practical expert to exercise their skill in different situations, whereas the non-expert who can perfectly imitate the expert will not be able to adapt to contexts outside of those in which they learnt the skill, unless they

have someone to copy. It is common for theorists to think that understanding something in this way constitutes an intellectual achievement and is thus non-instrumentally valuable.²⁹ Moreover, if being able to achieve the goals of an activity in different situations requires understanding the activity, then this would make understanding the activity instrumentally valuable. Thus, on Annas' account the fact that the practical expert has come to understand the activity is apt to fulfil the *special requirement* (SR), i.e., the desideratum that an account of practical expertise has the resources to account for what makes practical expertise valuable.³⁰ If so, it seems to follow that coming to understand the activity serves as a threshold that marks the transition from being a learner who can skilfully imitate the expert to becoming a practical expert.

4. Acting-in-flow

Annas' goal in *Intelligent Virtue* is not to give an account of what cultivating practical expertise or virtue requires for any conceivable creature.³¹ Rather, she is trying to give an account of what cultivating practical expertise and virtue typically requires for real people as they actually are. 'The account,' Annas says, 'is not presented as a theory which can be refuted by a single counterexample,' but 'answers to our experience as a whole' (pp.3-4). In line with this, Annas appeals to studies in the contemporary psychological literature, wherein highly skilful action has been widely observed to manifest a particular phenomenology, commonly known as a 'flow experience' (p.70).³² The term 'flow' was introduced by Mihaly Csikszentmihalyi³³ and refers to a phenomenology characterised by focussed attention on what one is doing in the present moment, a merging of action and awareness, a loss of self-consciousness, effortless acting, a sense of total control, and being enjoyable for its own sake. The conditions found to be apt to promote a flow experience are when the agent sees the activity as challenging, but

²⁹ See Grimm (2012) for an overview of the recent debate on the value of understanding.

³⁰ Chapter 1, §3.

³¹ Hills (2015:19).

³² See also Annas (2008, 2011b).

³³ Csikszentmihalyi (1990, 2014), Nakamura & Csikszentmihalyi (2009).

not beyond their abilities (and it is not actually beyond their abilities), and the activity has clear goals and immediate feedback. Csikszentmihalyi and colleagues found reports of flow experiences were robust across differences in culture, class, gender and age. In paradigmatic cases, this kind of phenomenology is manifested during the exercise of practical skills, but under the right conditions one can experience flow in almost any activity.

It is important to note, however, that having a flow experience is not, as the term might suggest, something passive that just happens to the agent—a kind of ‘going with the flow’—as Annas makes clear:

‘[A] ‘flow’ experience is best taken to pick out the direct engagement with the task characteristic of expertise, unmediated by deliberation, something active rather than passive’ (p.71).

It is more appropriate, then, to think of flow as a way of acting, and of flow experiences as resulting from the exercise of an ability to ‘act-in-flow’.³⁴ Annas appeals to the practical expert’s ability to act-in-flow for two reasons. The first is to rebut an objection to the skill model of virtue. Namely, the putative dis-analogy that:

‘[A]t least in some clear cases... the success of the skill rests on the excellence of the finished product whatever the affective state of the person producing it whereas this is not true of virtue’ (p.66).

The idea behind this objection is that full virtue requires not only doing the right thing, but also having the appropriate (virtuous) motivations and attitudes. For instance, if someone gives money to charity just because they will receive tax relief, then their action does not qualify as virtuous because the motives behind the action are not in line with the reason one should give money to charity, viz. to help others in need.³⁵ However, so long as the ‘product’ of their actions is excellent, the objection goes, a practical expert’s actions can still qualify as an

³⁴ Cf. Romdenh-Romluc (2012: 205).

³⁵ Some virtue theorists dispute this. For example, in *Uneasy Virtue* (2001), Julia Driver defends a consequentialist view of virtue that holds that any trait that has systematically good effects is a virtue. On this kind of view, making a single donation to charity for tax purposes would not be virtuous, but a trait that systematically causes the agent to donate to charity might well be, even if the accompanying motivations are not intuitively virtuous ones.

exercise of their expertise, no matter what their psychology is like. For instance, an expert potter might make a beautiful vase on commission, and the actions they perform in producing the vase will still count as an exercise of their skill, even if they begrudge making it and do so only because they want the money they will earn. According to this objection, this putative difference is an important dis-analogy that makes modelling virtue on skill unattractive.

The appeal to flow can help blunt this objection because acting-in-flow requires focussed attention on the activity at hand such that one becomes absorbed in what one is doing, which is not possible without having the appropriate attitudes and motivations. This does not mean that an expert potter cannot perform expertly if they are working on a commissioned piece. Rather, the idea is that since acting-in-flow requires focussed attention on what one is doing, if the potter has a begrudging attitude, or is thinking about the money they will earn instead of attending to the job at hand, they will not be able to act-in-flow because they are not fully concentrating on what they are doing. So, even if they produce a beautiful vase, the potter's actions in producing it would not count as a manifestation of their expertise in the full sense. Thus, the fact that when manifesting their expertise practical experts act-in-flow helps diffuse the objection that the practical expert does not need to have the appropriate affective states to qualify as exercising their expertise. There is no dis-analogy because, just like virtue, practical expertise does not depend solely on the excellence of the product, but also on psychology of the agent in producing it.

The second reason Annas appeals to the practical expert's ability to act-in-flow is to account for the claim that the virtuous person finds acting virtuously 'pleasant' (pp.68-76). The idea here is not that the virtuous person has pleasurable feelings or takes joy in all their virtuous actions. Situations which require virtuous actions might well be terrifying, or tragic. In such situations, it would be at best inappropriate and at worst vicious to have feelings of joy or pleasure. Instead, the enjoyment that comes from acting-in-flow is the kind of enjoyment one has when completely absorbed in an activity that is neither too difficult to accomplish, nor too easy that one gets bored, when one's actions are unimpeded by doubts or

recriminations, and all of one's 'relevant goals are harmoniously structured' (p.71).

As Annas puts it:

'It is very plausible that the enjoyment of virtuous activity does not consist in felt twinges of pleasure. It consists, rather, in the way the activity is done; this is not something extra to be added on but just is the ready and unselfconscious way the activity is performed, 'flowing' effortlessly from the person's overall harmoniously arranged goals unchecked by effortful self-questioning or conscious figuring out' (p.76).

Thus, by appealing to the phenomenology of flow, Annas tries to account for the idea that the virtuous person finds acting virtuously pleasant in the sense of experiencing a feeling of harmony, rather than sensory pleasure. For Annas, this feeling of harmony accompanies the fully virtuous person's actions even in circumstances when it would be inappropriate to feel pleasure or joy as normally understood. One may well have doubts about how successful each of these replies is. However, I will not be addressing such doubts here. The point I want to make is just that the fact that the practical expert acts-in-flow does some important theoretical work on Annas' account. Given this, we can attribute to Annas a third claim about practical expertise:

- 3) S has practical expertise of ϕ -ing only if when S ϕ 's to the best of their abilities S ϕ 's-in-flow.

5. Conclusion

In this chapter, I have given an overview of the key features of Annas' account of practical expertise. I have attributed to Annas the following claims:

- 1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.
- 2) S has practical expertise of ϕ -ing only if S has come to understand ϕ -ing.
- 3) S has practical expertise of ϕ -ing only if when S ϕ 's to the best of their abilities S ϕ 's-in-flow.

As I understand it, this forms the heart of Annas' account of practical expertise and the basis of her skill model of virtue. Claims 1) and 2) constitute intellectualism about practical expertise. Claim 3) is drawn directly from empirical studies of practical experts. Applying these claims to the virtuous person we arrive at:

- 4) S has full virtue only if, S can give reasons that explain why some particular action was, is, or would be the virtuous one in the circumstances.
- 5) S has full virtue only if S has come to understand acting virtuously.
- 6) S has full virtue only if when S acts virtuously S acts-in-flow.

On this reconstruction of Annas' account, claims 4) and 5) constitute intellectualism about virtue. On this view, it is characteristic of the fully virtuous person that they can explain and justify their actions, and that they understand the principles of virtue in the sense of seeing the reasons why a particular course of action is what courage, or justice, or honesty, etc., prescribe. Intellectualism about virtue is a controversial position, and Annas' account is supposed to lessen the controversy by showing how these features are inherited from everyday practical expertise. According to 6), when the ideally virtuous person is acting in character, they act-in-flow in the sense that they are so focussed on what they are doing that they lose a sense of themselves, whilst still having a sense of total control over their actions, which are unmediated by deliberation and flow effortlessly such that action and awareness seem to merge, and bring a sense of harmonious enjoyment. If this is a fair reconstruction of Annas' position, then the reason we should think that 4), 5) and 6) are true is because 1), 2) and 3) are true. Thus, Annas' strategy is to descriptively ground these features of virtue by appeal to the analogy with practical expertise.

In the next chapter, I examine the first of these claims, namely, the articulation requirement (AR) on an account of practical expertise. Since Annas is trying to give an account of practical expertise that is empirically adequate, the claims she makes are open to empirical scrutiny. I will present empirical evidence that suggests that when it comes to practical expertise, AR should be rejected.

3. The Articulation Requirement

1. Introduction

In the last chapter, I set out three key claims of Annas' account of practical expertise. In this chapter, I examine the first of these, namely:

- 1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

I labelled this claim the *articulation requirement* (AR) on an account of practical expertise. According to AR, experts in practical activities (like playing the piano, tennis and butchery) can tell us why performing some relevant action is appropriate in the situation. As mentioned in the last chapter, one reason Annas endorses AR is because she thinks '[t]he ability... to teach... a skill... depends on the ability to convey an explanation by giving... reasons' and that practical experts can teach their skill to others (p.19). We can formulate this argument as follows:

- 7) S can teach others skill at ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.
- 8) S has practical expertise at ϕ -ing only if S can teach others skill at ϕ -ing.

So,

- 1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

Call this is the *argument from education*. According to the argument from education, the requirements of practical expertise entail that the expert (qua teacher) must be

able to give the learner reasons that explain why some course of action is the appropriate one in the circumstances, i.e., the expert must be able to give the learner ‘reasons-explanations’. If this reconstruction is fair, then it appears that the fact that practical experts can teach others the skill at which they are expert entails AR. I assume it is a fair reconstruction, so the question is whether the argument from education is any good. The argument appears valid, so the question turns on whether the premises are true. I take it that 8) is true, and that 1) follows straightforwardly from 7) and 8). So that leaves 7). Is it true that teaching a practical skill requires being able to give reasons-explanations? In this chapter, I will argue that for most of the skills Annas mentions 7) is false.

One way Annas tries to support 7) is by pointing out that learning a practical skill requires coming to understand the activity in question, and that in order to come to understand an activity the learner needs to be given reasons that explain why different actions are appropriate in different circumstances. For instance, Annas says that for an apprentice plumber to truly learn her trade, she needs ‘to know not just *that* you do the... pipe-laying in such and such a way, but *why*’ (p.19). The primary source of such ‘whys’ is, for Annas, the expert who is teaching the learner (pp.17-20).³⁶ The main reason the learner needs to come to understand the activity is, according to Annas, because having such understanding is what enables them to apply their skill in new situations (*Ibid.*). The underlying idea here seems to be that having such understanding allows the practical expert to abstract the general principles that underlie the activity, which can then be applied to the particulars of each new situation. That is, understanding the activity enables the expert to reason their way to the appropriate action in new cases. I will save detailed discussion of the role of understanding in acquiring and exercising practical expertise for the next chapter. Foreshadowing that discussion, I accept that cultivating practical expertise requires coming to understand the activity, but, as I see it, the kind of understanding one needs to acquire does not entail AR. That is,

³⁶ There may well be other ways the apprentice plumber could acquire this kind of articulate understanding of plumbing, e.g., they might just figure it out for themselves. In what follows, I set such possibilities aside.

I reject the claim that in order to come to understand an activity the learner needs to be given reasons-explanations.

In this chapter, I focus on a related but different aspect of Annas' account, i.e., the fact that she is offering an empirical account of practical expertise. Annas is not claiming that the process of skill cultivation she describes is true of any conceivable being. Rather, her account is supposed to be of what is required for human beings as they typically are to cultivate practical expertise. Hence, we should think of 7) as offering an empirical generalisation, rather than a conceptual truth. As such, 7) is open to empirical scrutiny. As Annas points out, a few rare counterexamples do not suffice to refute an empirical generalisation (pp.3-4). But empirical generalisations are subject to some sufficiency condition. So, if there are enough counterexamples to 7), then we should reject it. How many is enough? On the face of it, it seems plausible that in order to qualify as a true empirical generalisation a claim must hold for the majority cases. In Chapter 1 §3, I called this the *empirical constraint* (EC). On this reading, for 7) to be empirically adequate it would need to hold in a minimum of 50% of cases. In this chapter, I present two alternative accounts of how practical expertise can be cultivated that do not require that teachers give learners reasons-explanations. I claim that these alternative accounts describe ways of cultivating practical expertise that apply to more than 50% of the practical skills that Annas claims need to be cultivated in the way she describes. If so, then for most of the practical skills Annas lists there are multiple ways of learning them, and most of these ways do not require the teacher to articulate reasons-explanations. In light of this, we can say that teaching a practical skill does not generally require being given reasons-explanations. That is, when taken as an empirical generalisation, 7) is false. If this is correct, the argument from education is unsound and AR does not follow.

It is important to note two things about my argument. The first is that I am not claiming that the account Annas offers is never the way practical skills are learnt. My claim is the weaker one that cultivating most of the practical skills Annas mentions does not *generally require* that the teacher explain to the learner why some particular course of action is the appropriate one in each new circumstance. This

weaker claim is compatible with it being the case that in fact many practical experts did acquire their expertise in the way Annas suggests. It is also compatible with the claim that there are some practical activities for which the process of skill acquisition does require being given reasons-explanations. (Indeed, below (§4) I discuss a class of practical skills for which this seems to be the case.) However, the claim is strong enough to cast doubt on the descriptive accuracy of Annas' account of practical expertise.

The second point is that even if we reject 7), this does not entail that AR is false. After all, even if it is true that qua teachers practical experts do not need to give learners reasons-explanations, because the process of cultivating such expertise does not require being given reasons-explanations, it does not follow that they could not give such explanations. However, since Annas' account is supposed to be empirically adequate, AR should also be taken as an empirical generalisation. So, just as before, if there are enough actual cases wherein practical experts cannot articulate reasons that explain why they did what they did, then we should reject AR. As I will show below, practical experts often cannot articulate an explanation of their actions. Importantly, these cases are relevantly similar to those for which the learning process does plausibly require the giving and receiving of reasons-explanations. Thus, the empirical evidence suggests that even in cases where 7) is more plausible AR is likely false.

The plan for this chapter is as follows. In §2, I outline Hubert Dreyfus & Stuart Dreyfus's model of skill acquisition (hereafter, the Dreyfus model). In §3, I describe the phenomenology of practising martial arts, and abstract a four-stage model of skill acquisition. Unlike Annas' account, neither the Dreyfus model nor the four-stage model imply 7). Further, I claim that for the majority of skills Annas lists as of the right kind to ground a conception of virtue either the Dreyfus model or the four-stage model are typically sufficient for cultivating expertise. Thus, for those skills it is not the case that cultivating practical expertise generally requires being given reasons-explanations. For the remainder of the skills Annas mentions, it is more plausible that the learning process does require being given explanations of why certain actions are appropriate in certain situations. Therefore, for this

subset of skills 7) looks more likely to be true. However, in §4, I present empirical evidence that suggests that even in such cases AR is still often false. This is evidence that tells directly against the articulation requirement. In §5, I conclude that in the case of most practical skills Annas mentions we should reject AR because it does not stand up to evidence provided by personal experiences of cultivating those skills, or empirical studies of practical experts. Rather, we should accept the claim that the ability to articulate reasons-explanations is *not* an ability *generally required* for one to qualify as a practical expert.

2. The Dreyfus model of practical expertise

In this section, I outline the account of skill acquisition proposed by Hubert Dreyfus and Stuart Dreyfus.³⁷ In contrast to Annas' account, the Dreyfus model describes an anti-intellectualist account of the process by which we acquire practical expertise in that there is no requirement for the teacher or instructor to give the learner reasons-explanations. According to the Dreyfus model:

'[T]he beginner develops into an expert who sees intuitively what to do without applying rules and making judgements at all... normally, an expert does not *solve problems*. He does not *reason*. He does not even act deliberately. Rather he spontaneously does what has normally worked and, naturally, it normally works.'³⁸

The Dreyfus model was developed by attending to the phenomenology of cultivating skill in activities such as driving, flying fighter jets, learning a second language, nursing and playing chess.³⁹ The Dreyfus model breaks skill acquisition down into five stages: novice, advanced beginner, competent, proficient and

³⁷ Dreyfus and Dreyfus (1986, 1991). Matthew Stichter (2007) also appeals to the Dreyfus model to argue against the articulation requirement as advocated by Annas on her earlier Socratic account of the skill model of virtue. However, there are important differences between Stichter's critique and mine. For one, Stichter claims that Aristotle is plausibly read as advocating an anti-intellectualist (or, as he puts it, 'empiricist') picture of both skill *and* virtue. Though I have some sympathy with it, I am not committed to the first conjunct, and, like Dreyfus and Dreyfus (1991), I deny the second conjunct. Compare Angier (2010: 128) and Swartwood (2013). Further, Stichter's (2007) project and mine differ in at least two further ways. First, Stichter is only concerned with the skill model of *virtue*, whereas I am interested in a more general skill model of *ethical expertise* that might account for both *wu-wei* and virtue. Second, though Stichter relies heavily on the Dreyfus model he does not appeal to or discuss the Merleau-Pontyian foundations upon which it is built.

³⁸ Dreyfus and Dreyfus (1991: 235, emphasis in original)

³⁹ Dreyfus & Dreyfus (1986).

expert. In what follows, I briefly outline each of these stages, using Dreyfus and Dreyfus's description of learning to drive.

*Novice*⁴⁰

When the learner first begins learning to drive, the instructor points out 'context-free features' of the activity to the learner and gives them 'simple rules' that relate to those features. For example, the novice driver 'learns to recognise such interpretation-free features as speed (indicated by his speedometer)' and the instructor specifies when to change gear in terms of speed. However, the inherent inflexibility of such simple rules often leads to trouble in different situations. For example, following the rule—at 20mph change into third gear—will work in some situations but will lead to stalling the car in others, for example, when travelling up hill.

Advanced beginner

With more experience the novice learns to recognise or is told of other 'situational aspects' and uses 'maxims' that refer to these aspects to decide what to do. The advanced beginner driver 'uses (situational) engine sounds as well as (non-situational) speed. He learns the maxim: shift up when the motor sounds like it is racing and down when it sounds like it is straining. No number of words can take the place of a few choice examples of racing and straining sounds.'

Competent

Over time, the learner learns to recognise more and more features and aspects of the situation, such that they can no longer simply apply a rule or maxim, but must 'adopt a hierarchical view of decision-making' wherein they choose 'a plan, goal or perspective' which allows them to narrow down the number of features and aspects they need attend to. At this stage in the process, the learner's actions are the result of 'analytical rule-governed choice' which involves 'detached planning' and 'conscious assessment' of the salient features of the situation in light of the

⁴⁰ All quotations used in the description of the five stages of the Dreyfus model are from Dreyfus & Dreyfus (1991: 232-5).

plan, goal or perspective chosen. This is followed by ‘an emotionally involved experience of the outcome’. There is an emotional response because learners are not given any explicit guidance about how to choose a particular plan, goal or perspective and so they have to learn which work by trial and error: a successful outcome induces ‘euphoria’ whereas failures are ‘felt in the pit of the stomach’. For instance, the competent driver ‘leaving the freeway on a curved off-ramp may, after taking into account speed, surface condition, criticality of time, etc., decide he is going too fast. He then has to decide whether to let up on the accelerator, remove his foot altogether, or step on the brake. He is relieved when he gets through the curve without mishap and shaken if he begins to go into a skid.’

Proficient

After much experience performing competently, the learner no longer needs to ‘detach’ themselves from what they are doing, or consciously choose a plan, goal or perspective, but instead simply ‘notices’ or ‘is struck by’ one, whilst remaining absorbed in the activity. However, even though at this stage the learner no longer needs to consciously choose a particular plan or perspective, they ‘still have to think about what to do’ by applying rules and maxims. This is because there are ‘generally fewer “ways of seeing” than “ways of acting”’. For example, ‘a proficient driver fearfully approaching a curve on a rainy day may sense that he is traveling too fast. Then, on the basis of such salient elements as visibility, angle of road bank, criticalness of time, etc., he decides whether to let up on the gas, take his foot off the gas or to step on the brake. (These factors were used by the *competent* driver to *decide that* he was speeding.)’

Expertise

Once the proficient performer has sufficient experience of a wide variety of different situations, ‘all seen from the same perspective but requiring different tactical decisions’, they gradually ‘decompose this class of situations into subclasses, each of which share the same decision, single action, or tactic’. This allows ‘an immediate intuitive response to each situation.’ Once they have attained this level of skill, a driver, ‘generally without any attention, not only knows by feel

and familiarity when an action such as slowing down is required; he knows how to perform the action without calculating and comparing alternatives. He shifts gears when appropriate with no awareness of his acts. On the off ramp his foot just lifts off the accelerator. What must be done, simply is done.’

According to the Dreyfus model, cultivating practical expertise is a process which starts with the novice applying simple rules provided by an instructor and progresses through to the expert who doesn’t think about what to do, but intuitively responds to the needs of the situation. As mentioned above, Dreyfus and Dreyfus developed their model by reflecting on their own experiences of what it is like to learn a skill, as well as first person reports from experts. That is, by attending to the phenomenology of cultivating expertise at driving, flying fighter jets, speaking a second language, nursing and playing chess.⁴¹ This list partially overlaps with Annas’ list of skills that she says fit her account. Since the Dreyfus model is drawn from the phenomenology of cultivating these skills, it is an account arrived at by careful reflection on what it was like to actually go through the cultivation process. As such, it provides empirical evidence that in some cases cultivating the right kind skills to ground a conception of virtue does not involve the teacher giving the learner reasons-explanations. Thus, for the overlapping activities at least, the Dreyfus model is a counterexample to 7).

There is empirical evidence, then, which tells against the articulation requirement. Of course, the Dreyfus model is not the only empirical account of practical expertise to be had. But it is a salient one. Its salience lies partly in the theoretical framework within which it is situated. The Dreyfus model is built on a foundation partly provided by Merleau-Ponty’s views of action and perception. In Chapters 4 and 5, I discuss Merleau-Ponty’s views in more detail. In Chapter 6, I suggest that Annas can adopt a Merleau-Pontyan framework and accept an anti-intellectualist account of practical skill, whilst retaining intellectualism about virtue. But there is much to be said before then. In the next section, I outline another way of

⁴¹ It is worth mentioning that Annas would probably object to my use of the Dreyfus model as an empirical counter-example to her account on the basis that this kind of model associates practical skill with ‘know-how’, which ‘obliterates’ the distinction between the learner and the expert (2011a: 16, n. 1). I discuss this issue in more detail in Chapter 5.

cultivating a practical skill that, like the Dreyfus model, does not require that the teacher gives the learner reasons-explanations.

3. A four-stage model of practical expertise

In this section, I describe the phenomenology of cultivating expertise in martial arts. In doing so, I abstract a four-stage model of skill acquisition. As we will see, this process does not require that teachers give learners reasons-explanations. I submit that for many practical skills that Annas mentions as practical skills of the right kind to ground a conception of virtue, the four-stage model describes a way that many people can cultivate expertise. If so, it is not the case that teaching someone these skills generally requires giving them reasons that explain why some particular action is the appropriate one in the circumstances. Thus, the four-stage model abstracted from practising martial arts serves as another counterexample to 7).

It is worth noting that the choice of martial arts here is not arbitrary. First, in ancient Chinese writings the practise of martial arts is itself seen to be a way of ethical self-cultivation.⁴² Moreover, it is likely that martial arts was (and still is) considered to be a way of cultivating ethical expertise because through the practise of martial arts one can attain the ideal of *wu-wei*.⁴³ Second, as should become evident in what follows, martial arts practise also fits several descriptions Annas gives of the kind of skills that are analogous to virtue such as requiring the motivations to continually improve, to self-direction in performance, and to understand the activity. That is, to become an expert martial artist one must have the ‘drive to aspire’.⁴⁴ However, as we will see, the kind of understanding one must acquire in martial arts is crucially different to the kind of understanding Annas thinks is needed for practical expertise. Third, I have been practising martial arts

⁴² See e.g., Allen (2015).

⁴³ In Japan, martial arts are seen as closely aligned with Zen Buddhism. A central concern for both disciplines is the attainment of *mu-shin* (‘no-mind’), which has been compared to flow by Krein & Ilundáin (2014). I think my discussion of *wu-wei* and flow in Chapter 6 complements Krein & Ilundáin’s analysis of the analogy between *mu-shin* and flow.

⁴⁴ Chapter 2, §2.

for many years, and so I can give a first-personal account of the cultivation process I experienced.⁴⁵

The way training in martial arts starts is by cultivating the physical techniques of combat. On this front, expertise in martial arts consists in both having good techniques and being able to use the appropriate techniques in the situation at hand. ‘Techniques’ here refers to hand and leg techniques, including both attacking and blocking techniques, and stances, including posture. A good technique is one that is well executed and effective. Since martial arts practise is training in hand-to-hand combat, expertise involves being able to react quickly enough to subdue one’s opponent. The speed necessary to do this requires a level of spontaneity that precludes conscious deliberation. The martial arts expert is able to immediately and effectively execute an appropriate technique without having to think about it.

I have abstracted four stages from the process of learning martial arts. These can be thought of in terms of four broad kinds of activity that are involved in martial arts training. In prescriptive terms the student must first *copy* the techniques of the instructor, they must *repeat* these techniques many times, and *perfect* them, once they have perfected them, they must come to *understand*. This four-stage process is not strictly linear, but rather cumulative and overlapping, reliant on recursive feedback. Broadly speaking though, these stages can be thought of as constituting a sequential process—a martial arts student cannot understand before perfecting, she cannot perfect without repeating, and she cannot repeat without first copying.

Stage 1: *Copy*

At first, the student copies the techniques of the instructor. This may include a highly simplified description of what the technique involves, but need not.⁴⁶ At

⁴⁵ I started my training in Taekwondo in 1993 with Cargin Moss, who was then a 5th *Dan*, and is now an 8th *Dan* master. I attained my 3rd *Dan* before I stopped grading. To a lesser extent I have also trained with experts in Wing-Chun, Aikido, Shotokan karate, Ju-jitsu and Kickboxing.

⁴⁶ For example, in the relatively modern Korean martial art of Taekwondo, principally developed by Gen. Choi Hong Hi (1918–2002), and represented worldwide by the International Taekwondo Federation (ITF), classes begin and end with students lining up in rows in order of skill (designated by the colour of their belt, see n.45 below). Those with the most skill and experience form the first row, where they are closest to where the instructor usually stands. The students in subsequent rows are distinguished by their degree

this stage, the student does not recognise whether they are performing techniques correctly. Typically, the instructor will make the learner aware of any major errors, and what they need to do to rectify them.

Stage 2: *Repeat*

The student repeats the techniques they have copied from the instructor many times.⁴⁷ Repetition serves to make the techniques habitual, and their performance more uniform and controlled. This process is sometimes referred to in sporting contexts as achieving ‘muscle memory’. Repeating the techniques also serves to increase the student’s fitness by developing strength, stamina and flexibility.

Stage 3: *Perfect*

There are two senses of ‘perfect’ that both apply to this stage. The first is the verb sense of perfect, e.g., ‘you need to perfect your techniques’, which we can understand as meaning refine. The second is the adjective sense of perfect, e.g.,

of skill relative to the students in the first row. Thus, the degree of skill increases as one moves through the rows towards the front, which (typically) correlates with amount of practise students have done. In a mature ITF club, the students in the first row will typically be experts in the comparative sense I discuss in the main text (Stage 3: *Perfect*). Instructors either call out the name of a particular technique, or perform it, or both, and then the students in the front row perform the technique, and the other students try to imitate them. The beginner is in the last row, and if they can’t see the senior grades or the instructor, then they copy the students they can see. A good deal of my martial arts training has involved standing in rows copying other students, with an instructor calling out techniques, and weaving in between us to adjust our finishing positions.

⁴⁷ In the Japanese classic *The Book of Five Rings*, master swordsman Miyamoto Mushashi writes, ‘See to it that you temper yourself with one thousand days of practice, and refine yourself with ten thousand days of training’ (2002: 57). A common misconception amongst non-martial artists and beginners is that once one becomes a ‘black-belt’ they are a martial arts expert. This is a misconception on two fronts. First, some martial arts do not have a formal grading system at all, let alone one distinguished by the colour of one’s belt (sash, etc.). Second, even for those martial arts that do employ such a grading system, there are often numerous black-belts one can attain. For example, in Taekwondo there are ten coloured belts (or *Kup* grades) and nine black-belts (or *Dan* grades). In this kind of system, it takes a minimum of three years consistent training to attain one’s first black-belt. These first years of training are to develop one’s performance of the basic techniques proficiently enough to begin the more rigorous and nuanced training of becoming an expert. That is, at this point the student has achieved enough skill to *start* on the path to becoming an expert. This is exemplified in Taekwondo by the fact that the *Kup* grades are ranked in descending order, i.e., the first belt worn by the beginner (white) is the 10th *Kup*, whereas the last colour belt grade before black-belt (red-tag) is the 1st *Kup*. The nine black-belt grades by contrast are numbered in ascending order (1st *Dan* – 9th *Dan*). This way of numbering the different formal grades shows that the transition from holding a coloured belt to receiving one’s first black-belt marks the beginning of the progression to expert, rather than the end. Of the nine *Dan* grades attainable in Taekwondo, only those with 7th, 8th or 9th *Dan* are considered to be masters. In the ITF, assuming consistent training, the minimum amount of time one must remain at a particular *Dan* grade before being eligible to progress to the next *Dan* grade is equal to the rank of the grade, i.e., 1 year for 1st *Dan*, 2 years for 2nd *Dan*, etc. Given this, it would take 39 years of consistent training and progression to be eligible to receive a 9th *Dan*. This is of course a formal structure that is not universal to all martial arts. The general point is that to become a martial arts expert takes decades of dedicated practise.

‘your performance of that technique was perfect’, which we can understand as faultless. Taking the first sense of perfect, the student refines their techniques in two ways. The first involves performing the first two stages again. This is where the recursive, overlapping nature of the process comes into play. As students continue to practise, the performance of techniques becomes more habitual and precise. They start to discern finer levels of detail in the instructor’s techniques and learn to more accurately adjust their own techniques to match. The more skilled the martial arts student becomes the more aware they become of what the instructor is doing. Moreover, as the student develops, so does their awareness of their own techniques. This is correlated with an increase in how accurately they can imitate the instructor. The student must then repeat the performance of the refined techniques many times so that they become habitual. This stage is sometimes supplemented by the instructor providing instructions regarding the execution of techniques, but need not be.

The second way the student refines their techniques is by testing them in simulated combat by sparring with other students. Amongst other things, sparring develops speed and spontaneity, and helps improve the effectiveness of techniques. During sparring techniques are performed in an organic way that flows from the engaged nature of the activity. Whereas the initial learning and habituation of techniques is done in a highly structured and formal way, to improve at sparring requires the student to learn to ‘free-up’ their techniques so that they are not constrained by this formal structure. By doing so they eventually come to naturally adapt their techniques to each unique situation and opponent. Through this process of repeatedly sparring with many different opponents, students begin to make their techniques their own. We can understand what this means by attending to the fact that physically we are all unique, as is the way our bodies develop when we train. Because of this uniqueness there is not really any one single way to perform a technique correctly, and the formal structured way the student initially learns can be seen to exemplify an idealised and unrealistic way of performing techniques. Sparring helps the student overcome the rigid way they were shown techniques, enabling them to change the way they perform techniques to one that is more

natural. Finally, sparring is a valuable way of testing one's skill in realistic conditions. Sparring at this stage of the process qualifies as real combat inasmuch as it involves being seriously hurt and seriously hurting your opponent. In martial arts, testing one's skills in real life situations is an important way of refining them.

As for the second sense of perfect, i.e., faultless, this refers to the stage in the student's training when they can perform the techniques they have been shown faultlessly. By training for many years, the student has developed speed, stamina and spontaneous reactions, and can perform techniques accurately and effortlessly. At this stage of their training, the student can perfectly imitate the instructor in that their techniques are faultless. In being able to physically perform techniques perfectly in this sense, it might be thought that at this stage we should say that the student has become a martial arts expert. There is a sense in which this is correct. If we think of expertise as a spectrum, a matter of degree, then one is an expert in comparison to someone else who is not. Given this comparative sense of expertise, by the time the student has attained a level of faultless performance of techniques, they will have more expertise than those who have not, the beginner for example. Further, at this stage students are explicitly involved in teaching. Being able to perfectly imitate the instructor qualifies the student to be a role-model for others to imitate.

It is important to emphasise that students might attain a level of perfection without there being any (or extremely little) verbal communication from the instructor. This is illustrated by the fact that it seems in principle possible for an ordinary person to cultivate perfect martial arts technique without being able to speak the same language as the instructor. In such cases, corrections and refinements could be communicated by physical demonstration and physical manipulation, i.e., by the instructor demonstrating what the student is doing and how they need to move in order to correct the technique, or by the instructor actually moving the student's body into the correct position, thereby negating the need to talk.⁴⁸

⁴⁸ I have experienced these methods of correction as both a learner and an instructor.

Stage 4: *Understand*

The three stages above constitute one way we can cultivate perfect martial arts technique. However, becoming an expert martial artist, as opposed to having expertise in the comparative sense, also requires coming to understand what one is doing. In the next chapter, I discuss the nature of understanding in more detail. There, I claim that there are both practical and theoretical kinds of understanding. The difference between the martial arts student who can perfectly imitate the instructor, and the master martial artist is, I claim, fully accounted for by the master's *practical understanding* of what they are doing. What practical understanding amounts to, in short, is the ability to act-in-flow. To perform some activity (appropriately described) in-flow is to manifest one's practical understanding of that activity. This is because acting-in-flow requires being aware of what one is doing in a very specific way, a way that amounts to having a certain kind of epistemic access to the world. Chapters 4 and 5 fill out the details of this conception of practical understanding.

The phenomenology of martial arts practise suggests a model of skill acquisition that has four stages: copy, repeat, perfect, understand. On the four-stage model of practical expertise that emerges, expertise is cultivated by 'instructor-guided-practise', rather than learning the theoretical principles that underpin the activity. This way of cultivating practical expertise does not require teachers to give learners reasons-explanations. Thus, the four-stage model abstracted from martial arts provides a second counterexample to 7).

Several points are worth noting about the two models outlined above. First, a key difference between these two models is their starting points; whereas the Dreyfus model starts with instructions in the form of prescriptive rules and maxims, the four-stage model starts with copying the instructor. But, although not necessary, the four-stage model can allow for instructions of the kind that the Dreyfus model starts with. In similar fashion, it seems in principle possible that someone might learn to play chess or to drive, say, by careful observation and imitation.⁴⁹ Second,

⁴⁹ Cf. Ryle (1949: 41).

as mentioned above, the Dreyfus model is built, in part, on a Merleau-Pontyian foundation. In Chapter 5, I expand the Merleau-Pontyian framework in a way that fits the four-stage model as well as the Dreyfus model.⁵⁰ Third, though the Dreyfus model is an anti-intellectualist model in the sense that there is no requirement for teachers to give reasons-explanations, and so does not support AR, it is, at least up to the stage of proficiency, concerned with how learners come to *decide* or *choose* the actions they are going to perform. However, this is not the case for the four-stage model. On the four-stage model there is no mention of student's deciding or choosing to perform certain actions. Thus, the four-stage model appears to provide a more robust contrast to Annas' account, and a better fit with *wu-wei*, than the Dreyfus model.

4. Expertise in recreational activities

As mentioned in the last chapter, Annas gives several examples of the kinds of practical activities for which she thinks cultivating expertise requires being given reasons-explanations. These include: playing the piano, pottery, rock-climbing, dancing, ice-skating, playing tennis, playing golf, playing football, throwing the javelin, running, athletics, farming, building, plumbing, fixing a computer, surgery, translating Homer, speaking a second language and driving. A few words on these activities. It seems likely that Annas intends her account to apply not just to expert pianists, but more generally to expert musicians. Similarly, it seems likely that the account is supposed to apply not just to expert tennis players, footballers, golfers and athletes, but to sporting experts more generally. These activities seem to form a natural group with rock-climbing, ice-skating and dancing. First, they are activities that most practitioners find intrinsically enjoyable. People often take part in these activities just for the fun of it, and it is not uncommon for people to devote a large proportion of their income and free time to the pursuit of such

⁵⁰ I think the expanded Merleau-Pontyian framework I develop also nicely accommodates other prominent accounts of practical expertise, e.g. Klein (1998, 2009), which I discuss below, and Narvaez & Lapsley (2005).

activities. For this reason, I will call these ‘recreational’ skills.⁵¹ Second, all these recreational skills requires intense physical practise, and expertise in such activities often manifests in graceful, precise or nimble actions that are aesthetically pleasing to watch. Third, expertise in these activities rests solely on how one *does* the activity, i.e., one’s performance is the product of the skill.

I think that both the Dreyfus model and the four-stage model describe ways of cultivating expert skill in many recreational activities such as those Annas lists. In most cases, it is typically possible for us to cultivate practical expertise in a recreational activity without having to learn the theory that underlies it. If this is right, then we can say that the process of learning a recreational skill does not generally require being given reasons-explanations. Further, since the Dreyfus model was developed in part by attending to experiences of learning to drive, speak a second language, fly fighter jets, and nurse, it describes a way of cultivating those particular skills that does not require being given reasons-explanations. Given this, it follows that for the majority of the activities Annas cites—viz. playing the piano, rock-climbing, dancing, ice-skating, playing tennis, playing golf, playing football, throwing the javelin, running, athletics, speaking Italian and driving—it is not generally the case that 7) is true. For most of the skills on Annas’ list, cultivating them does not require teachers give learners reasons-explanations. That leaves a sub-set of practical skills that Annas cites that do not initially appear to naturally fit the Dreyfus model or my four-stage model, i.e., farming, building, plumbing, fixing a computer, surgery and pottery.⁵² I discuss the acquisition and exercise of skills like these in the next section.

In summary, for most of skills that Annas mentions as practical skills of the right kind to ground a conception of virtue either the Dreyfus model or my four-stage model, or both, describe a way of becoming an expert that does not imply 7). It is worth emphasising that my claim is not that either of these alternative ways of

⁵¹ Note, the distinction I am currently drawing in the main text is purely functional. I am not in any way claiming it ‘carves nature at the joints.’

⁵² From now on, I am going to exclude the expert at translating Homer. The reason for this is that I don’t think it qualifies as a *practical* skill. Instead, the ability to translate Homer seems more accurately thought of as a *theoretical* skill. I discuss the difference between practical and theoretical skills in more detail in the next chapter.

cultivating practical expertise described by the Dreyfus model or my four-stage model is necessary for cultivating practical expertise in these activities. My claim is only that they provide different ways one *can* cultivate practical expertise in these cases. That is, the Dreyfus model and/or the four-stage model describe an educational process that is sufficient for cultivating expertise in those activities. Nonetheless, given that 7) is supposed to provide a general requirement on learning all the skills Annas cites, the existence of these alternative models serves to undermine its veracity. So, at least for the recreational activities listed above, we should reject 7). If this is right, then, for the majority of the skills Annas mentions, the argument from education is unsound, and AR does not follow.

5. Expertise in vocational activities

So far in this chapter I have presented two different models of how practical expertise might be cultivated, both drawn from the phenomenology of the learning processes involved. My claim is that for the majority of the activities that Annas lists, these models serve as empirical counterexamples to the claim that learning the skill in question requires being given reasons-explanations. However, this leaves a sub-set of skills that Annas mentions as appropriate to ground a conception of virtue—farming, building, plumbing, fixing a computer and surgery—for which it seems more plausible that learning the skill in question does require learning some of the principles that underlie the activity.⁵³ These activities have further features in common, which helps to group them. They are all commonly undertaken as paid occupations, rather than for fun, and they all have a distinct product apart from the performance of the activity. Moreover, whether one qualifies as an expert in these activities partly depends on the quality of this

⁵³ The reader may have noticed I have not included pottery in either the list of recreational activities or the list of vocational activities. This is not because pottery is not a practical skill, clearly it is. It is because it doesn't naturally fit into either group. For example, on the one hand, it seems plausible that one could develop pottery skills by way of either the four-stage model or the Dreyfus model, and, these days, people often take up pottery for enjoyment and relaxation purposes. On the other hand, there is a product distinct from the performance of the skill, and one's expertise is partly dependent on the quality of this product, and, at least at one time, being a potter was a way to make a living. In any case, if am right that pottery expertise is achievable by way of the Dreyfus model or the four-stage model, then in the case of pottery, 7) should be rejected.

product. After all, we would not call someone an expert plumber if we ended up with leaks all over the place. Call these kinds of activities ‘vocational’ activities.

In contrast to the recreational activities discussed above, it seems more plausible that the learning process involved for vocational activities might well require being given reasons-explanations.⁵⁴ For instance, in economically prosperous countries, it is common for learners of vocational skills such as plumbing to take further and higher education courses in relevant academic subjects, alongside hands-on practise (and this is mandatory for trainee surgeons). Let’s grant then that for this sub-set of activities 7) is true, and that cultivating practical expertise in these kinds of vocational skills generally requires a teacher who can explain why a certain course of action is the right one. If so, then in the case of vocational skills the argument from education looks sound, and so AR follows. In which case, the argument from education appears to show that to qualify as an expert in a vocational activity requires having the ability to give reasons-explanations. However, appearances can be deceiving, and in the case of vocational activities I think they are. In this section, I report on empirical studies that bear directly on AR. These studies show that experts in vocational activities relevantly similar to those Annas mentions often cannot give reasons-explanations for their actions. Moreover, it seems that the cases in which these experts fail to be able to explain why they acted as they did are cases in which they were acting expertly. The implication is that even though the learning process involved in cultivating vocational skills might generally require teachers to give learners reasons-explanations, it is often the case that experts in these activities cannot explain why they acted as they did.⁵⁵

⁵⁴ Only ‘more plausible’ because the Dreyfus model was in part developed by studying the phenomenology of such vocational skills, e.g., nursing (cf. Brenner 1984) and flying fighter jets.

⁵⁵ Although she has doubts about the skill model of virtue, Hills (2015) suggests that there is empirical research into practical expertise that supports Annas’ view (e.g., Ericsson et al. 1993) but concludes that ‘there is no agreement among psychologists that... any rival theory of practical expertise, is correct. The evidence is complicated and does not support any... theor[y] of moral virtue... conclusively’ (2015: 19). However it bears pointing out that it is not necessary to have conclusive support for any particular theory to show that an empirical hypothesis is false. All that is needed is a sufficient number of counterexamples to the hypothesis.

In a series of studies, Gary Klein and colleagues developed a model of expert decision-making by studying firefighters, pilots, nurses, military leaders, nuclear power plant operators, chess masters, and experts in a range of other domains.⁵⁶ A failing of empirical research on practical expertise up to that point, Klein thought, was that it took experts out of their ‘natural decision-making’ (NDM) habitats and relocated them within the laboratory.⁵⁷ NDM researchers set out to investigate how experts actually decide what to do in real life situations. The results of NDM studies are various and complex.⁵⁸ In the current context, the key result is that often in real life situations the experts studied could not explain why they performed the actions they did and had no inkling of how they knew to do so.

In one set of studies, NDM researchers found that expert nurses on a neo-natal intensive care unit (NICU) were able to tell that certain new-born babies had an infection before they had blood test results.⁵⁹ As Klein reports it, these nurses ‘could look at a baby, even a micro-baby, and tell the physician when it was time to start the antibiotic’.⁶⁰ The nurses’ ability to detect infections was trusted to such an extent that even when initial tests results came back negative, doctors would often still start antibiotics based on the nurse’s testimony. In most cases, a second set of tests came back positive, confirming the decision. With the help of the researcher, some nurses managed to describe some of the ‘cues’ that had caught

⁵⁶ Klein (1998, 2009) collates and summarises the key results of these studies. Swartwood (2013) also appeals to NDM studies (see below) in developing his ‘expert skill model of wisdom’ (*phronesis*), claiming the fully virtuous person sometimes uses explicit reasoning and sometimes uses intuition, depending on the features of the case. However, as I argue in the main text, given that the experts studied do sometimes make mistakes, they do not have *full* expertise in an ideal sense. They are all experts in the comparative sense. By analogy this seems to make Swartwood’s model a model of ethical expertise in a comparative sense rather than in an ideal sense.

⁵⁷ Klein (1998, Ch.1), cf. Klein & Kahneman (2009).

⁵⁸ There are two key elements to the *recognition-primed decision making* (RPD) model that NDM researchers developed (Klein 1998, Ch.3, Ch.4, Ch.10). The RPD model suggests that most of the time experts just react to subtle patterns of ‘cues’ in each new situation, based on experiences of many similar situations in the past, which they often cannot articulate. This is an ability Klein refers to as ‘situation awareness’ (pp.152-154). The RPD also suggests that in complex situations practical experts do explicitly think about what to do, but usually not by applying principles or comparing options ‘side-by-side’. The most prevalent way of thinking found in the experts studied was to simulate the most intuitive course of action, and accept or reject it depending on the nature of any foreseen problems. If the first course of action simulated is rejected, the expert then moves onto the next most intuitive course of action and assesses its viability, and so on. I think both these elements of the RPD are nicely accommodated within the Merleau-Pontyan framework I describe later.

⁵⁹ Klein (1998: 41-43).

⁶⁰ *Ibid.* p.42. A micro-baby is a premature baby that weighs less than 2 lbs/907g.

their attention, but without this help they often could not explain how they knew a particular baby had an infection, simply putting it down to ‘experience’ and ‘intuition’.

In another set of studies, NDM researchers found that expert firefighters often could not explain why they acted as they did, or how they knew that a course of action was appropriate.⁶¹ In one case, a fire crew attended what appeared to be a routine kitchen fire in a one-storey domestic residence. When they arrived, the crew went in the front door, through to the kitchen in the back, and started to spray water on the fire. But this had little effect. The lieutenant in charge suddenly felt something was wrong, and on the basis of this feeling ordered the crew to evacuate the building. As soon as the crew had left, the floor where they had been standing collapsed into the basement, where the seat of the fire actually was. The lieutenant had had no idea that the house had even had a basement and could not explain how he knew they should get out, putting it down to ‘ESP’ and ‘a sixth sense’. Again, with careful questioning, researchers were able to help the lieutenant reconstruct some of the cues that might have led to his feeling that something was not right, but this was not something the lieutenant had been able to do by himself.

Both nursing and firefighting are examples of vocational activities in that (typically) they are paid occupations that have a distinct product (e.g., a healthy neonate or a building not burnt to the ground), and the education process involved often requires academic study and learning the principles of the trade. Given this, it follows that there is empirical evidence that even in vocational activities that require studying theory to learn, there are times when experts cannot give reasons-explanations. Thus, there is empirical evidence that tells directly against AR.

A question that arises from this result is: How is it possible that AR is false in these cases if we grant that for vocational activities the argument from education is sound? The answer is that the context of the argument limits the scope of AR. The argument from education only applies to the learning process, not to the exercise of expert skill in real-life situations. If this is right, it would explain how

⁶¹ Klein (1998: 33-37).

the argument from education can be sound in the face of the results of NDM studies. On this reading, we can say that when teaching a vocational skill, an expert must be able to give the learner reasons-explanations, but that for recreational skills this is not the case, and, when performing at their best, practical experts in both vocational and recreational activities often cannot give reasons-explanations.

At this point, Annas might claim that the fact that the nurses and firefighters studied could not provide explanations on these occasions is just an anomaly.⁶² On the presupposition that most of the time the experts studied could explain their actions, we could claim that what matters is how practical experts act generally, and so we can ignore the above cases as rare outliers. On this suggestion, we could maintain that, in the cases that matter, practical experts can always give reasons-explanations. The problem with making this move is that the two examples cited are not isolated incidents, but illustrative of a more widespread phenomenon.⁶³ The prevalence of this kind of inarticulate expert skill generalises to the extent that it forms a core part of the model of practical expertise NDM researchers describe. Thus, these cases cannot be easily dismissed.

Another response Annas might make is to claim that although the experts studied have expertise in the comparative sense discussed in the last section—they are more skilled than most of their colleagues—they do not have *full* expertise in the non-comparative way suitable to ground an ideal of full virtue, i.e., someone who ‘always gets it right’.⁶⁴ The reason why the practical experts studied cannot articulate reasons that explain their actions, according to this response, is because they are not fully expert. Thus, we could account for the inability of the nurses and firefighters cited above to explain their actions as a deficit in their otherwise considerable skill. Given this, we could still hold that the *ideal* practical expert would always be able to articulate the appropriate kind of explanations.

⁶² Hills (2015: 21).

⁶³ Klein reports a similar inarticulacy in expert tank commanders and experienced sailors (1998: 33), expert paramedics (*Ibid.*: 45), airline pilots (2009: 17), and land mine ‘lifters’ (*Ibid.*: 42-43).

⁶⁴ Hills (2015: 23).

But while it is true that the practical experts studied do not have full expertise in an ideal sense—they do sometimes get it wrong—it does not follow that if they did have full expertise that they would always be able to articulate reasons-explanations. Another possible explanation goes in the other direction. On this explanation, it is the lack of full expertise that explains why sometimes the experts studied need to think about what to do.⁶⁵ On this view the ideal practical expert would never need to reason their way to the appropriate action, because they would always intuitively see the best thing to do in the circumstances. I think that the fact that the occasions on which the experts studied could not explain their actions were occasions when they were manifesting their expertise provides support for this direction of explanation. My claim, then, is that the inability of practical experts to be able to explain their expert actions in real life situations is best explained by the fact that in manifesting their expertise practical experts do not reason their way to the appropriate action, but instead intuitively see what to do without having to think about it. In the next two chapters, I suggest that Merleau-Ponty's views of perception and action can provide a fitting theoretical framework for this view.

6. Conclusion

In this chapter, I have examined the articulation requirement, i.e., the claim that:

- 1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

I evaluated a key argument Annas gives in support of 1), which holds that teaching a practical skill requires the practical expert to give reasons that explain why some particular action is the appropriate one in the circumstances. I outlined two alternative ways of cultivating practical expertise based on the phenomenology of the learning process involved—the Dreyfus model and the four-stage model—

⁶⁵ Dreyfus and Dreyfus (1991).

and claimed that for the majority of practical skills that Annas mentions the existence of these alternatives serves to undermine her argument, because on these alternative models, even in cases when the instructor needs to tell the learner what to do, the instructor does not need to explain to the learner why a particular action is the appropriate one in the circumstances. I then presented empirical evidence from NDM studies that tells directly against the articulation requirement, even in those cases where the argument from education looks sound. On the strength of this evidence, I submit that we should reject 1). Given this, if it is true that the virtuous person can always articulate reasons that explain their actions, this is not a feature inherited from the analogy with practical skill.

In light of these results, in the next chapter I examine the claim that practical expertise requires coming to understand the activity. I suggest that because for Annas coming to understand the activity implies being able to give reasons-explanations, the empirical evidence presented above also tells against the claim that practical expertise requires having an articulate understanding of the activity. Given this, having theoretical understanding of the activity is not apt to account for what is special about the practical expert. I propose instead that the defining featuring of practical expertise is having an embodied practical understanding of the activity, and that this can account for what makes the expert special.

4. The special requirement

1. Introduction

In the last chapter, I examined the articulation requirement (AR)—i.e., the claim that practical experts can articulate explanations that give reasons why a particular action is the appropriate one in the circumstances. There, I argued that when taken as an empirical generalisation, AR should be rejected because it does not stand up to the evidence provided by experiences of cultivating practical skills, or empirical studies of practical experts. In this chapter, I turn my attention to the second claim I attributed to Annas, namely:

2) S has expertise of ϕ -ing only if S has come to understand ϕ -ing.

According to the four-stage model of skill acquisition I abstracted from martial arts in the last chapter, coming to understand the activity is the final stage in the process of becoming an expert. Understanding, I claimed, is what distinguishes the practical expert from the skilled but non-expert practitioner who can perfectly imitate the actions of the expert. Moreover, if my account is correct, then having such understanding can account for what makes practical expert special. In this way, that practical expertise requires coming to understand the activity fulfils the *special requirement* (SR)—i.e., the desideratum that an account of expertise should have the resources to account for what is valuable about expertise.⁶⁶ On the face of it, then, my account of practical expertise appears to be in line with Annas' account in this respect. For Annas also thinks that cultivating practical expertise requires coming to understand the activity, and that this is what makes for the difference between the genuine expert and the 'clone-like impersonator' (p.17).

⁶⁶ Chapter 1, §3.

Despite initial appearances however, there is a crucial difference in the way I conceive of the understanding of the practical expert, as compared to how Annas thinks of it. As Annas puts it, the practical expert has acquired ‘an understanding that can be represented as a unified grasp of the principles that underlie [her] actions and decisions.’⁶⁷ And they are able ‘explicitly to explain and justify [their] particular decisions and judgements, and to do so in terms of some general grasp of the principles which define that skill.’⁶⁸ These remarks suggest that on Annas’ account the kind of understanding required for practical expertise involves having a unified grasp of the principles that underlie the activity, and that, generally, this enables the expert to give reasons that explain and justify their actions. Thus, it appears that for Annas the fact that practical expertise requires understanding the activity entails AR. We can schematise this argument as follows:

2) S has expertise of ϕ -ing only if S has come to understand ϕ -ing.

9) If S has come to understand ϕ -ing, then, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances

So,

1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances

Call this the *argument from understanding*. According to the argument from understanding, having practical expertise entails that, in relation to the activity in which one has expertise, one can explain why one acted in one way rather than in another. If this reconstruction is fair, then for Annas the fact that practical expertise requires coming to understand the activity entails AR. However, since there is empirical evidence that suggests that the ability to give such explanations is not required to qualify as a practical expert (the conclusion of the last chapter),

⁶⁷ Annas (1993: 67, note omitted).

⁶⁸ Annas (1995: 233). It is significant that Annas does not give an explicit statement of either of these claims in *Intelligent Virtue*. However, I think Annas’ remarks in *Intelligent Virtue* (esp. pp.17-19) suggest that this view has not significantly changed.

it appears that AR is false. In which case, either the empirical evidence is misleading, or there is something wrong with the argument from understanding. I assume the empirical evidence is not misleading. So, there must be something wrong with the argument. So, what's wrong with it?

First, it is important to notice that 2) is ambiguous between two different kinds of expertise one might have in relation to some activity—*practical* expertise and *theoretical* expertise. For example, if someone is, e.g., an expert ice-skater, their expertise is practical in nature in that it requires performing physical actions. On the other hand, if someone is, e.g., an expert mathematician, their expertise is theoretical in nature in that it requires performing actions in thought. As we will see, on this view someone can qualify as either a theoretical expert or a practical expert (or both) in regard to some practical activity. Since on both Annas' account and my four-stage model the difference between the practical expert and the skilled but non-expert practitioner who can perfectly imitate the expert is that the practical expert has come to understand the activity, it seems natural to think that the defining feature of expertise is that it requires understanding.⁶⁹ If so, then since according to the argument from understanding someone is an expert in regard to some activity only if they have come to understand the activity, one natural place to look for what makes for the difference between having theoretical expertise and having practical expertise of some activity is the kind of understanding that relates to each. Given this, we might distinguish two versions of 2):⁷⁰

2a) S has *theoretical* expertise of ϕ -ing only if S has come to *theoretically* understand ϕ -ing.

And:

⁶⁹ Cf. Wilkenfeld et al. (2016).

⁷⁰ Technically there are more than two versions of this claim. One might hold, e.g.,
2c) S has *theoretical* expertise of ϕ -ing only if S has come to *practically* understand ϕ -ing, or
2d) S has *practical* expertise of ϕ -ing only if S has come to *theoretically* understand ϕ -ing, or
2e) S has *theoretical* expertise of ϕ -ing only if S has come to *theoretically* and *practically* understand ϕ -ing, or
2f) S has *practical* expertise of ϕ -ing only if S has come to *theoretically* and *practically* understand ϕ -ing.
I set these possibilities aside in this section as 2c) and 2d) seem to be non-sequiturs, and 2e) and 2f) initially seem overly demanding. However, see sec. 3.2 and chapter 6.

2b) S has *practical* expertise of ϕ -ing only if S has come to *practically* understand ϕ -ing.

I discuss the details of these two kinds of understanding below. For now, we can think of having theoretical understanding of ϕ -ing as having certain cognitive abilities with regard to ϕ -ing, and of having practical understanding of ϕ -ing as having certain physical abilities in regard to ϕ -ing. Both 2a) and 2b) seem initially plausible. But importantly they lead to two different versions of the argument from understanding.

2a) S has theoretical expertise of ϕ -ing only if S has come to theoretically understand ϕ -ing.

9a) If S has come to theoretically understand ϕ -ing, then, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

So,

1a) S has theoretical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

And:

2b) S has practical expertise of ϕ -ing only if S has come to practically understand ϕ -ing.

9b) If S has come to practically understand ϕ -ing, then, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

So,

1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

In terms of the skill model, we are interested in what practical expertise requires, so the target conclusion is 1). But, whereas 9a) seems plausible, my claim is that 9b) is not. That is, we have no reason to think that someone who has come to practically understand some activity will be able to explain their actions. If this is correct, then *theoretical* expertise does entail the articulation requirement, whereas *practical* expertise does not. Even though cultivating practical expertise requires coming to understand the activity, this does not entail that the expert must be able to give reasons-explanations. Thus, my aim in this chapter is to show that in the case of practical expertise it is possible to hold that understanding is what fulfils the special requirement, whilst still rejecting the articulation requirement.

The plan for this chapter is as follows. In §2, I consider two cases of expertise and disambiguate theoretical expertise of ϕ -ing from practical expertise of ϕ -ing. If understanding is what makes expertise valuable, then disambiguating these two ways one can be an expert in relation to a practical activity opens up the possibility that there are two corresponding ways of understanding an activity. To a greater or lesser extent, it seems plausible that those who actually are experts in regard to some practical activity will have some measure of both practical expertise and theoretical expertise, and correspondingly some measure of the relevant kinds of understanding. That is, in many actual cases these two kinds of understanding will often come together. However, it seems equally plausible that in principle these ways of understanding an activity can come apart. In §3 I outline a ‘working account’ of theoretical understanding by drawing on leading accounts in the contemporary epistemological literature. And in §4, I introduce two ways our understanding of an activity might be practical: the conceptual metaphor view suggested by George Lakoff and Mark Johnson and the embodied understanding view suggested by Merleau-Ponty’s writings. I claim that the kind of practical understanding practical expertise requires is the kind of embodied understanding suggested by Merleau-Ponty’s writings. If this is right, then since this kind of embodied practical understanding does not entail the ability to give reasons-explanations, we can explain where the argument from understanding goes wrong. Finally, in §6, I briefly discuss several ways that having theoretical understanding

is considered valuable and suggest that each way transfers over to having embodied practical understanding.

2. Kinds of Expertise

According to the argument from understanding, cultivating expertise in regard to some practical activity requires coming to understand the activity. However, I think that there are different ways in which one might qualify as an expert in regard to some practical activity, and that the argument from understanding equivocates between these two kinds of expertise. We can see this by considering the following cases of golf expertise.

Case 1: Golf expertise I

Claire is a world class golfer, and she has won many international competitions, accumulating both fame and fortune. Claire can perform just about any golf shot possible. She has some natural talent, but her golf skills are primarily the result of years of dedicated practise and training. Though she is an expert golfer, Claire knows next to nothing about golf theory. For example, she knows nothing about the physical theory relating to golf, or how different weather conditions might affect matters. Though often Claire picks the appropriate club for the shot, she does so on instinct, and when asked, cannot say why it is the right club.

Case 2: Golf expertise II

John is Claire's caddy. From an early age John spent his weekends at the local golf club. In his teens he got a part-time job as a caddy, carting clubs around for the players, watching them play, listening to their discussions about club selection, the appropriate shot, weather conditions, back-swing, and more. He reads every golf book he can get his hands on. And he even wrote his undergraduate dissertation on 'the perfect golf swing'. Despite this, John is a terrible golfer.

Reflecting on these cases suggests that there are two kinds of expertise one might have in relation to some practical activity.⁷¹ In these examples, both Claire and John are golfing experts. Claire is an expert golfer who acquired her expertise by practising golf. John is an expert in golf theory who acquired his expertise by observing golfers and studying the principles of the game. Thus, it appears that Claire's expertise is practical in nature, whereas John's expertise is theoretical in nature.⁷² Claire's expertise is practical in nature in virtue of the fact that the abilities it consists in are bodily abilities cultivated by doing the activity; whereas John's expertise is theoretical in nature because it consists in cognitive abilities, cultivated by studying the theory that underlies the activity. Claire can play almost any golf shot possible, and it is predominantly this ability which makes her an expert. John, on the other hand, is able to explain why a particular club is the appropriate one for a particular shot in terms of the theoretical considerations relevant to golf, e.g., angle of the club-face, distance from the pin, and environmental factors like humidity, and the speed and direction of the wind. It is largely this ability that qualifies John an expert in regard to golf. If this is right, then we can say that while Claire and John are both experts in regard to golf, Claire's golfing expertise is practical in nature, whereas John's golfing expertise is theoretical in nature.

As another example, consider that one could be an expert in snooker theory—including, e.g., being an expert at geometry and ballistics—without being able to play snooker very well. And it seems clear that one does not need to be an expert at geometry to qualify as a great snooker player. I think this line of reasoning probably generalises to most recreational activities.⁷³ But the distinction between practical expertise and theoretical expertise does not seem to be limited to just recreational activities. Vocational activities like teaching also admit of practical and theoretical experts. An amazing teacher may not be well versed in education theory. Conversely, someone might be exceptionally well versed in education

⁷¹ Julia Driver (2013) argues for a similar distinction in the moral case. According to Driver there are at least three ways of being be a moral expert: in practice, in judgement and in analysis.

⁷² This is not to say that caddying is not a practical skill, but just that the kind of expertise John has of the activity of 'playing golf' is not practical in nature.

⁷³ Chapter 3, §4.

theory, but not be able to put this theory into practice very well. Finally, consider the expertise of the professional critic. For many practical activities—e.g., cooking, making films, painting, novel writing, making music, sculpting—there are expert critics. But it seems unlikely that all expert food critics are good cooks, or that the world’s premier art critic can paint well. If this generalises, then it follows that in regard to some practical activity one can have practical expertise or one can have theoretical expertise, or both. Further, there are some activities, e.g., doing mathematics, for which someone could have theoretical expertise, but for which there is no corresponding practical kind of expertise to be had. Call these ‘purely theoretical’ activities. Purely theoretical activities in this sense require sophisticated cognitive skills, but no particular physical skills. Annas’ example of being an expert at translating Homer seems to me to be a case of a purely theoretical activity in this sense.⁷⁴ Finally, there seem to be activities for which to qualify as an expert one would need to have both theoretical expertise and practical expertise in this sense. Performing brain surgery seems a clear example of such an activity.

In summary, the skill model of ethical expertise draws an analogy between practical expertise and some ethical ideal. But in terms of *practical* activities, like golf, we might have in mind either practical expertise or theoretical expertise (or both). For some activities, both kinds of expertise are possible, though for others it may be that only theoretical expertise is possible.⁷⁵ In the context of the argument from understanding, if there is a difference in kinds of expertise one can have in regard to a practical activity, a natural explanation for this difference is that there is a corresponding difference in the kind of understanding that relates to each. Following this line of thought, we can disambiguate:

- 2a) S has theoretical expertise of ϕ -ing only if S has come to theoretically understand ϕ -ing.

From:

⁷⁴ See note 52 above.

⁷⁵ Could there be other activities in which only practical expertise is possible? I can’t think of any likely candidates. It seems that if we can talk about an activity, someone might have theoretical expertise of that activity.

2b) S has practical expertise of ϕ -ing only if S has come to practically understand ϕ -ing.

According to both these premises, cultivating expertise requires the agent to go through the process of coming to understand the activity. As suggested by the golf expertise cases above, even though they share a common object (ϕ -ing), the features of theoretical expertise and practical expertise, and the way in which each is cultivated, are quite different. Practical expertise involves cultivating the ability to perform a particular set of skilful actions to a very high standard, whereas theoretical expertise involves learning the theory that underpins such a performance. The way in which each kind of expertise manifests the expert's understanding is also quite different. Claire's expertise is manifested in her ability to play almost any golf shot possible. John's expertise, on the other hand, is manifest in his ability to give accurate explanations in terms of the theory that underpins golf. I propose that there are two distinct kinds of understanding that correspond to these two different kinds of expertise: *practical understanding* and *theoretical understanding*. As I construe them, both practical understanding and theoretical understanding necessarily have features in common that unite them as kinds of understanding. But there are also important differences. In the context of the argument from understanding, chief amongst these differences is that practical understanding does not imply the ability to articulate reasons-explanations, whereas theoretical understanding does.

In the next section, I give an account of theoretical understanding based on leading views in the recent literature. It is important to emphasise that I am not claiming that any particular theorist mentioned holds the view explicated. Rather, my intention is to give a plausible 'working account' of theoretical understanding that is generally faithful to the contemporary debate on understanding. The account is a 'working account' in that it is explicated with an eye to providing a basis upon which to build an analogous account of practical understanding in coming sections. Nonetheless, I think the account of theoretical understanding I explicate is both plausible and a fair representation of what we might call 'the standard view' emerging from the contemporary literature.

3. Theoretical understanding

Until very recently, the topic of understanding was not of central concern in contemporary epistemology. Most epistemologists were primarily concerned with the concept of knowledge and associated concepts like justification and testimony. Lately, however, there has been a shift in focus and understanding is increasingly being seen as an important epistemic notion, worthy of serious attention.⁷⁶ Although there are many points of disagreement among theorists, a general picture of understanding is beginning to emerge. For example, many epistemologists conceive of understanding as an intellectual virtue.⁷⁷ In light of Annas' view of the metaphysics of virtues, we can say that intellectual virtues are good traits of character, each constituted by a suite of 'trainable' cognitive dispositions (abilities) acquired by practise.⁷⁸ Intellectual virtues, in this sense, are a kind of complex cognitive skill.⁷⁹

Theorists working on the epistemology of understanding often draw distinctions between different kinds of understanding based on the way we use the term 'understanding' and cognates in ordinary language. The two primary kinds of understanding distinguished are 'understanding-why' some statement is true, and 'objectual understanding' of a subject-matter, system, domain, or activity—i.e., something far more complex than a single statement. For instance, if Andy's house burnt down because some of the wiring was faulty, then if Andy knows this and has a grasp of how faulty wiring can lead to house fires, then Andy understands why his house burnt down.⁸⁰ In this case, Andy's understanding relates to a

⁷⁶ See Zagzebski (2001) and Kvanvig (2003) for two important early movers in this shift. Zagzebski (2001) is a rich discussion that includes seeds of both theoretical and practical understanding. Kvanvig (2003) was one of the first authors to explicitly distinguish between 'understanding that', 'objectual understanding' and 'understanding how' (pp.188-190).

⁷⁷ Zagzebski (2001), Kvanvig (2003), Riggs (2003), Pritchard (2009), Grimm (2019). See Zagzebski (1996) and Sosa (1991) for seminal contrasting general accounts of intellectual virtue.

⁷⁸ Chapter 2, §2.

⁷⁹ Cf. Hills (2016).

⁸⁰ This example is borrowed from Pritchard (2009: 31).

statement: that his house burnt down because of faulty wiring.⁸¹ By contrast, if Melanie understands American football, then her understanding relates not to a statement, but to a game, viz. American football. Likewise, if Elmer understands the civil service, then the object of Elmer's understanding is an institution, viz. the civil service. In both Melanie's case and Elmer's case, the kind of understanding they have is objectual understanding. Both understanding-why and objectual understanding are kinds of what I am calling theoretical understanding.⁸² Since the argument from understanding concerns understanding an activity, the kind of theoretical understanding involved is objectual understanding. Henceforth, when I talk of theoretical understanding this should be understood as objectual understanding, unless otherwise stated.

Drawing on leading accounts of understanding in the contemporary literature, we can pick out three component dispositions or abilities of the agent generally thought constitutive of their theoretically understanding something. The first is often cashed out as the ability to cognitively 'grasp' the structure of the object of one's understanding.⁸³ This is the ability to see, or the potential to become aware of, how certain features of the subject-matter, system or activity depend upon and relate to one another. Second, theoretical understanding involves the ability to cognitively 'manipulate' the structure one has grasped.⁸⁴ And third, theoretical understanding entails, other things being equal, the ability to explicitly articulate informative explanations, accurate predictions, insightful questions, and so on, in regard to the object of one's understanding.⁸⁵ In the next three subsections, I discuss each of these abilities.

⁸¹ Of course, from Andy's perspective the statement would be: that *my* house burnt down because of faulty wiring. Or maybe: that *the* house burnt down because of faulty wiring.

⁸² There are interesting and difficult questions surrounding the relationship between understanding-why and objectual understanding, and also surrounding the relations between these kinds of understanding and different kinds of knowledge. Thankfully, for the purposes of this thesis, I need not get entangled in these questions. But see Chapter 5, §6.

⁸³ Kvanvig (2003), Grimm (2010, 2011, 2012, 2017), Strevens (2013), Hills (2016). It is important to note that Hills' focus is on understanding-why as opposed to objectual understanding, however, I think much of what she claims of understanding-why applies equally well to objectual understanding. Grimm's discussions of understanding often move back and forth between understanding-why and objectual understanding, and I take it that much of what he says is intended to apply, *mutatis mutandis*, to both.

⁸⁴ Grimm (2011), Wilkenfeld (2013), Hills (2016).

⁸⁵ Zagzebski (2001), Kvanvig (2003), Grimm (2010), Hills (2016), Lynch (2017).

3.1 Grasping the structure

The first component of theoretical understanding is the ability to cognitively ‘grasp’ the structure of the object of one’s understanding. Although an illuminating and often used metaphor, there is no consensus on what exactly this ability involves. As Wayne Riggs characterises it, understanding involves ‘a deep appreciation’ or ‘awareness’ of ‘order, pattern’ and how things ‘hang’ or ‘fit’ together, and the ‘role each [part] plays in the context of the whole’.⁸⁶ On this kind of view, cognitively ‘grasping’ the structure of something is equated with ‘a deep appreciation or awareness’ of that structure. Further, the ‘order’ or ‘pattern’ that makes up the structure one is aware of in theoretically understanding something is often taken to be a set of ‘principles,’ ‘theorems’ or ‘propositions’ that constitute a theory or model of whatever it is that one understands.⁸⁷ Given this, we can say that the structure that one is aware of in theoretically understanding something consists in conceptual representations.⁸⁸ It follows that what one is *immediately* aware of (the structure one cognitively grasps) when manifesting one’s theoretical understanding is a theoretical model that (partly) consists in a network of conceptual representations. In many cases, this model is *of* some concrete bit of the world. In which case, we can say that if the model is accurate, then the *ultimate* object of one’s theoretical understanding is the bit of the world the model represents.⁸⁹

However, it is not enough for theoretical understanding in this sense that one is merely aware of the conceptual representations in which the model consists.

⁸⁶ Riggs (2003: 217).

⁸⁷ Kvanvig (2003), Riggs (2009), Grimm (2011), Wilkenfeld (2013), Hills (2016), Lynch (2017); cf. Annas (2011a: 18-19, 1993: 67-8, 1995: 231-3). Note Riggs is explicit that the object of understanding need not be propositional in nature on his view (2003: 218). Similarly, Zagzebski says, ‘a person can know the individual propositions that make up some body of knowledge without understanding them. Understanding involves seeing how the parts of that body of knowledge fit together, where the fitting together is not itself propositional in form’ (2001: 244). However, given that most, if not all, theorists take theoretical understanding to entail the ability to give good explanations in terms of the model, and on the assumption that good explanations express the conceptual content of a thought, it seems that a good deal of the model must be made up of conceptual representations.

⁸⁸ In what follows, I take it that a ‘conceptual representation’ is the part (or property) of (the content of) an explicit thought that is *about* something *and* is in principle expressible in (some) language, i.e., a conceptual representation is an idea *of* something that can in principle be expressed in words.

⁸⁹ There are various ways ‘accurate’ here might be cashed out, e.g., in terms of factivity or functionality.

Rather, one must also see how they ‘hang together’. This involves having a ‘deep appreciation’ of the ways in which different parts of the model depend upon and relate to each other. As Jonathan Kvanvig puts it, what we cognitively grasp in understanding something are the ‘coherence-making relationships in a large and comprehensive body of information’.⁹⁰ This implies that for someone to have theoretical understanding, the model in question—the ‘large and comprehensive body of information’—must be coherent, and further, that one understands the model only if one is aware of what *makes* it coherent, viz., the ‘coherence-making relationships’.⁹¹ For Kvanvig, these relationships include ‘explanatory relations’, ‘logical relations’, and ‘relations of probability’.⁹² Let’s shorten this to just *explanatory relations*.⁹³

In these terms, the explanatory relations of the model one grasps contrast with the relations that hold in the worldly system that is the ultimate object of one’s understanding, and which the model is supposed to represent. Call the relations that hold in the world *dependency relations*. Depending on one’s other theoretical commitments, the content of explanatory relations might be thought to include, for example, *causation*, the central metaphysical relation in play in the natural sciences; *composition*, the metaphysical relation that holds between the parts of a thing and the whole thing; *supervenience*, the modal relation of necessary covariance; or *grounding*, the metaphysical dependence relation that underpins (non-causal cases of) supervenience, and no doubt more.⁹⁴

Scientific understanding is a paradigm case of theoretical understanding in this sense. In the sciences, we often look to conceptual models to make sense of worldly systems that are either inaccessible or too complicated for us to understand in themselves. To give a vivid example, imagine that the world is one vast system of complexly interacting processes. On the working account I am

⁹⁰ Kvanvig (2003: 192).

⁹¹ Just how coherent the model needs to be I leave open.

⁹² Kvanvig (2003: 193)

⁹³ Note that by making this contraction, I do not intend to suggest that ‘explanatory relations’ includes the other two kinds of relation.

⁹⁴ This list is purely an illustration of the kinds of relations that explanatory relations might represent. I have no fixed view about what relations there may be in the world.

proposing, coming to theoretically understand the world would require becoming aware of how the network of dependency relations in which this system of processes consists depend upon and relate to one another, by becoming aware of how the explanatory relations that hold within an accurate conceptual model of the world depend upon and relate to one another. Generalising, we can say that cognitively grasping the structure of some worldly system is equated with being aware of how the explanatory relations that hold within an accurate conceptual model of that system ‘fit together’.

3.2 Manipulating the model

This is still not enough to qualify as having theoretical understanding, however. The second component required for theoretical understanding is the ability to cognitively manipulate the model one grasps in systematic ways.⁹⁵ As Alison Hills puts it, the ability to cognitively manipulate amounts to having a kind of ‘cognitive control’ over ‘the relationship between two propositions’, which includes, amongst other things, ‘knowing how to draw conclusions, give explanations [and] put [the] propositions into your own words’.⁹⁶ Cognitive manipulation in this sense entails an articulate ability. But it is important to note that there may be contingent reasons why an agent cannot express their theoretical understanding in words. For instance, an agent may be unable to do so if gagged and bound by a militant gang of disgruntled students. In such circumstances, we might still want to say that the agent has theoretical understanding even though they cannot articulate it. Thus, it seems the ability to cognitively manipulate a conceptual model comes apart from the actual expression of explanations and predictions, and so on. In such cases, an agent would still qualify as having theoretical understanding in virtue of having the ability to draw good inferences from the model about the system it models, even though they cannot express them. Given this, we can say that the ability to cognitively manipulate a conceptual model is the ability to draw good inferences about the actual and counter-factual consequences implied by the

⁹⁵ See e.g., Grimm (2011), Wilkenfeld (2013), Hills (2016).

⁹⁶ Hills (2016: 3, n. 6). Cf. Grimm (2017).

model and, in favourable conditions, be able to express them.⁹⁷ In this way, the ability to ‘manipulate’ required for theoretical understanding is primarily equated with the ability to draw good inferences about a worldly system from a conceptual model of that system.

Given the above account of what the abilities to cognitively grasp and manipulate amount to, we can stipulate the following formulation of what it is to have theoretical understanding: To count as theoretically understanding some part of the world, *w*, a subject, *S*, must be aware of, and have the ability to draw good inferences from an accurate model, *m*, of *w*. Schematically:

$$\text{TU: } S > m \rightarrow w$$

Here $>$ represents *S*’s awareness of, and ability to draw good inferences from, *m*, and \rightarrow represents accurate representation.

3.3 Articulating explanations

The third component of theoretical understanding is the ability, in favourable conditions, to articulate good explanations about the bit of the world the model represents. Almost all theorists working on theoretical understanding endorse a claim of this nature.⁹⁸ As we have just seen, this ability is often taken to be entailed by the abilities to cognitively grasp and manipulate mentioned above. A related view that, other things being equal, what makes an explanation good is that it provides theoretical understanding is popular in philosophy of science.⁹⁹ If good explanations are what provide theoretical understanding, it makes sense to think that having achieved such understanding the expert will be able to give similar explanations.

Moreover, it seems we are often only happy to judge that someone has theoretical understanding in this sense, if they can articulate such explanations, i.e., they can tell us why certain parts of the world are related to other parts of the world in the

⁹⁷ Cf. Lynch (2017).

⁹⁸ See e.g., Zagzebski (2001), Kvanvig (2003), Elgin (2007), Strevens (2013), Hills (2016), Grimm (2010, 2017). An outlier here is Lipton (2009).

⁹⁹ Kitcher (1981), Grimm (2010).

way they are, given the underlying model.¹⁰⁰ For example, in the natural sciences the adequacy of a theoretical model is often judged in relation to its success in predicting effects in the actual system, given certain changes in initial conditions. As a corollary, theoretical understanding in the sciences is often only attributed to those who can figure out, and then articulate, what effects such changes will produce and why. Of course, for different subject matters there will be different kinds of explanations.¹⁰¹ For instance, in the natural sciences explanations will generally cite causal factors, whereas this will not be so for explanations in mathematics. I will briefly discuss the kinds of explanations relevant to practical activities in the next section. Before that, two final points are worth mentioning.

The first is that the working account of theoretical understanding implies that theoretical understanding is gradable; one can have more or less theoretical understanding, and one can have better or worse theoretical understanding. One way we might cash this out is in terms of how detailed the conceptual model is, and the extent of one's ability to cognitively manipulate it. In these terms, the more detailed one's conceptual model is, the more one theoretically understands of the target phenomenon, and the greater one's ability to draw good inferences from the model, the better one's understanding is. For instance, a history professor will (I would hope) have more theoretical understanding, of a better quality, of, say, how the assassination of Franz Ferdinand led to the outbreak of World War I, than the history undergraduate, because they have a more detailed conceptual model of the events surrounding the assassination,¹⁰² and a greater ability to draw good inferences from the model. There are no doubt other ways to cash out the gradable nature of theoretical understanding. The point for now is just that to theoretically understand something in this sense is not an all or nothing affair, but rather a matter of degree.

The second point is that there is a difference between *having* theoretical understanding and *manifesting* that understanding. For example, consider the expert

¹⁰⁰ Wilkenfeld et al. (2016).

¹⁰¹ The philosophical literature on explanation is vast, but this is not the place to delve into that literature. For a flavour, see Ruben (2004), Leibowitz (2011), Leibowitz & Sinclair (2016).

¹⁰² As Catherine Elgin nicely puts it their 'web of belief is more tightly woven' (2007: 36).

geologist on their day off. It seems clear that when she is enjoying a weekend swim with her children, the expert geologist is not manifesting her theoretical understanding of geology. But it is not as if she somehow loses that understanding on the weekend, only for it to return when she gets back into the lab or office on Monday. Rather, the geology expert manifests her theoretical understanding of geology by *exercising* the constitutive abilities, i.e., by drawing inferences or giving explanations about the bit of the world in question. In order to do this, she must be aware of the relevant conceptual representations and the explanatory relations that hold between them. Thus, we can say that an agent qualifies as *having* theoretical understanding so long as she has the constitutive abilities and, under the right conditions, she can *exercise* those abilities and, thereby, *manifest* that understanding.

In summary, to have theoretical understanding of some part of the world, one must be (potentially) aware of the explanatory relations that hold within an accurate conceptual model of that part of the world, and have the ability to cognitively manipulate the model such that one can draw good inferences about that part of the world from the model. In favourable circumstances, the complex of cognitive dispositions that underlies these first two abilities plausibly entail the ability to give good explanations about the object of one's understanding in terms of the model that represents it. The discussion above has glossed over many nuances and controversies about theoretical understanding. However, for my purposes, this account of theoretical understanding should do. In what follows, I will simply stipulate that these three abilities are the essential constituents of theoretical understanding and use this as a working account of the kind of understanding that corresponds to theoretical expertise in activities like doing mathematics.

3.4 Theoretically understanding an activity

As outlined above theoretical understanding takes as its immediate object a model that consists in a network of conceptual representations, and it takes as its ultimate object the part of the world the model represents (if there is one). In order to

apply this to the argument from understanding, we need to supplement this account slightly. In the context of the argument from understanding, the bit of the world one is modelling is an activity at which it is possible to become an expert. Activities in this sense differ from some other worldly systems in that they have *goals*. Because of this the explanatory relations that constitute a conceptual model of an activity will include teleological relations, e.g., ‘means-ends’ relations. For instance, the goal of golf is to complete the course taking as few shots as possible within the rules of the game.¹⁰³ Given this, for someone to have theoretical understanding of golf they must be able to cognitively grasp and manipulate an accurate conceptual model of the elements of golf as they relate to this goal. This might include, for example, being aware that using a 3-iron is the best way to ensure the ball lands on the green in *this* situation, given the strength and direction of the wind and the distance from the green, etc. If this is correct and generalises, then we can say that to theoretically understand an activity one must be able to cognitively grasp and manipulate a conceptual model of the worldly system in which the activity consists, with reference to the goal(s) of the activity.

On this view, the kinds of explanations someone with theoretical understanding of an activity will be able to give will often (implicitly) refer to the goal of the activity. For instance, if I ask Ted, an expert baker, why he put yeast in his dough, he may simply reply: ‘to make the bread rise’. In this case, the explanation Ted gives just states the goal of the action. However, since explanations have a pragmatic aspect, which explanation is appropriate in any particular situation will depend on features of the context, not least the audience and the purpose of giving the explanation. For explanations of action, the purpose often just is to make sense of the action to the audience. Given this, Ted the baker’s explanation of why he puts yeast in his dough only makes sense of his action if the audience has certain background beliefs, for example, that well-risen bread is pleasant to eat. If Ted had been talking to a small child, he may well have said something like: ‘to make the bread nicer to eat’. And if Ted had been talking to a teenager who wanted to

¹⁰³ Of course, someone might play golf purely for fun, or to get some exercise, etc., but when we attribute expertise to a golfer it is in relation to their ability to complete the course in as few shots as possible, rather than their ability to, say, enjoy themselves.

learn about the chemistry of bread-making for a school project, he might have answered: ‘because yeast releases carbon dioxide into the dough’. None of these three potential answers is false and depending on the context all could be appropriate. Ted could put all three answers together to give a more complete explanation: ‘because the yeast releases carbon dioxide into the dough, which makes the bread rise, and well-risen bread is nicer to eat than dense, heavy bread.’ In this more complete answer, we can see the explanatory links between adding yeast to dough and ending up with tasty bread—the final goal of the activity of bread-making qua activity. Thus, the choice of which part of this answer is offered as ‘the’ reason for Ted’s action of putting yeast in his dough will depend on features of the context.¹⁰⁴

On the proposed working account, then, having theoretical understanding of an activity will entail the ability to give reasons that explain why some action is the appropriate one in the circumstances given the goal(s) of the activity. Given Annas’ emphasis on articulating reasons-explanations as part of the process of learning a skill, and her claim that the practical expert can articulate good explanations in virtue of having a unified grasp of the principles that underlie the activity, I take it that the kind of understanding she is attributing to the practical expert is theoretical understanding in this sense. However, since the empirical evidence presented in the last chapter suggests that when actually manifesting their expertise practical experts often cannot give reasons that explain why the course of action they took was appropriate, even when there is no obvious impediment to doing so, it seems that when acting expertly practical experts are not exercising the abilities that constitute theoretical understanding of the activity. If they were, they should be able to articulate reasons-explanations. If this is correct, then it seems we are not warranted in thinking that what makes practical expertise valuable, i.e., what fulfils SR, is that practical experts theoretically understand the activity. In the next section, I try to motivate the idea that there is a distinctly

¹⁰⁴ Cf. Alvarez (2010, Ch. 6).

practical kind of understanding that can fulfil SR whilst still allowing us to reject AR.

4. Practical understanding

In this section, I discuss two ways we might think our understanding of an activity is practical, both of which appeal to the fact of our embodiment. The first way is by holding that the content of our concepts in general depends on the ways our physical embodiment allows us to interact with the world. The second way is by holding that exercising the physical skills required for an activity can itself manifest an embodied understanding of what one is doing. I suggest that this second kind of embodied understanding is the kind of practical understanding that the practical expert has achieved. In the next subsection, I briefly discuss the first way that our understanding of the world is dependent on our embodiment. I do this both to distinguish it from, and to motivate, the second kind of practical understanding, which I turn to in the following subsection.¹⁰⁵

4.1 Conceptual metaphors

The first way some claim our understanding of the world is practical is by holding that the content of our concepts is generally informed by embodiment and action metaphors. In their work, George Lakoff and Mark Johnson use analysis of a vast range of everyday linguistic expressions to argue that this kind of metaphorical structuring is a ubiquitous feature of human conceptual systems.¹⁰⁶ Call this the *conceptual metaphor view*. It is helpful here to briefly note the way metaphors work. Metaphors are a kind of figurative analogy that work by comparing a *source* domain with a *target* domain, alleging that some relevantly important feature or features belongs to both domains. The idea is that we can elucidate some tricky or obscure

¹⁰⁵ There is a growing literature on ‘embodied cognition’ and the ‘extended mind’ which acknowledges Merleau-Ponty’s views of action and embodiment. See e.g., Varela, Thompson and Rosch (1991), Clark and Chalmers (1998), Clark (2008), and Sterelny (2010). This includes the ‘conceptual metaphor view’ discussed in the main text. The account of awareness in action I develop in the next chapter might be seen as part of this larger movement.

¹⁰⁶ Lakoff & Johnson (1980). Slingerland (2003) makes heavy use of the conceptual metaphor view in his discussion of *wu-wei*.

concept or phenomenon by attending to some more familiar concept or phenomenon. For instance, the phrase TIME IS MONEY expresses a metaphor that brings out, amongst other things, the value of time (target) by identifying it with money (source), which is conventionally conceived of as essentially valuable. Of course, the identification used in such metaphors is a rhetorical device, not meant literally.¹⁰⁷

According to the conceptual metaphor view, the content of many of our more abstract concepts depends on contingent features of our physical bodies, and the ways they allow us to interact in the world. For example, we often talk as if love is a journey, by using locutions such as: ‘We’re at a *crossroads* in our relationship’ or ‘This relationship isn’t *going* anywhere.’ In similar fashion, we often talk as if our minds are containers, saying things like: ‘He’s *out* of his mind!’ or ‘That’s not what I had *in* mind.’ For Lakoff and Johnson, the ubiquitous use of such locutions suggests that the way we make sense of certain abstract concepts (love/minds) is in terms of more familiar concepts from our everyday embodied experience of interacting in the world (taking journeys/using containers). The content of these source concepts thus depends on the contingent ways we physically interact in the world. On this view, in a world where we no longer needed to journey anywhere, we would have a different idea of loving relationships. Thus, the conceptual metaphor view holds that because the source concepts of these metaphors rely on contingent features of our embodied experience, the content of the corresponding target concepts is defined, in large part, by contingent facts about our morphology (bodily structure and physiological make-up) and our physical (action-orientated) capacities and skills.

The working account of theoretical understanding explicated in the last section is an example of the conceptual metaphor view in action. As we saw, epistemologists often invoke the action metaphors of ‘grasping’ and ‘manipulating’ when explicating their notion of theoretical understanding. The use of these metaphors

¹⁰⁷ Another class of figurative analogies are similes which generally draw a comparison explicitly in terms of similarity, as opposed to identity, with the use of terms such as ‘like’ or ‘as’. Another difference between metaphors and similes is that metaphors are generally ‘open’ in the sense that they invite further exploration and development of potential comparisons, whereas similes are ‘closed’ in this sense.

is intended to shed light on the cognitive abilities that are required for theoretical understanding (target) by mapping the familiar abilities of physically grasping and manipulating things in the world (source) into the intellectual realm, where physically grasping and manipulating theoretical models is not actually possible. Thus, the concept of theoretical understanding had by many epistemologists itself appears to be an example of an abstract concept that is contingently dependent on the particularities of our embodied experience.

However, since on the conceptual metaphor view it is the content of our concept of theoretical understanding (and many other concepts) that depends on the way our bodies are and how they enable us to interact with our environment, the view does not really posit a distinctly practical, embodied kind of understanding. At most, what it suggests is that the way we make sense of theoretical understanding qua intellectual achievement is highly contingent on our nature as embodied subjects. Although I find this plausible, it is not the kind of practical understanding that I think corresponds to practical expertise. Nonetheless, the conceptual metaphor view does help motivate the kind of embodied understanding I think corresponds to practical expertise. To see this, we need to consider what might account for the fact that the physical actions of grasping and manipulating in particular help to elucidate the concept of theoretical understanding. On the conceptual metaphor view, abstract concepts are often informed by action and embodiment metaphors, but there are many different embodiment and action metaphors up for grabs. So, why is it that these particular metaphors are illuminating in this case? Why is the elucidatory metaphor not ‘understanding is a container’ or ‘understanding is a journey’—which also work by appeal to our embodied experience?

It seems to me plausible that the reason why these particular action metaphors shed light on what theoretical understanding requires is because the exercise of our abilities to physically grasp and manipulate things in the world can, under the right conditions, themselves manifest an embodied practical understanding of

what one is doing.¹⁰⁸ If successfully grasping and manipulating physical objects does (sometimes) manifest one's practical understanding of what one is doing, this would nicely explain why the actions of grasping and manipulating help to make the concept of theoretical understanding intelligible. If this is right, then positing an embodied kind of practical understanding does interesting explanatory work independently of any role it might play in an account of practical expertise. The inspiration for this notion of embodied understanding comes from the writings of Maurice Merleau-Ponty. I discuss Merleau-Ponty's view of action and perception in detail in the next chapter. In the next subsection, I discuss some examples of the kind of practical understanding that emerges from his view.

4.2 A Merleau-Pontyan framework

Merleau-Ponty's project in *The Phenomenology of Perception*¹⁰⁹ is to argue that human beings are essentially embodied subjects, embedded in their environments. As he puts it at one point, we are directly 'geared into the world' via our embodiment.¹¹⁰ Merleau-Ponty argues that as essentially embodied subjects, embedded in the world our actions have a kind of 'motor intentionality'.¹¹¹ In doing so, Merleau-Ponty appeals to empirical studies and the work of Gestalt psychologists. In his discussion of 'Schneider'—a war veteran who suffered a traumatic brain injury and had some odd pathologies as a result—Merleau-Ponty says:

'In this patient... a dissociation between the act of pointing and the reactions of taking or grasping can be observed: the same subject who is incapable of pointing to a part of his body on command quickly reaches with his hand for the point at which a mosquito is biting him. ... [A]sked to point to a part of his body, such as his nose, [he] only succeeds if he is allowed to grasp it. If the patient is directed to interrupt the movement before it reaches its goal... the movement becomes impossible. It must be admitted [therefore] that 'grasping' ... [is] different from 'pointing'... From its very beginnings, the grasping movement is magically complete; it only gets under way by anticipating its goal, since the ban on grasping is enough to inhibit the movement'.¹¹²

¹⁰⁸ Another reason why these terms might be illuminating is because in exercising our theoretical understanding we are literally cognitively grasping and manipulating abstract objects in a Platonic realm. However, I set this possible explanation aside because it goes against the neo-Aristotelian spirit of my investigation, and does not comply with EC (Chapter 1, §3).

¹⁰⁹ Merleau-Ponty (2013).

¹¹⁰ *Ibid.* p.261.

¹¹¹ *Ibid.* p.113.

¹¹² *Ibid.* pp.105-106.

Merleau-Ponty's mention of 'magic' here is eyebrow-raising, but the important claim in this analysis is that the activity of grasping 'gets underway by anticipating its goal', where (presumably) this is not so for the activity of pointing. This comment suggests that the activity of grasping is teleological in a way that the activity of pointing is not.

In his discussion of Merleau-Ponty's analysis of Schneider, Sean Kelly suggests that Schneider is able to grasp things that he is unable to point at because the activities of grasping and pointing involve two different kinds of understanding.¹¹³ Kelly suggests that pointing at something in order to identify it (as one does when ordered to) is a form of 'demonstrative reference' wherein the thing being identified is understood in contrast to everything else by way of an 'objective cognitive map'.¹¹⁴ Grasping something, on the other hand, Kelly says, only involves identifying it in relation to oneself, a kind of 'egocentric' identification that only requires one to have a 'bodily sensitivity' to it.¹¹⁵ Kelly concludes that there is a distinction between 'essentially bodily understandings' of things in the world, on the one hand, and 'essentially cognitive or reflective understandings' of them on the other.¹¹⁶

'[T]he understanding... that informs... skilful, unreflective bodily activity—activity such as unreflectively grasping a door knob in order to go through the door, or skilfully typing at the keyboard—is not the same as, nor can be explained in terms of, the understanding... that informs... reflective, cognitive or intellectual acts—acts such as pointing at the doorknob in order to identify it'.¹¹⁷

¹¹³ Kelly (2002).

¹¹⁴ *Ibid.* p.380.

¹¹⁵ *Ibid.*

¹¹⁶ *Ibid.* p.378.

¹¹⁷ *Ibid.* p.377. In his discussion, Kelly appeals to Milner and Goodale's (1995) neuro-scientific studies of 'D.F.' who, after suffering carbon monoxide poisoning, exhibited a similar visual pathology to Schneider. As Milner and Goodale describe it:

'D.F.'s ability to recognise or discriminate between even simple geometric forms is grossly impaired... [Her] pattern of visual deficits [however] ... is largely restricted to deficits in form perception. D.F. ... recovered, within weeks, the ability to reach out and grasp everyday objects with remarkable accuracy. We have discovered recently that she is very good at catching a ball or even a short wooden stick thrown towards her... She negotiates obstacles in her path with ease ... These various skills suggest that although D.F. is poor at perceptual report of object qualities such as size and orientation, she is much better at using those same qualities to guide her actions' (1995: 126-128, cited in Kelly 2002: 378).

According to Kelly, Milner and Goodale suggest that the neurophysiological basis for D.F.'s condition is that 'there are two different streams of visual information flow into the brain, one of which is geared to perceptual judgement, the other which is geared directly to action' (Kelly 2002: 379). As we will see, this

On Kelly's reading of Merleau-Ponty, to exercise the ability to physically grasp something manifests an embodied, practical understanding of it in relation to oneself qua grasper, rather than a theoretical understanding of it in relation to everything else. In her discussion of Merleau-Ponty's views on embodiment and skill, Komarine Romdenh-Romluc provides the following helpful example:

'Annette knows that her great-grandfather was a miner. She knows of the long hours he spent deep underground, the cramped conditions in which he worked, the ever-present danger of explosions. One day, however, Annette is out walking when she falls down a pothole. She has to make her way along a narrow tunnel to an opening further down the hill where she can escape into the sunlight. The experience of being trapped below the surface, in cramped, claustrophobic conditions, gives Annette a new insight into her great-grandfather's life. She does not learn any new facts about his existence. But she now has a different appreciation of what it was like to work down the mine'.¹¹⁸

Romdenh-Romluc tells us that:

'For Merleau-Ponty, Annette has gained a bodily 'understanding' of her great grandfather's life and the demands for action that his situation would have made upon him. [...] Annette's bodily grasp of her great-grandfather's situation is an exercise of her motor skills, including those required to move through the pothole to free herself after she is trapped'.¹¹⁹

It is important to note that even though Romdenh-Romluc talks here of 'motor skills', the notion of practical skill she associates with Merleau-Ponty is broader than this term normally implies. As I understand it, 'motor skills', 'habits' and 'practical skills' all lie on a spectrum. On this interpretation, 'skilful actions' include mundane abilities like walking, clicking our fingers, opening doors, biting our nails, making a cup of tea, showering, riding a bike, swimming and driving to work, as well as the elite skill of the master martial artist, Formula 1 driver, expert pianist, master carpenter, etc. This provides a spectrum that increases in skilfulness from clumsy exercise of one's motor skills to expert action.

Following Kelly and Romdenh-Romluc, Merleau-Ponty's views seem to suggest that our skilfully interacting with our environment both requires an embodied

analysis of what is happening in perception is very close to the analysis of awareness in action I give in the next chapter.

¹¹⁸ Romdenh-Romluc (2012: 210).

¹¹⁹ *Ibid.*

understanding of the world, and can provide us with an embodied understanding of the world. Further, in light of the conceptual metaphor view, and the working account of theoretical understanding, we can characterise practical understanding of some activity as requiring i) being aware of the dependency relations in which the activity qua system of the world consists, and ii) having the ability to manipulate that system by interacting with it, and thereby change how the activity of the system unfolds in real time. For instance, in applying back-spin to the ball as she hits it, the skilled tennis player causes the tennis ball to drop shorter than it would have had she not applied back-spin, and since skilful action is intentional in the sense that it is not done accidentally, inadvertently, or unknowingly,¹²⁰ if the tennis player applies back-spin *skilfully*, then she must be (in some way) aware that hitting the ball *in that way* would have the effect of making it drop short (even if she cannot articulate why applying back-spin has this effect).

However, since practical understanding is supposed to account for what is special about practical expertise, having the mere ability to physically manipulate some system is not enough to qualify as a practical expert. For instance, I could manipulate a car engine by bashing it with a hammer, but this would hardly make me a skilled mechanic, let alone an expert one. Rather, the thought is that experts manifest their embodied practical understanding of an activity by manipulating the relevant system *well*, i.e., in line with the goals of the activity. For example, the master mechanic's practical understanding is manifested in their skilful ability to successfully service and repair car engines. If this is right, then practical understanding is fundamentally an embodied achievement, which consists in habituating a certain physical skillset relevant to the goals of some activity. On this proposal, S practically understands some part of the world, w, if S is aware of and can manipulate w, *well*.

PU: S > w

¹²⁰ Cf. Mele (1992:201).

Here, > represents S's ability to grasp and manipulate, w. I submit that this kind of embodied practical understanding of an activity is the defining feature of practical expertise.

In these terms, a key difference between theoretical experts and practical experts in regard to some activity is that theoretical experts understand the activity in an indirect way because their understanding is necessarily mediated by cognitive interaction with the explanatory relations within a conceptual model of the activity, whereas practical experts understand the activity in a direct way because their understanding necessarily involves physically interacting with the dependency relations that hold within the activity (qua system of the world) itself. But these different ways of understanding an activity are not exclusive, one might have both theoretical understanding and practical understanding of an activity. In the maximal case, we can say that such a person has *full understanding* of the activity. Crucially, since practical understanding in this sense does not require cognitively grasping and manipulating a conceptual model, there is no reason to think that having practical understanding of an activity will entail being able to explain anything. This fits with the rejection of AR from the last chapter.

There is more to be said about this kind of practical understanding. However, this will have to wait until I have discussed the third of Annas' claims about practical expertise, i.e., the claim that the practical expert's expert actions result in a flow experience, which is the task of the next chapter. There, I develop the resources needed to fill out the Merleau-Pontyan view of embodied practical understanding that I have outlined in this section. My central claim will be that one manifests one's practical understanding of an activity by exercising the ability to perform it in-flow, where this requires being aware of what one is doing in a very specific way. In the next and last section of this chapter, I discuss several proposals for what is valuable about theoretical understanding and argue that when practical understanding is thought of as the ability to act-in-flow, these ways of being valuable can all be applied to practical understanding.

5. The potential value of understanding

In this section, I briefly discuss three proposals in the contemporary literature about what is non-instrumentally valuable about theoretical understanding. These are that coming to theoretically understand something i) has a distinctively valuable phenomenology, ii) is a creative act, and iii) is a cognitive achievement. Although I think all three proposals have merit, my aim in this section is not to give a comprehensive analysis of any of them, nor to evaluate them in their own right, but rather to suggest that, if cogent, each of them can also be applied to the case of practical understanding.

The first proposal is that theoretical understanding is non-instrumentally valuable because coming to theoretically understand something has a distinctively valuable phenomenology. As Kvanvig puts it:

‘It is the perceived achievement of objectual understanding (though not under that description, of course) that produces the “aha” and “eureka” experiences that provide closure of investigation into the subject at hand. This last point suggests that it is the same thing that is the target of enquiry. It is putting the pieces of the intellectual puzzle together that does that’.¹²¹

The idea here is that coming to theoretically understand something provides ‘a kind of first-person intellectual satisfaction’ at having made sense of some part of the world.¹²² As Stephen Grimm notes, Kvanvig’s use of ‘perceived achievement’ in the quotation above suggests that this kind of intellectual satisfaction applies to the ‘subjective’ or ‘internal’ aspect of theoretical understanding in that we might experience the ‘aha’ feeling of having put all the pieces of the puzzle together, even when the model we end up with does not accurately represent the world.¹²³ For instance, it seems plausible that Georg Ernst Stahl might have felt some legitimate intellectual satisfaction at hitting upon ‘phlogiston theory’ even though it was ultimately proved it false. Given this, the intellectual satisfaction that comes with becoming aware of a coherent model does not entitle us, epistemically speaking, to close down inquiry altogether as Kvanvig’s remark suggests. Nonetheless,

¹²¹ Kvanvig (2011: 89).

¹²² Grimm (2012: 107).

¹²³ *Ibid.*

Grimm concedes that even in such cases, the intellectual pleasure inherent in the subjective aspect of theoretical understanding is ‘legitimate’ insofar as it indicates ‘we have hit upon what looks like a legitimate contender for the truth’ given our own experiences and point of view.¹²⁴

There may well be good objections to this proposal for the value of theoretical understanding. However, as mentioned, my aim is not to evaluate the proposal, but to ask whether, if cogent, such an explanation would transfer over to the case of practical understanding. The answer is yes. On my proposal, one’s practical understanding is manifested by exercising the ability to act-in-flow. An essential characteristic of acting-in-flow is that agents find it intrinsically enjoyable, even when the results of their actions are themselves ‘objectively’ unsuccessful. Sports provide the clearest example of this. For many, the enjoyment of playing a sport is not dependent on the outcome, but comes from playing-in-flow. For instance, many people enjoy playing basketball even when they are losing the game. Thus, if what makes theoretical understanding valuable is that it has a specific kind of satisfying phenomenology, then this value appears to transfer over to the case of coming to practical understand an activity conceived of as being able to perform the activity in-flow.

The second proposal for what is non-instrumentally valuable about theoretical understanding is put forward by Michael Lynch, who claims that coming to theoretically understand something is a creative act.¹²⁵ If one thinks creativity is non-instrumentally valuable, then if coming to theoretically understand something is a creative act, it too will seem to be non-instrumentally valuable in virtue of the creativity involved. It is important to note, though, that the idea on Lynch’s proposal is not that coming to theoretically understand something is a creative act in the sense of being historically original. Rather, it is creative in the sense of being an original act for the person. For instance, consider Layla, who is studying for her GCSE exams. Suppose that one afternoon, after weeks of studying biology textbooks, something ‘clicks’ and Layla can suddenly just see how certain cells

¹²⁴ *Ibid.* p.109.

¹²⁵ Lynch (2017).

convert light energy into chemical energy, and thus understands how photosynthesis works. The thought is that Layla's coming to theoretically understand photosynthesis in this way is a psychologically novel act *for her*, irrespective of the fact that many people have come to theoretically understand photosynthesis before her. The insight involved in theoretically understanding something may not always be sudden, but it does, on this proposal, always constitute a creative act of the agent. On this view, it is the creativity inherent in the act of coming to theoretically understand something that makes theoretical understanding valuable.

Once again, I am not claiming that coming to theoretically understand something constitutes a creative act, or even that creativity is non-instrumentally valuable, though both claims do seem plausible. Rather, my question is just whether this proposal, if cogent, transfers over to practical understanding. And again, I think the answer is yes—in a way. In the case of theoretical understanding the creativity which purportedly makes it valuable is part of the process of coming to understand. However, in the case of practical understanding the creativity involved is part of the exercise of the abilities that manifests one's practical understanding of the activity. This creativity is grounded in the fact that acting-in-flow is itself acting creatively, because to perform an activity in-flow is to spontaneously adapt to the features of a situation in a creative way. For instance, consider an expert guitarist improvising a solo. The idea is that the guitarist's spontaneous riffing is a creative act that results from their ability to play-the-guitar-in-flow. If so, then if what makes theoretical understanding non-instrumentally valuable is that it involves creativity, this value also appears to transfer over to the case of practical understanding.

The third proposal for what makes theoretical understanding non-instrumentally valuable that I will consider is that coming to theoretically understand something constitutes an intellectual *achievement*.¹²⁶ I think this proposal is highly plausible. Indeed, in the proceeding discussion, I have referred to both theoretical

¹²⁶ See e.g., Pritchard (2009).

understanding and practical understanding as achievements. The thought here is that since achievements are themselves non-instrumentally valuable, then if coming to theoretical understanding something constitutes an intellectual achievement, this makes (achieving!) theoretical understanding non-instrumentally valuable. As Duncan Pritchard describes it, in general an achievement is a success because of ability.¹²⁷ However, as Gwen Bradford points out, we do not normally consider every success because of ability an achievement.¹²⁸ For instance, in normal circumstances an able-bodied adult can successfully open a door, and the success is because of their physical abilities, but it seems odd to think that an able-bodied adult opening a door constitutes an achievement. The extra ingredient needed, Bradford suggests, is that the success must be *difficult* where the difficulty of an activity is cashed out as ‘requiring some sufficient degree of *effort*’ and is relativized to an appropriate contrast class.¹²⁹ So we should supplement the account to say that a *genuine* achievement, as opposed to a mere accomplishment, is a difficult success because of ability. Continuing with the example above: on this view, when Layla comes to theoretically understand the process of photosynthesis this qualifies as a genuine achievement because the success (theoretically understanding photosynthesis) is because of her abilities (the abilities to cognitively grasp and manipulate a coherent conceptual model of the process of photosynthesis), and theoretically understanding photosynthesis is difficult in the appropriate sense (many GCSE students studying biology do not theoretically understand photosynthesis).

I think this proposal for what is valuable about theoretical understanding also transfers over to practical understanding. On the account being offered, cultivating practical expertise is the process of coming to practically understand the activity. As we saw in Chapter 2, the process of cultivating practical expertise is long and effortful, requiring dedication and years of practise. Moreover, many

¹²⁷ *Ibid.* pp.22-23.

¹²⁸ Bradford (2015).

¹²⁹ Bradford (2015: 27-28). Bradford’s account is more complicated than I am making out, for instance, rather than the generic ‘because of’ relation Pritchard uses, Bradford has a more complicated notion of the relation involved, which she calls ‘competent causation’ (2015, Ch. 3). However, for my purposes I need not delve into this or the other details of her account here.

people who start out learning a practical skill give up before they become an expert due to the arduous nature of this process. Thus, if what makes theoretical understanding valuable is that it is a genuine achievement, then this value clearly applies to the case of practical understanding.

All three proposals for what is non-instrumentally valuable about theoretical understanding I have considered appear to transfer over to practical understanding conceived of along Merleau-Pontyan lines as having the ability to act-in-flow. Further, it is worth noting that even if we reject all these proposals for what is non-instrumentally valuable about practical expertise, being able to perform well in familiar circumstances seems clearly instrumentally valuable because doing so helps the agent to achieve the goals of the task at hand. My claim is that part of what makes the practical expert special, as compared to the perfect imitator, is the additional ability to perform well in *un*familiar circumstances, which is itself made possible by the practical expert's ability to perform the activity in-flow. In the next chapter, I try to make good on this claim by arguing that the ability to act-in-flow enables the expert to spontaneously adapt to unfamiliar situations because when acting-in-flow an appropriate course of action is always presented and taken up. If this is right, then since having the ability to perform an activity in-flow constitutes having practical understanding of the activity, having practical understanding of the activity makes having practical expertise instrumentally valuable. Given this, we can say that the fact that the practical expert has come to practically understand the activity is apt to fulfil the *special requirement* (SR).

6. Conclusion

In this chapter, I have examined Annas' claim that coming to understand an activity is generally required for someone to cultivate practical expertise in it. According to the argument from understanding, coming to understand an activity entails AR. This is a problem because the empirical evidence presented in the last chapter suggests that AR is false. However, we can dissolve the problem by distinguishing between:

2a) S has theoretical expertise of ϕ -ing only if S has come to theoretically understand ϕ -ing.

And:

2b) S has practical expertise of ϕ -ing only if S has come to practically understand ϕ -ing.

While 2a) is in conflict with the rejection of AR, 2b) is not. By appealing to Merleau-Ponty's views on embodiment, I have suggested that the practical expert manifests their embodied understanding of an activity in exercising their ability to act-in-flow. In the next chapter, I fill out the details of this suggestion.

5. The phenomenology of expert action

1. Introduction

In the last chapter, I argued that coming to *theoretically* understand an activity is *not* generally required for cultivating practical expertise in that activity (though it need not hinder and may sometimes help)—at least, for the majority of activities that Annas mentions, viz., playing the piano, pottery, rock-climbing, dancing, ice-skating, playing tennis, playing golf, playing football, throwing the javelin, running, athletics, speaking a second language and driving. Further, I suggested that there is a *practical* way of understanding an activity that is distinct from theoretical understanding, and that having this practical kind of understanding is what distinguishes the expert from the non-expert who can perfectly imitate them. That is, having practical understanding is what fulfils the special requirement on an account of practical expertise. In this chapter, I develop the resources needed to fill out this conception of practical understanding.¹³⁰ In order to do so, I first need to examine the third claim I attributed to Annas, namely:

- 3) S has practical expertise of ϕ -ing only if when S ϕ 's to the best of their abilities S ϕ 's-in-flow.

For the purposes of my project, Annas' appeal to the practical expert's ability to act-in-flow in delineating her skill model of virtue is highly salient. This is because scholars of early Daoism often draw an analogy between Zhuangzi's notion of *wu-wei* and the ability to act-in-flow. The ability to act-in-flow thus provides a conceptual bridge between Annas' account of virtue and contemporary accounts of *wu-wei* via the analogy with practical expertise. Examining this ability is, therefore, central to understanding the nature of ethical expertise. This is the task of this chapter. I will argue that 'the standard view' of human action in the

¹³⁰ Versions of §§ 3-5 of this chapter appear in Lee (forthcoming).

contemporary literature struggles to account for the practical expert's ability to act-in-flow, but that, when suitably extended, Merleau-Ponty's view of action does have the resources to account for the ability to act-in-flow. All else being equal then, we should prefer the Merleau-Pontyan view over the standard view of action when giving an account of practical expertise.

The chapter proceeds as follows. In §2, I briefly outline the standard view of action in contemporary philosophy and explain why this view struggles to account for the ability to act-in-flow. In §3, I set out an alternative Merleau-Pontyan framework of action, and in §4 I apply this framework to ordinary everyday cases of skilful action, distinguishing different ways we can be aware of what we are doing. In §5, I apply this distinction to a case of acting-in-flow and suggest that acting-in-flow is essentially characterised by *attentive engagement* with one's environment, where this is a specific way of being aware of what one is doing. In §6, I use the idea that there are different ways we can be aware of what we are doing to distinguish between practical knowledge and practical understanding, and in §7 I conclude that having the ability to perform some activity in-flow is to have practical understanding of that activity in the sense that exercising one's ability to act-in-flow provides a specific kind of epistemic access to the world.

2. The standard view of action

In this section, I briefly outline what I call 'the standard view' of action in the contemporary literature, and explain why this view struggles to account for the practical expert's ability to act-in-flow. According to the standard view, the defining feature of an action is that it is caused by an intention, i.e., a mental event that conceptually represents the performance of the action. On the standard view, it is the presence of an intention that distinguishes an action—something we *do*—from something that merely happens to us. There are broadly speaking two kinds of account that fall under the standard view: *event-causal* accounts and *agent-causation* accounts. Event-causal accounts can be characterised as holding that some bodily movement qualifies as an action if, and only if, it is caused (in the right way) by an

intention, which is itself caused by an appropriately related belief-desire pair of the agent.¹³¹ The ‘in the right way’ qualification is needed to rule out deviant causal chains, whereby an intention causes a bodily movement in such a way that we do not think it qualifies as an action. I will return to this qualification in a moment. For now, the point I want to emphasise is that on event-causal accounts for a bodily movement to be an action, it must be caused (in the right way) by an intention of the agent.

Event-causal accounts utilise a standard scientific conception of causation wherein inherently passive states and events are the relata within a dynamic causal network, i.e., causation is an external relation. On this way of understanding action, I desire that x be the case, and I believe that my ϕ -ing will bring about x, this belief-desire pair causes me to form the intention to ϕ , and the intention to ϕ causes my ϕ -ing. Schematically:

$$[B+D] \rightarrow [I] \rightarrow [A]$$

Here, B = belief, D = desire, I = intention, A = action, \rightarrow = dynamic causal link, and the square brackets enclose inherently passive states or events. On this view, my belief and desire *rationalise* my intention, i.e., make sense of my intending to ϕ , and so my ϕ -ing. As such, the belief-desire pair constitute my reason for ϕ -ing.¹³² And my ϕ -ing counts as the action in this sequence just because of the rationalising belief-desire pair that feature in its event-causal history.

There are two main objections to event-causal accounts. The first concerns the ‘in the right way’ qualification, i.e., the possibility of deviant causal chains. An example here is the case of the assassin whose decision to shoot her ‘mark’ is so unnerving that it causes her trigger-finger to twitch, causing her rifle to fire, and thereby kills her target.¹³³ In this case, the intention to shoot her target caused the assassin to

¹³¹ Velleman (1992: 461) calls this kind of event-causal account the ‘standard story’, and Wiland (2012: 17) calls the view that reasons for action ‘are psychological attitudes’ the ‘standard view’. I extend this to include agent-causation accounts, because I take it that what is standard about both kinds of account is that they give *causal* explanations of action where the explaining factor is a mental event of some kind.

¹³² Davidson (1963).

¹³³ This example is adapted from Romdenh-Romluc (2011b: 80). Cf. Davidson (1973: 152-153), Mele (1983: 346), Velleman (1992: 463).

shoot her target, but in the wrong way. That is, it seems odd to think that a twitch counts as an action. The second objection to event-causal accounts is that they leave no room for agency, but instead construe human actions as simply another kind of passive event within the causal nexus, i.e., in ‘acting’ nobody ever *does* anything.¹³⁴

Given such objections, some theorists find event-causal accounts of action untenable and advocate an alternative account. These are commonly known as *agent-causation* accounts.¹³⁵ According to agent-causation accounts agents play an ineliminable role in producing action. On this kind of view, the agent has an irreducible causal influence on the world, i.e., they are a source of causal power. Instead of an action being just another passive event within a chain of external causal relations, agent-causation accounts hold that there is a fundamentally different kind of causal relation at play in human action, an internal causal relation between intention and action. Utilising the same schema as above:

$$[B+D] \rightarrow [I \Rightarrow A]$$

Here \Rightarrow indicates agent-causation. On this view, I desire that x be the case, and I believe that my ϕ -ing will bring about x , this belief-desire pair causes me to form the intention to ϕ , but my intention to ϕ internally causes my ϕ -ing. That is, “*I*” am part of the cause. According to agent-causation accounts, my intention to ϕ and my act of ϕ -ing cannot be separated into distinct, passive events. The action in this case is constituted by my ϕ -ing *and* my intention to ϕ , i.e., by both a mental event and a bodily movement. In this way agent-causation accounts make room for agency and can exclude ‘actions’ caused in a deviant way. The assassin’s nervous twitch causing the rifle to fire does not include the assassin *herself* as a cause of the firing, and so does not count as an action. However, even though agent-causation accounts purport to provide a place for the agent in action and exclude deviant causal chains, it is difficult to make sense of how an intention ‘internally causes’ an action and of how an agent can be a ‘source of causal power’.

¹³⁴ Nagel (1986: 114).

¹³⁵ See e.g., Chisholm (1976), O’Connor (2000).

What exactly is internal causation? And how does it differ from the standard scientific external conception of causation? Does being a source of causal power require being outside the causal nexus? And if so, how do we make sense of the idea that something outside the causal nexus can causally affect things within the nexus? Without satisfying answers to such questions, agent-causation views leave us with an explanatory lacuna that is arguably more worrying than the problem they set out to solve.¹³⁶

This is just a quick sketch of the two kinds of account that fall under what I'm calling the standard view of action in contemporary philosophy. Many interesting details have been omitted, but for my purposes this simple sketch should do. Neither kind of account is without its specific problems.¹³⁷ However, the crucial point is that both kinds of account hold that something is an action if, and only if, it is caused (in the right way) by an intention that conceptually represents the performance of the action. That is, on the standard view, actions are caused by our thoughts about what to do. Now, when it comes to accounting for the practical expert's ability to act-in-flow this is problematic. This is because thinking about what to do is inimical to acting-in-flow. For instance, in his autobiography ex-international cricketer Ken Barrington says the following:

'When you're playing well you don't think about *anything* and run-making comes naturally. When you're out of form you're conscious of needing to do things right, so you have to think first and act second. To make runs under those conditions is mighty difficult.'¹³⁸

And, similarly, a US championship basketball player tells us:

'When I get hot in a game... Like I said, you don't think about it *at all*. If you step back and think about why you are so hot all of a sudden you get creamed.'¹³⁹

Csikszentmihalyi summarizes this aspect of acting-in-flow:

'Typically, a person can maintain [flow] for only short periods interspersed with interludes [...] These interruptions occur when questions flash through the actor's mind

¹³⁶ Taylor (1992: 53), Clarke (2003, Ch.10).

¹³⁷ See Mele (1992) for discussion of several more issues with such accounts.

¹³⁸ Barrington (1968: 97, emphasis added).

¹³⁹ Reported in Csikszentmihalyi (2014: 139), emphasis added

such as “Am I doing well?” or “What am I doing here?” or “Should I be doing this?” When one is in a flow episode..., these questions simply do not come to mind.”¹⁴⁰

When an agent acts-in-flow they are acting intentionally in that their actions are not done unknowingly, inadvertently or accidentally.¹⁴¹ But, as these comments suggest, even though they are acting intentionally the agent is not thinking about what they are doing in the sense of conceptually representing the performance of their actions. Therefore, because the standard view claims that actions are caused by intentions that conceptually represent the performance of the action, it struggles to account for the practical expert’s ability to act-in-flow. It is important to note at this point, that since the issue is with thinking about what to do, the objection to the standard view generalises to what Maria Alvarez calls the ‘minority view’ of action, which eschews causal explanations and instead holds that citing the reason consideration of which motivated the agent to act is sufficient to explain their action.¹⁴² This is because on the minority view too, actions necessarily involve thinking about what to do in the sense of conceptually representing the performance of the action.

There might seem to be an obvious reply available to the both the advocate of the standard view and the advocate of the minority view here, i.e., that the intentions involved when acting-in-flow are ‘implicit’ in the sense that they are below the surface of conscious experience. According to this reply, when acting-in-flow the practical expert’s actions are still caused or motivated by thoughts that conceptually represent the performance of their actions, it is just that the agent is not consciously aware of them. However, there is the question of why we should think that there are such implicit thoughts in the first place? After all, parsimony suggests that, all else being equal, if we have no evidence for something, then we should not commit to its existence. Since implicit thoughts are below the surface of conscious experience, we have no evidence for them. So, parsimony suggests that, all else being equal, we should not commit to the existence of implicit thoughts. An obvious reply here is that we should think there are such implicit

¹⁴⁰ Csikszentmihalyi (2014: 138).

¹⁴¹ Mele (1992: 201).

¹⁴² Alvarez (2010: 2, Ch.5).

thoughts because when we ask the agent why they did what they did, they can tell us the reasons.¹⁴³ However, as we saw in Chapter 2, it is often the case that practical experts cannot tell us why they acted as they did, so this reply is not convincing.

A better reply is that we should think there are such implicit thoughts because otherwise we have no way of explaining the expert's actions, and parsimony allows that we can posit the existence of entities for which we have no evidence if those entities are needed to explain some phenomenon. That is, advocates of the standard and minority views of action can try to justify the claim that when the practical expert acts-in-flow their actions are caused or motivated by implicit thoughts by inference to the best explanation. However, this reply would only be convincing if there are no better alternative explanations available. In the next section, I outline Merleau-Ponty's view of action, which does not require us to conceptually represent the performance of the action in thought for it to count as an action. My claim is that given the availability of this alternative explanation, we need not commit to the claim that when acting-in-flow the agent's actions are caused or motivated by implicit thoughts.

3. Merleau-Ponty's view of action

In trying to account for the practical expert's ability to act-in-flow, both the standard view and the minority view need to appeal to implicit conceptual representations causing or motivating the agent's actions. In this section, I outline an alternative view of action which does not require the presence of intentions, and so does not need to appeal to implicit conceptual representations. This is a view of action based on the work of Merleau-Ponty, and developed by Dreyfus and Romdenh-Romluc, amongst others.¹⁴⁴

¹⁴³ This is the kind of response Annas (2011a: 29) gives. I discuss this in more detail in the next chapter.

¹⁴⁴ Merleau-Ponty (2013), Dreyfus (2000, 2002, 2005), Romdenh-Romluc (2011a, 2011b, 2012, 2014, 2015), Kelly (2002), Webber (2002), Berendzen (2010), Allen (2018).

According to Merleau-Ponty, actions are essentially brought about as the result of *perception*. This is made possible because perception presents the world to us as offering opportunities to act. That is, we see features of the world not as objective things with objective properties, but rather ‘in terms of the value they have for our behaviour’.¹⁴⁵ On this view, an agent might directly perceive, say, her car as offering the opportunity to drive somewhere, or a glass of water as offering the opportunity to quench her thirst. But importantly these are not judgements she makes on the basis of seeing the car or the glass of water; she directly perceives them as offering her a way to act—i.e., she sees the car as for-driving and the water as for-drinking.

There are two key notions that need to be distinguished on this view of action. The first is the notion of an *affordance*. Affordances are the possibilities for action that actually exist in the world.¹⁴⁶ Which aspects of the world afford an agent the opportunity to act will depend on features of the agent, i.e., their morphology and motor capacities, as well as the circumstances that pertain. In this way, a car affords the opportunity to drive somewhere to an adult who is able to drive, but not to a ten-year-old, or someone who cannot drive. Affordances are objective in that there is a fact of the matter about whether or not some feature of the world affords an agent the opportunity to act. But they are also subjective in that whether some feature of the world affords an agent the opportunity to act will depend on features of the particular agent, e.g., their morphology and motor capacities. Moreover, some feature of the world may afford some agent the opportunity to act without them ever perceiving that affordance.¹⁴⁷

The second key notion is the notion of a *solicitation*. This concerns how affordances are actually experienced by an agent.¹⁴⁸ In the right conditions, agents experience

¹⁴⁵ Romdenh-Romluc (2012: 199).

¹⁴⁶ Gibson (1977), Dreyfus (2005: 55-56), Romdenh-Romluc (2012: 199).

¹⁴⁷ It is worth noting that an agent’s perception of some feature of the world *as* an affordance might be mistaken. For example, when the tasty looking mushroom I find in the woods does not actually afford a viable source of food because (unbeknownst to me) it is poisonous. Thus, there is an important difference between an affordance, the perception of something *as* an affordance, and the ‘veridical’ perception of an affordance. In what follows, I will ignore the complication of the possibility of mistaken perception of a feature as an affordance, and will limit my talk to veridical perception of ‘real’ affordances that present the agent with genuine opportunities to act.

¹⁴⁸ Merleau-Ponty (2013: 222), Dreyfus (2005: 55-56), Romdenh-Romluc (2012: 199).

affordances as soliciting them to take up the opportunities they offer—i.e., as inviting them to act. If I am thirsty, then simply seeing a glass of water draws me to drink it. And, if circumstances are favourable, my perceiving the glass of water will solicit the action of drinking. Solicitations can vary in the strength or urgency with which they draw an agent to act. The more salient an affordance is to the agent, the more urgently perception of that affordance will solicit them.¹⁴⁹ If I am very thirsty, seeing a glass of water will draw me more urgently than if I am not very thirsty. And seeing a glass of water might not solicit me at all, if I have just had a large drink. In some cases, an affordance will be so salient to an agent that they will, if they perceive it, experience it as requiring or even demanding that they take up the opportunity it offers. For example, if I see my daughter is about to drink from a bottle of bleach, I will experience the situation as demanding that I stop her.

On the Merleau-Pontyan view, then, intentions are neither necessary nor sufficient for action. But the view does allow that an agent's thoughts and intentions do in many cases have a big impact on which actions they perform.¹⁵⁰ There are at least two ways this might happen. First, our thoughts and intentions can affect which actions we perform by affecting which affordances become salient for us, and the urgency with which they solicit us to act. For instance, imagine I am trying to decide whether to go for a drive or go for a run. If I decide to go for a drive, my intention to go for a drive makes affordances relevant to driving more salient than affordances relevant to running; seeing my car keys draws forth the action of picking them up, whereas seeing my running shoes draws forth no action. Second, our thoughts can help initiate and guide our actions by allowing us to conceptually represent opportunities to act that we don't actually

¹⁴⁹ Mark Schroeder (2007) develops a notion of 'salience-striking' whereby when a consideration strikes an agent as salient to the promotion of their desires; this constitutes the agent 'being in a state with the content that that consideration is a reason for [them]' to act (2007: 174, cf. pp.156-158, 161-162). Schroeder's view is complex, and I won't discuss it here further than to note that insofar as Schroeder holds that intentional action requires considering the conceptual content of our thoughts, it appears to be a version of the minority view discussed above. As I discuss below, the Merleau-Pontyan framework is generally compatible with the claim that considering reasons might prompt or interrupt episodes of acting-in-flow.

¹⁵⁰ Romdenh-Romluc (2012: 206-208).

perceive. I will discuss an example of this in detail the next section. So, even though intentions are by themselves neither necessary nor sufficient for action, often perception and thought *jointly* bring about action.

On the Merleau-Pontyan view, an agent perceives the affordances of her environment and immediately responds to those affordances that solicit her most strongly. This is made possible by the agent's capacity for skilled behaviour. As we have seen, becoming skilled at some activity involves becoming more and more familiar with that activity by repeated practise. The more familiar we are with an activity the more we see the world in terms of affordances relevant to the performance of that activity. For a pianist, a piano affords the opportunity to play it, whereas it does not afford a non-pianist the same opportunity. And, if the pianist is in the mood to play, then her perception of the piano will solicit her to play it. In this way, the practical skills we have inform how we perceive the world, and in turn which features of the world solicit us to act.

Although we have been generally concentrating on the practical expertise of athletes, musicians, artists, nurses, plumbers, etc., as mentioned in the last chapter, the Merleau-Pontyan conception of skills is far broader than just these kinds of 'elite' skills. In this broader sense we are all skilled at some things, whether it is driving, dressing or simply walking. These are all practical skills on this view because they take repeated practise to master. Of course, many of us won't remember developing the skill of, say, walking, but as many parents will attest, learning to walk takes practise. As such, the ability to walk is a practical skill. Some of these more mundane skills are such that they can be elevated to an elite level. Driving is an example of such a skill. Although many people can drive quite well in typical circumstances, there are situations in which it takes a far greater degree of expertise to drive well, e.g., in snow, on rough terrain, or in Formula 1. But the fact that driving skill admits of degree does not diminish the fact that even the ordinary everyday ability to drive is a practical skill. The mechanism by which our actions manifest in such cases is our capacity for skilled behaviour.

The appeal to skills makes the Merleau-Pontyan framework apt for making sense of the practical expert's ability to act-in-flow. However, it does raise the question of what distinguishes acting-in-flow from ordinary skilled action on this view. My proposal is that we can start to answer this question by distinguishing two 'layers' in our awareness of what we are doing. On this proposal, action drawn forth by perception is action that is drawn forth by the way we are aware of our environment, where this includes ourselves. If there are different ways we can be aware of our environment, then this opens up space for there to be different ways perception can draw forth action. If so, this gives us a way to distinguish acting-in-flow from ordinary skilled action. In the next section, I discuss three examples which together strongly suggest that there are different layers of awareness in action, which can lead to different ways we might act. In the following section, I apply this layered conception of awareness to a case of acting-in-flow.

4. Ways of being aware

In trying to understand how acting-in-flow is a way of acting, and how it differs from ordinary skilful action drawn forth by perception, we can draw a distinction between two layers of awareness in action. In normal everyday experience, these layers of awareness blend together seamlessly such that the distinction between them is all but undetectable to the agent. The plausibility of such a distinction is illustrated by considering examples.

Case 3: The distracted driver

Jane has had a hard day at work. At the end of it, she packs up her things and heads out of the office. As Jane is walking towards her car, she is thinking about the events of the day. She gets into her car and starts to drive home. As she drives, Jane reflects on the meeting she just had, and the things she needs to do tomorrow, in preparation for a big presentation next week. Even though her mind is preoccupied with thinking about these various things, still she manages to drive safely home, navigating roundabouts, stopping at red-lights, avoiding cyclists, and

so on. Eventually, Jane finds herself at her destination, and with a slight sense of concern she realises she has no recollection of the journey.

Reflecting on this example suggests that there are two layers of awareness in action. In Case 3, Jane must be aware of the salient features of her environment such that she is able to respond to them appropriately and make it safely home. But even though Jane must be aware of her surroundings, this is not what she is thinking about whilst she is driving. Instead, it is as if she is driving on ‘autopilot’ and not paying attention to what she is doing. She can do this because she has developed the habit of driving this particular route home. Thus, Jane’s driving home is manifested through her *habitual* awareness of what she is doing. Being habitually aware of her environment is how Jane is able to perceive and respond to features salient to driving home. On the Merleau-Pontyan view of action, we can say that habitual awareness of these features directly solicited the appropriate actions, without the need for Jane to conceptually represent any of her actions in thought. This explains how she can drive home whilst thinking intently about work, and why she can’t remember the journey. Whilst Jane is habitually aware of her environment, what she is actually paying attention to are her thoughts of work. If to pay attention to something requires being aware of it, then in Case 3, Jane must also be aware of her thoughts about work. Being attentively aware of her thoughts about work is how Jane is able to reflect on her meeting, and make plans accordingly. If thinking, planning and reflecting are things we *do*, then on the Merleau-Pontyan view of action, we can say that Jane’s thinking about work is manifested through her *attentive* awareness of what she is doing. The suggestion, then, is that Jane’s driving actions are brought about by her habitual awareness of her environment, whilst, simultaneously, her thinking is brought about by her attentive awareness of the content of her thoughts.¹⁵¹ Let’s consider a contrasting case.

¹⁵¹ Compare McClelland (2015, forthcoming).

Case 4: The learner driver

Allison is taking her third driving lesson. She is listening carefully to the instructions of the driving instructor and trying to do everything the instructor is telling her. She concentrates hard on coordinating depressing the clutch and putting the car into gear. Then, she concentrates on easing the pressure off the clutch slowly, whilst at same time applying a small amount of pressure with her other foot to the accelerator. All the while, Allison is consciously repeating the instructions she has been given back to herself. She does this throughout the whole lesson: *now* change gear, *now* indicate, *feed* steering wheel, *check* mirrors, *ease* brakes, and so on. By the end of the lesson Allison is exhausted, having expended so much mental effort concentrating on what she was doing.

Applying the distinction in awareness to Case 4, we find that Allison's driving actions are produced by her being attentively aware of what she is doing. Because she is not yet skilled at driving, she must continually think about what to do next.¹⁵² If Allison had let her attention drift onto something else, she would soon have found herself in trouble, stalling the car or worse. Instead, she had to concentrate on every tiny movement she made, telling herself what to do every step of the way. This suggests that Allison's driving actions are initiated and guided by her being attentively aware of certain driving-related affordances by conceptually representing those affordances in thought. This contrasts with Case 3, wherein Jane's driving actions are initiated and guided by her habitual awareness of her environment. Let's expand Case 3.

Suppose that after she got home, Jane was a bit upset that she had driven home without paying attention to what she was doing. Suppose she made a cup of tea, and in an effort to make herself feel better, tried to remember her journey. After considering it carefully, she can recall glimpses of the journey home, e.g., the fluorescent yellow waterproofs a cyclist was wearing, and negotiating a particularly busy roundabout. But this is all she can come up with—fleeting, blurry memories—not nearly enough to constitute a clear narrative of the journey.

¹⁵² On the Dreyfus model (Chapter 3, §2) Allison is at the stage of novice.

Considering these cases suggests that during sustained activity there can be a merging of layers of awareness. If during a routine habitual activity something unfamiliar, or particularly complicated happens, we pay attention to what we are doing. This allows us to cope with the unfamiliar or complicated feature of the activity. For instance, at the busy roundabout Jane is still habitually aware of her surroundings, and her driving actions (taking off the handbrake, applying pressure to the accelerator pedal, etc.) are still drawn forth by this habitual awareness. But she is also attentively aware of her environment, and *decides* when to go, after the green car (say) by representing ‘the-gap-after-the-green-car’ as offering her the opportunity to go. In this way, Jane is attentively aware of an affordance that she is not habitually aware of. The merging of habitual and attentive awareness in such cases might be so subtle that we don’t even notice it happening. This would explain why Jane can, with some effort, recall the busy roundabout and cyclist’s waterproofs, as these features drew her attention away from her thoughts for an instant. It would also explain why, at the time, it seemed to Jane as though there had been no break in her thinking at all. After the unfamiliar or complicated bit of the task had been successfully negotiated, Jane’s attentive awareness smoothly returned to her thoughts of work, without her noticing they had been interrupted. It is as though Jane’s stream of thought paused mid-flow while she momentarily focused her attention to deal with an unfamiliar feature of her task, then re-started once she was back on familiar territory and no longer needed to pay attention to her driving.

We can further bring out this feature of awareness in action by re-imagining the case such that roadworks mean that Jane cannot take her usual route home and must instead follow a diversion. In this case, the unfamiliarity of the route will require more sustained merging of layers of awareness. Once she has completed the diversion, and is back on her familiar route, Jane might need to remind herself of what she was thinking about: “oh yes, next week’s presentation.” In this version of events, the interruption to Jane’s thoughts of work is significant enough to be clearly noticeable.

Insofar as habitual awareness and attentive awareness can be simultaneously directed at different objects or activities, they appear independent. However, since awareness is ‘transparent’—in the sense that we are (typically) only aware of the objects of experience, and not our *awareness* of those objects *itself*—when both layers are directed at the same object (in this context, the activity of driving) they merge together. However, the layering of Jane’s habitual and attentive awareness at the roundabout does not amount to a complete or total merging. This is because Jane’s attentive awareness of some of the relevant affordances is mediated by her thoughts. In this way, she is *indirectly* attentively aware of those affordances. In the next section, I consider a case of acting-in-flow where there appears to be a complete merging of layers of awareness in that both layers of awareness are directed at the same activity, but the agent is not conceptually representing any affordances in thought. Before that let’s consider a third case, where it seems habitual awareness is operative without attentive awareness.

Case 5: Sleepwalking

Peter is prone to bouts of sleepwalking. One night, whilst still asleep, Peter gets out of bed, puts on some clothes, opens the bedroom door and walks downstairs to the kitchen. Once there he proceeds to make and eat a cheese sandwich. At no point does Peter wake up. After he has finished his sandwich he lies down on the sofa. In the morning Peter awakes to find himself on the sofa, his shirt on inside out, but with no memory of what he did or how he got there.

In this case Peter’s ‘actions’ are the result of his habitual awareness. He was able to perform routine tasks like dressing, opening the bedroom door, walking downstairs, and making a sandwich—whilst asleep. As he was asleep, Peter was not attentively aware of what he was doing and consequently he had no memory of the episode. Nevertheless, insofar as he could navigate doors and stairs etc., Peter was able to respond appropriately to features of his environment, albeit in limited ways. If in order to respond to something, we must be aware of it then, in Case 5, Peter must have been aware of his environment. If he were not aware of, e.g., the door handle, he would not have been able to open the door. Thus, even

though he is ‘unconscious’ Peter is still habitually aware of his surroundings, and it is this layer of awareness that draws forth his behaviour. Of course, sleepwalking is an exceptional case, and we do not normally attribute agency to the sleepwalker.¹⁵³ But, taken with the examples above, Case 5 strongly suggests that rather than one fully integrated mode of awareness, there are distinct layers to our awareness of the world that can and do come apart.

5. Attentive engagement

Cases 3 - 5 suggest a distinction in the way we can be aware of what we are doing. I think we can use this distinction to help distinguish acting-in-flow from ordinary skilled action. Let’s consider another example.

Case 6: The expert driver

Ken is a professional driver with many years of experience driving in Formula 1. Having raced this particular track many times, Ken knows the layout intimately, and instinctively slows down, changes gear, and so on, as the lay of the track demands. Whilst driving Ken is not thinking about his actions, or the movements of the car, or anything else. Instead, he is completely absorbed in the activity of driving. His senses are perfectly attuned to the task of getting around the track as quickly (and safely) as possible. When he sees the opportunity to overtake he doesn’t decide to do so, but immediately responds by accelerating into the gap he has seen.

In Case 6, Ken is acting-in-flow. On the layered conception of awareness, we can say that Ken’s actions are initiated and guided by his awareness of the relevant affordances. However, (unlike Case 3) Ken is not preoccupied with thoughts of other things, nor (unlike Case 4) is he thinking about what he is doing, but (unlike Case 5) he is both habitually aware and attentively aware of what he is doing. Like Jane at the busy roundabout, both Ken’s habitual awareness and his attentive

¹⁵³ For this reason, killing someone whilst sleepwalking would not be murder as there would be no *mens rea*.

awareness are directed at the same object, i.e., driving. As such, Ken's habitual and attentive awareness have merged. However, unlike Jane, Ken's attentive awareness of what he is doing is not mediated by way of conceptual representations of relevant affordances. Since Ken's attentive awareness of what he is doing is not mediated by the content of his thoughts, he is *directly* attentively aware of what he is doing. Thus, in Case 6, there is a complete merging of layers of awareness. Call this being *attentively engaged* in what one is doing. In this way, Ken is attentively engaged in driving. This is what distinguishes Ken's performance from Jane's and Allison's: Jane's actions are a form of distracted-habit, Allison's a form of conscious-striving, whereas Ken is acting-in-flow. By generalising, we can say that when acting-in-flow the practical expert's actions are drawn forth by their attentive engagement in what they are doing.

On this proposal, human behaviour is partly the result of explicitly considering what to do, and partly the result of being directly solicited by awareness of salient affordances. Our thoughts, feelings, commitments, values and goals can all affect which affordances become salient to us. But most of our actions come about without our explicitly thinking about what to do, e.g., whenever we are acting habitually while thinking about something else, or when we are acting-in-flow.¹⁵⁴ On the layered conception of awareness, routine habitual action is brought about by habitual awareness of the relevant opportunities to act, and acting-in-flow is brought about by attentive engagement in the activity.

The layered conception of awareness in action provides a phenomenologically grounded account of the ability to act-in-flow, without the need to appeal to implicit thoughts and intentions that cause or motivate the agent's actions. Thus,

¹⁵⁴ Barbara Montero (2015) argues that practical experts can and do often 'think in the zone' and that this often aids performance. Montero cites evidence that suggests that practical experts often make aesthetic judgements about their body position during a performance and, in stressful situations, consciously re-double their efforts and focus on task, and that these are both ways of thinking that can help performance. I think the account I am offering is compatible with Montero's claims. On the layered conception of awareness, we can say that while the expert is having such thoughts they are (at least momentarily) not attentively engaged in what they are doing. In the first case, altering one's technique in response to an aesthetic judgement about one's movements can be seen as action drawn forth by attentive awareness of an opportunity to improve one's technique by conceptually representing the movement required. In the second case, consciously increasing one's focus and efforts on task can be seen as an attempt to attentively (re)engage in the activity.

suitably extended, Merleau-Ponty's framework provides a more parsimonious account of the practical expert's ability to act-in-flow than the standard or minority views of action discussed above. Given this, then, all else being equal, when giving an account of practical expertise we have reason to prefer the Merleau-Pontyian inspired view of action over the standard or minority views. Moreover, since the claim that when manifesting their expertise practical experts act-in-flow is empirically robust, if the expanded Merleau-Pontyian view of action I have argued for above best explains the ability to act-in-flow, then any empirically adequate account of practical expertise should endorse it. Since Annas is trying to give an empirically adequate account of practical expertise, it follows that she should accept the expanded Merleau-Pontyian view of action.

In the next section, I use this layered conception of awareness in action to fill out the idea that having the ability to act-in-flow constitutes having practical understanding of what one is doing.

6. Awareness, know-how and practical understanding

I have suggested that the defining feature of practical expertise is that the practical expert has an embodied, practical understanding of what they are doing. But the idea of practical epistemic states is hardly new. For instance, it has long been accepted that there is a practical kind of knowledge. As John Dewey puts it:

'We may... be said to *know how* by means of our habits... We walk and read aloud, we get off and on street cars, we dress and undress, and do a thousand useful acts without thinking of them. We know something, namely, how to do them... [If] we choose to call [this] knowledge... then other things also called knowledge, knowledge *of* and *about* things, knowledge *that* things are thus and so, knowledge that involves reflection and conscious appreciation, remains of a different sort...'.¹⁵⁵

This kind of 'know-how' is a practical kind of knowledge. Given that many theorists already accept that there is a practical kind of knowledge, it should not be too much of a stretch to think that there is a practical kind of understanding—

¹⁵⁵ Dewey (1922: 177-178), cited in Dreyfus and Dreyfus (1991: 236), cf. Ryle (1949).

an ‘understanding-how’ as it were.¹⁵⁶ In *Intelligent Virtue*, Annas rejects accounts of virtue that appeal to ‘know-how’ because she says they often ‘obliterate’ the distinction between learner and expert.¹⁵⁷ However, elsewhere, Annas phrases the difference between a routine-habit and a practical skill in terms of a difference in the kind of ‘know-how’ involved. Citing examples of her routine-habit of driving to work, and her practical skill of playing the piano, Annas says:

‘[We] can... say both that I know how to drive to my garage and that I know how to play the Schubert *Wanderer* fantasy. Given the sharp differences between these kinds of ‘knowing how,’ we should be curious as to the different notions here of ‘know-how’.¹⁵⁸

In trying to satisfy this curiosity, Annas gives an account of the kind of ‘know-how’ associated with practical skills by invoking her claim that practical skills require the drive to aspire, which includes being motivated to theoretically understand the activity. But, as I have argued in previous chapters, if what we want is an empirically adequate account of practical expertise, the epistemic state involved cannot be theoretical understanding, because having theoretical understanding entails the articulation requirement, which the empirical evidence suggests is not met by many real experts. Instead, I think the layered conception of awareness developed above allows us to give an account of practical epistemic states that can accommodate both practical knowledge and practical understanding.

It is useful at this point to recall a distinction drawn in Chapter 2, between the exercise of a skill in a routine habitual way and the exercise of *the same skill* in an expert way. For example, in Case 6 above, Ken is exercising his ability to drive in an expert way, but we can easily imagine a scenario in which Ken exercises his ability to drive in a non-expert, routine way. For instance, he might do this when dropping his children off at school. My suggestion, then, is that in exercising his ability to drive in a routine way Ken manifests his practical knowledge of driving,

¹⁵⁶ Kvanvig (2003: 189-190).

¹⁵⁷ Annas (2011a: 16, n1).

¹⁵⁸ Annas (2011b: 101).

whereas in exercising his ability to drive in an expert way, i.e., in-flow, Ken manifests his practical understanding of driving.

The key assumption that underpins this conception of practical epistemic states is that our awareness is what provides us epistemic access to the world. If, as I have argued, the way we are aware of what we are doing when acting-in-flow is different to the way we are aware of what we are doing when our actions are a matter of routine habit, then there are different ways we can be aware of what we are doing. Given the assumption that our awareness is what provides us epistemic access to the world, the layered conception of awareness proposed thus makes room for the possibility that there are different ways of having epistemic access to the world. If so, then we can say that the kind of epistemic access we have when acting-in-flow is different to the kind of epistemic access we have when our actions are the result of routine habit. The idea then, is that the agent's practical knowledge corresponds to their being habitually aware of affordances relevant to performing that activity, whereas the expert's practical understanding corresponds to their being both habitually aware and *directly* attentively aware of the affordances relevant to the activity, i.e., being attentively engaged in the activity. In this way, we have a notion of practical understanding that incorporates practical knowledge, but goes beyond it.

There are at least three differences between practical knowledge and practical understanding on this proposal. First, and most obviously, there is a difference in the phenomenology involved. I have discussed this at length above. When we act in a routine habitual way we need not pay attention to what we are doing, whereas when acting-in-flow we are totally absorbed in what we are doing. Second, having practical understanding of an activity, i.e., being able to attentively engage in the activity, allows the expert to deal with unexpected features of the task better than just having practical knowledge. This is because attentive engagement allows the agent to perceive more affordances, which allows them to spontaneously adapt in unfamiliar situations. The general idea is that by paying attention to what we are doing, we notice more details of the situation, and thus more opportunities to act. By contrast, when we are thinking about something else whilst acting habitually

we are distracted, and affordances which may have solicited us to act go unperceived.

The way that attentive engagement enables the agent to respond appropriately in unfamiliar situations can be brought out by considering the way the practical expert responds to unexpected features of a situation. To give a martial arts example, sometimes when sparring with a less skilled opponent, I am able to match their abilities without attentively engaging in the bout. But because I do not need to concentrate on what I am doing, my mind would sometimes start to drift. Of course, occasionally my opponent would land a technique, and because I was not fully concentrating this would catch me by surprise. I was not expecting the blow because I was not paying attention. By contrast, when sparring with someone of equal or better skill, I have to attentively engage in the bout just to keep up. Again, sometimes my opponent would land a technique. However, on these occasions I am not surprised and could often see the technique coming, even though I could not respond quickly enough to counter it or move out of the way. This suggests that being taken by surprise indicates that the agent is not attentively engaged in what they are doing. If this is right, it would explain why there are ‘cattle crossing’ signs and the like on roads—to alert drivers to pay attention—to make the unexpected, expected. The general thought, then, is that when attentively engaged in what one is doing it is far less likely that one will be caught by surprise, as compared to when performing in a habitual way, because one is more sensitive to opportunities to act. This is the feature of practical expertise that enables the practical expert to respond appropriately and flexibly.

The third difference between practical knowledge and practical understanding on this view is that practical understanding is more difficult to manifest than practical knowledge. This is because to attentively engage in some activity the agent has to pay attention to what they are doing in a very specific way, and this is difficult.¹⁵⁹ When we are familiar with an activity such that we can perform it without thinking about it, it is very easy to become distracted and start thinking about something

¹⁵⁹ Indeed, in Chapter 7, we will see that this difficulty forms the basis of an objection to the skill model of *wu-wei*.

unrelated. It takes effort and focus to continuously pay attention to something, especially if it is not something that is particularly exciting or interesting to us.¹⁶⁰ To maintain attentive engagement in what we are doing for more than a short period of time is thus extremely difficult if the activity is easy and familiar because the cognitive load needed to perform in-flow is much higher than the cognitive load necessary to perform habitually, and this is the reason it is easy to become distracted when performing a routine task—it saves energy. If this is right, then to manifest our practical knowledge is much easier than to manifest our practical understanding because it is much easier to exercise a practical skill in a routine way than it is to attentively engage in the activity.

At this point, a question might arise in regard to my key assumption, i.e., that awareness provides us with epistemic access to the world: What is awareness exactly? However, I do not have anything very informative to say about the nature of awareness. I think awareness is probably fundamental in the sense that it cannot be analysed in terms of something more basic. For there to be something it is like for me, I must be aware of what it is like. But I doubt that we can define what awareness is in a non-circular way. Nonetheless, on the assumption that awareness is what provides access to the world, I submit that exercising the ability to act-in-flow manifests an embodied, practical kind of understanding of what one is doing, and that this kind of practical understanding is the defining feature of practical expertise.

7. Conclusion

In this chapter, I have argued that, all else being equal, the Merleau-Pontyan framework provides a better explanation of the practical expert's ability to act-in-flow than either the standard view or the minority view of action in the contemporary literature. Further, I have extended this framework by developing a layered conception of awareness in action, which I used to fill out the idea that

¹⁶⁰ Csikszentmihalyi (1990, 2014).

there is a practical embodied way of understanding an activity. I claim that this kind of embodied understanding is the kind of practical understanding that practical expertise requires. The account of practical expertise that emerges holds:

2a) S has practical expertise of ϕ -ing only if S has come to practically understand ϕ -ing.

3) S has practical expertise of ϕ -ing only if when S ϕ 's to the best of their abilities S ϕ 's-in-flow.

On this account, cultivating practical expertise requires the learner to become more and more familiar with doing some activity. This gradually informs her perception such that she becomes habitually aware of certain features of her environment as affording her the opportunity to exercise her skill. The more experienced the learner becomes with acting in a certain way, the more she sees the world as offering the opportunity to engage in that behaviour. By repeated practise she learns to perform the activity in-flow by attentively engaging in what she is doing. In so doing, she manifests her practical understanding of the activity, and this qualifies her as an expert at doing it. When the expert is attentively engaged in an activity the affordances most salient to her commitments continuously solicit her actions, without her having to think about what to do. This allows her to spontaneously adapt to new conditions as they arise. If she finds herself in a situation in which no action is forthcoming, then she has to think about what to do by attentively dis-engaging from the activity and reflecting about the options. At this point, she has already 'dropped out of flow'.

On the assumption that our awareness is what provides epistemic access to the world, I contend that having the ability to attentively engage in an activity can be said to constitute having an embodied practical understanding of the activity. As I see it, this is what is valuable about practical expertise.

This completes my defence of my main thesis, i.e., the claim that practical expertise essentially requires an inarticulate, embodied, practical kind of understanding of

what one is doing.¹⁶¹ In the next chapter, I apply this account of practical expertise to virtue and in Chapter 6, I apply it to *wu-wei*.

¹⁶¹ Chapter 1, §1.

6. The skill model of virtue

1. Introduction

Let's recap. Julia Annas has recently defended a neo-Aristotelian account of the essential features of the process of cultivating virtue, and the state of being a fully virtuous person.¹⁶² Her account emphasises the developmental nature of virtue, and the importance of the distinction between the learner in virtue and the fully virtuous person. Annas' goal is to descriptively ground her account of virtue by drawing an analogy with the process of cultivating practical expertise, and the state of being a practical expert that results. In Chapter 2, I reconstructed Annas' account of practical expertise in the following terms:

- 1) S has practical expertise of ϕ -ing only if, in relation to ϕ -ing, S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.
- 2) S has practical expertise of ϕ -ing only if S has come to understand ϕ -ing.
- 3) S has practical expertise of ϕ -ing only if when S ϕ 's to the best of their abilities S ϕ 's-in-flow.

Claims 1) and 2) constitute intellectualism about practical expertise. Claim 3) is drawn directly from empirical studies of practical experts. Applying these claims to the virtuous person we arrive at:

- 4) S has full virtue only if S can give reasons that explain why some particular action was, is, or would be the virtuous one in the circumstances.
- 5) S has full virtue only if S has come to understand acting virtuously.
- 6) S has full virtue only if when S acts virtuously S acts-in-flow.

¹⁶² Annas (2008, 2011a, 2011b).

The notion of ‘full virtue’ here is an *ideal* analogous to the notion of full practical expertise delineated in Chapter 2. There, I interpreted an ‘ideal practical expert’ as someone that has *full* practical expertise in the sense that when exercising their skill they never make a mistake, and always perform an appropriate action. In this way, someone with full practical expertise—viz. the ideal practical expert—is like the ideally virtuous agent mentioned in 4), 5) and 6) who always does the right thing. But once again, it is important to note that, just like the notion of practical expertise, the idea of ‘being virtuous’ (as opposed to ‘being fully virtuous’) is gradable in the sense that one can be more or less virtuous. Learners in virtue can be legitimately called courageous or honest even though they do not have full virtue. Thus, for someone to qualify as virtuous (to some degree) it is not necessary that they have *full* virtue in the maximal sense that (when acting in character) they never make a mistake and always performing the right action, but only that they have some measure of courage, honesty, justice, etc.

On Annas’ account, the reason we should think that 4), 5) and 6) are true is because 1), 2) and 3) are true. For Annas, one reason the practical expert must be able to articulate reasons-explanations, i.e., claim 1), is because the ability to give reasons-explanations is generally required to teach a skill to others, and the expert should be able to do this. According to Annas, a learner needs to be given reasons because cultivating practical expertise requires coming to understand the activity, and a learner comes to understand an activity, in part, by being given reasons-explanations. The reason that cultivating practical expertise requires coming to understand the activity is because understanding allows the agent to reason their way to the appropriate action in different situations, and this is what makes for the difference between the expert and someone who can perfectly imitate them, but is not an expert. Thus, on Annas’ account, that they have come to understand the activity is what makes the practical expert special. Since for Annas 1) – 3) are empirical generalisations, her account of practical expertise is open to empirical scrutiny.

In Chapter 3, I presented empirical evidence that suggests 1) is false. In Chapter 4, I distinguished:

2a) S has *theoretical* expertise of ϕ -ing only if S has come to *theoretically* understand ϕ -ing.

From:

2b) S has *practical* expertise of ϕ -ing only if S has come to *practically* understand ϕ -ing.

On this formulation, claims 1) and 2a) together constitute intellectualism about practical expertise. I argued that since coming to theoretically understand ϕ -ing enables an agent to give reasons-explanations in relation to ϕ -ing, then in light of the empirical evidence presented in Chapter 2 that suggests practical experts often cannot give such explanations, we should reject the claim that practical expertise of ϕ -ing generally requires coming to theoretically understand ϕ -ing.¹⁶³ In which case, having theoretical understanding of ϕ -ing cannot be what makes the practical expert special. Instead, I suggested that rather than theoretical understanding, what makes the practical expert special is that they have come to practically understand ϕ -ing, i.e., that 2b) is true. Given this, if the arguments in these chapters are correct, then in the case of practical expertise intellectualism is false. By applying the distinction in kinds of understanding to the case of virtue, we end up with:

5a) S has full virtue only if S has come to *theoretically* understand acting virtuously.

And:

5b) S has full virtue only if S has come to *practically* understand acting virtuously.

In Chapter 5, I examined claim 3). I argued that since thinking about what one is doing is inimical to acting-in-flow, prominent views of action in the contemporary literature struggle to make sense of acting-in-flow considered as an exercise of one's agency. I further argued that by endorsing a Merleau-Pontyan framework

¹⁶³ That is, if 2a) was generally true this would imply that 1) is generally true, therefore (by *modus tollens*), the rejection of 1) should lead us to reject 2a).

of action we can give a parsimonious and phenomenologically grounded account of the practical expert's ability to act-in-flow. Moreover, I claimed that when the Merleau-Pontyan framework is suitably extended it can also help make sense of how exercising one's ability to act-in-flow manifests one's practical understanding of the activity. Given certain assumptions, Merleau-Ponty's views provide the resources to say that the way we are aware of what we are doing when acting-in-flow provides a specific kind of epistemic access to the world. I argued that this framework best accounts for the practical expert's ability to spontaneously adapt to new situations. On my proposal, what is special about the practical expert as compared to the non-expert who can perfectly imitate them is the practical expert's ability to perform the activity in-flow, which I described as the ability to attentively engage in what they are doing. The upshot, I claim, is that we should reject Annas' account of practical expertise and instead we should endorse an account that holds:

2b) S has practical expertise of ϕ -ing only if S has come to practically understand ϕ -ing.

3) S has practical expertise of ϕ -ing only if when S ϕ 's to the best of their abilities S ϕ 's-in-flow.

If this account is correct, then, at most, the analogy between practical skill and virtue descriptively grounds the following:

5b) S has full virtue only if S has come to practically understand acting virtuously.

6) S has full virtue only if when S acts virtuously S acts-in-flow.

However, it is important to note that it does not follow from anything I have said so far that 4) and 5a) are false. As John Doris puts it, 'to show that an ethical theory is *descriptively* inadequate is not to show that it is *normatively* inadequate'.¹⁶⁴ Thus, even though the empirical evidence suggests that the skill model does not ground 4) and 5a), there may be other reasons to endorse them. In the next

¹⁶⁴ Doris (1998: 505, original emphasis).

section, I argue that there is indeed reason to hold that in the case of virtue intellectualism is true. In §3, I address a tension that remains in Annas' account and in §4 I discuss an important aspect of Annas' account so far neglected, viz., the requirement that the fully virtuous person has a commitment to goodness.

2. Hills' arguments for intellectualism

In this section, I look at three possible reasons to accept 5a), i.e., the claim that to be virtuous, a person must have a (significant degree) of articulate theoretical understanding of acting virtuously.¹⁶⁵ In a recent article, Alison Hills defends the idea that the fully virtuous person has theoretical understanding of acting virtuously, which she calls 'moral understanding'.¹⁶⁶ Although Hills does not reject the skill model of virtue, she does think that Annas' appeal to the analogy with practical skill does not secure intellectualism about virtue. Hills acknowledges that in the case of practical skills one need not be able to give reasons-explanations to qualify as an expert.¹⁶⁷ However, she says that this not so in the case of virtue. The intellectual view Hills defends and ascribes to Annas is:

I: 'It is characteristic of a fully morally virtuous person that she grasps explicitly why her action is right and can *always* explain the reasons why it is right'.¹⁶⁸

For Hills, explicitly grasping why an action is morally right and being able to explain why it is morally right constitutes having 'moral understanding'.¹⁶⁹ Hills argues that *I* is true on the grounds that the fully virtuous person (i) reliably does the right thing, (ii) can give advice to others about the right course of action, and (iii) can justify their actions to others, and that (i), (ii) and (iii) all require moral understanding. First, Hills provides us with an example of a 'naïve moral agent':

'Nomy always does the right thing. She just recognises what the situation calls for and acts accordingly. Yet if you ask her why she did what she did she just shrugs. "It seemed

¹⁶⁵ Since 5a) entails 4), if there is reason to endorse 5a), we get 4) for free.

¹⁶⁶ Hills (2015), cf. Hills (2009).

¹⁶⁷ Hills (2015: 28).

¹⁶⁸ *Ibid.* p.9, my emphasis.

¹⁶⁹ See Hills (2009) for a more detailed description of the explanatory abilities she thinks moral understanding involves.

like the thing to do,” is all she can say. She genuinely has no idea why her actions are right. [...] She is disposed to do the right action. And she is well motivated: she cares deeply about the things that matter.¹⁷⁰

Hills tells us that the view upon which Nomy has ‘all the qualities needed for full virtue’ is:

N: ‘It is not characteristic of a fully morally virtuous person that she grasps explicitly why her action is right and can explain the reasons why it is right.’¹⁷¹

According to *N*, full virtue does not require that one has theoretical understanding of acting virtuously. As Hills notes, *N* does not hold that the virtuous person can never explain why a particular course of action is the virtuous one, just that the ability to give reasons for one’s actions is not essential to being fully virtuous. With the example of Nomy in hand, Hills argues that *N* is false because there are situations when Nomy would not be able to do (i), (ii) and (iii), as well as someone who does have theoretical understanding of acting virtuously.

2.1 Reliably doing the right thing

According to Hills, there will be occasions when Nomy will fail to reliably do the right thing because she does not have theoretical understanding of acting virtuously. We can formulate this argument as follows:

10) S has full virtue only if S reliably performs the right action.

11) S can reliably perform the right action only if S has theoretical understanding of acting virtuously.

So,

5a) S has full virtue only if S has theoretical understanding of acting virtuously.

Claim 10) seems uncontroversial (if a bit indeterminate as to what ‘reliably’ amounts to). In support of 11), Hills presents several kinds of situation in which she claims Nomy is at risk of not doing the right thing.¹⁷² The first are complex

¹⁷⁰ Hills (2015: 10-11).

¹⁷¹ *Ibid.* p.11.

¹⁷² *Ibid.* pp.23-25.

situations where the stakes are significant and there are multiple interests in play. Hills gives the following example:

‘Nomy is the executor of a will of a small business owner. Family members, employees, a co-owner and other businesses all have claims on the estate. Their claims are of different kinds: some are contractual and legally binding; others are based on what seems fair. Unfortunately there is no agreement about what would be a fair distribution.’¹⁷³

Hills’ concern is that in this situation Nomy’s lack of theoretical understanding makes her susceptible to implicit biases, and that if she acts on the basis of biased judgements she fails to reliably act well. But, Hills suggests, if Nomy theoretically understood acting virtuously, she would be better able to correct for any such biases by carefully considering the relevant reasons.

The second kind of situation in which Hills suggests Nomy would be at risk of behaving badly are when someone has offered her ‘forceful arguments’ against the course of action she is characteristically inclined to take.¹⁷⁴ Hills’ concern here is that if Nomy is presented with an argument against doing what she intuitively thinks she ought to do, then the fact that she cannot explain where the argument goes wrong would lead her to think that the course of action she is inclined to must be irrational. For instance, suppose that Nomy is a doctor who has been asked by the government for access to her patient’s records.¹⁷⁵ Suppose further that Nomy’s first impulse is to refuse to do so. The government official tells her that if they had access to the records it would substantially improve local policy-making decisions, and so would be highly beneficial for the wider community. Suppose that Nomy believes the official and cannot think of anything further to say in support of not releasing her patients’ data, but is still inclined to refuse access. Hills suggests that in this situation Nomy may well see her inclination to refuse access to the government as irrational and, recognising this, she will be more likely to acquiesce to the government’s request. After all, the government official has presented her with a forceful argument that she cannot rebut. But, if Nomy had theoretical understanding, then she would be able to think of reasons to rebut

¹⁷³ *Ibid.* p.23.

¹⁷⁴ *Ibid.* p.24.

¹⁷⁵ This is the example Hills uses (2015: 24-25).

the argument, and so would not think that her refusal to grant the government access to her patients' records is irrational.

A third kind of situation in which Hills suggests Nomy might struggle to reliably act well are situations when she has 'strong motivations for doing something else'.¹⁷⁶ The example Hills gives is one where Nomy is characteristically inclined to tell the truth, but 'strongly tempted to lie to avoid a highly embarrassing revelation becoming public'.¹⁷⁷ If Nomy had theoretical understanding, says Hills, then she would be able to rehearse to herself the reasons why telling the truth is the right thing to do, and thereby make it more likely that she will resist the temptation to lie.

How much support do these cases provide for 11)? I think they lend very little support to 11). The reason why can be seen by recalling Annas' distinction between the learner in virtue and the fully virtuous person.¹⁷⁸ What seems to be underpinning all these examples is that without theoretical understanding, Nomy will be more susceptible to some negative psychological influence (implicit bias, weakness of will, embarrassment, wishful thinking, etc.) which might lead to her doing the wrong thing. In the first situation, the concern is that without theoretical understanding Nomy might make a poor decision because she will be more susceptible to implicit bias. But, it seems plausible that if someone suffers from implicit bias, this is a good sign that they are less-than-fully virtuous. In the third kind of situation, the concern is that without theoretical understanding Nomy will not be able to remind herself of the reasons why a certain course of action is the right one when she has strong conflicting motivations. But again, it seems natural to think that having such strong conflicting motivations is an indication that one is still a learner in virtue. Having to remind oneself of the reasons why telling the truth is the right thing to do does not seem like something the ideally virtuous agent would need to do. Given this, I do not think these cases support 11).

¹⁷⁶ *Ibid.* p.25.

¹⁷⁷ *Ibid.*

¹⁷⁸ Chapter 2, §2.

The second kind of situation is less straightforward. Hills suggests Nomy is likely to go wrong if she is presented with a forceful argument against doing what she thinks she ought to do. Since she lacks theoretical understanding, Nomy would have nothing to say in response to the argument. Hills' concern is that in such situations Nomy is likely to think herself irrational and because of this she is more likely to be persuaded not to do what she thinks she ought to do. But there are other plausible options here. For instance, Nomy might think herself irrational but rather than think of irrationality as some kind of defeater for the course of action she is inclined to, she might just reduce her credence in the importance of any link between rationality and doing the right thing. Or, Nomy might think of herself as irrational but just stand firm, especially if the situation is time pressured or the stakes are very high. Given these alternative options, this case does not seem to support 11) either.¹⁷⁹ In short, the problem with the argument from reliability is that having to fight against one's own psychology to do the right action looks like the mark of the learner in virtue, rather than the ideally virtuous person.¹⁸⁰ If this is right, then at most Hills' cases suggest that 11)—the claim that to reliably act well requires theoretical understanding—applies to the learner in virtue.¹⁸¹ But they do not seem to give us any reason to think the same is true for the fully virtuous person.

A response Hills might make here is that, as a matter of fact, all human beings are susceptible to some of these negative psychological influences, and that is why we need to have theoretical understanding of morality in order to mitigate their influence. However, it seems that even if it is true that all human beings are in fact prone to implicit bias, wishful thinking, weakness of will, and so on, the implication we should draw is not that the fully virtuous person is also susceptible

¹⁷⁹ As Chris Woodard has pointed out to me in personal communication, 'having theoretical understanding might [itself] lead the virtuous agent astray in other cases (unless it is stipulated that they never make reasoning mistakes) e.g. when confronted with very plausible seeming but fallacious counterarguments. Lacking theoretical understanding might make it more likely that the agent does the right thing in that sort of case'.

¹⁸⁰ Cf. Annas (2011a: 67-68).

¹⁸¹ Cf. Dreyfus & Dreyfus (1991).

to them, but rather that becoming fully virtuous is an unattainable ideal for human beings as they actually are. As Annas says,

‘being *fully* virtuous does seem to be an ideal that we aspire towards but can never achieve [...] an ideal which none of us can exemplify’.¹⁸²

For Annas, this is part of the appeal of the skill model. Even the most virtuous actual person is still a learner, but to the extent that they have cultivated some of the characteristics of the fully virtuous person they can still qualify as honest, courageous, just, etc. Indeed, some of Hills’ own comments seem to support this idea—for example, when she says that ‘none of us is truly fully virtuous’.¹⁸³ If becoming fully virtuous is not something human beings can actually achieve, but can only aspire towards, the inference from the fact that we are all susceptible to negative psychological influences, to the claim that the fully virtuous agent must have theoretical understanding, looks invalid. It seems more plausible that becoming fully virtuous is an unachievable ideal for humans, than that the ideally virtuous person might suffer from implicit bias, weakness of will or wishful thinking. In which case, it might be best to think of virtue as a regulative ideal in the sense that though ultimately unachievable, the process of cultivating a virtuous character is still an important and worthwhile human endeavour.¹⁸⁴

In any case, there is another reason to doubt the veracity of 11). On the account of practical expertise being offered reliably acting well is accounted for by having a level of familiarity with the activity such that one can attentively engage in doing it, i.e., reliability is accounted for by the agent having a practical embodied understanding of the activity. This is because attentive engagement allows the agent to perceive more relevant affordances, which allows them to spontaneously adapt in unfamiliar situations. Since on the skill model cultivating a *fully* virtuous character requires practising virtuous actions to the point that one never makes a mistake, the fully virtuous agent will, on my proposal, be habitually aware of any opportunities to act virtuously, and when conditions are favourable, this awareness

¹⁸² Annas (2011a: 64-65, original emphasis).

¹⁸³ Hills (2015: 19).

¹⁸⁴ Although I won’t pursue this idea here, in Chapter 7, §3 I return to the idea of regulative ideals in my discussion of *wu-wei*. I think much of what I say there would also fit nicely with this conception of virtue.

will immediately solicit them to act.¹⁸⁵ By paying attention to what they are doing, the fully virtuous person is able to notice aspects of the situation that the less-than-fully virtuous person, trying to figure out what to do, might have missed. Moreover, because attentive engagement requires one's goals to be harmoniously aligned, this helps further blunt Hills' concerns about negative psychological influences. Thus, rather than 11), the claim the account I am proposing suggests is:

11a) S can reliably perform the right action only if S has *practical* understanding of acting virtuously.

Given these considerations, I think we should reject Hills' argument for intellectualism on the basis of reliability.

2.2 Giving advice

A second argument Hills gives in support of intellectualism is that without theoretical understanding of morality an agent will not be able to give good advice.¹⁸⁶ Hills acknowledges that the fully virtuous agent may not always be persuasive, but thinks they would always be able 'to offer an argument that might get some purchase'.¹⁸⁷ The nub of Hills' argument is that full virtue itself requires being a 'moral teacher', which requires being able to give good advice and explain one's actions. We can formulate this argument as follows:

12) S has full virtue only if S can teach others about virtue.

13) S can teach others about virtue only if S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

So,

4) S has full virtue only if S can give reasons that explain why some particular action was, is, or would be the appropriate one in the circumstances.

¹⁸⁵ See Chapter 5, §4.

¹⁸⁶ Hills (2015: 27-34).

¹⁸⁷ *Ibid.* p.28.

This argument is similar to the argument from education discussed in Chapter 3. As we saw there, Annas tries to ground the articulation requirement in the case of practical expertise by claiming that teaching a practical skill requires giving the learner reasons-explanations. I argued that for most of the practical skills Annas mentions the argument from education is unsound. That left a subset of ‘vocational skills’—farming, building, plumbing, fixing a computer, and surgery—for which the cultivation process more plausibly does require being given reasons-explanations.¹⁸⁸ However, the empirical evidence suggested that even in the case of vocational skills often experts could not articulate why they did what they did. Hills acknowledges that in the case of practical expertise the skill of teaching comes apart from the skill of doing.¹⁸⁹ However, she insists that this is just one area where virtue is not like a practical skill. She says:

‘Moral teaching (in the broad sense, encompassing giving advice, justifying and explaining oneself to others) is part of good moral practice: it includes actions that a virtuous person is morally required to perform. If she cannot she does not have all the virtues, she is not a moral ideal: she is not fully morally virtuous.’¹⁹⁰

If Hills is right, then if virtue is like a skill, it is like a vocational skill, but one that requires the agent can always say something in favour of the action that they are characteristically inclined to. I think the vocational skill that Annas mentions that best fits this picture is surgery. For instance, for someone to qualify as an expert brain-surgeon it seems plausible that they must have acquired both theoretical understanding and practical understanding of brain-surgery. If so, then this suggests that the ideal brain-surgeon would have highly skilled hands and would always be able to explain their preferred course of surgical action both in the lecture theatre and the operating theatre. If this is right, and virtue is both intellectual and like a practical skill, it is probably best modelled on a vocational skill like surgery.

¹⁸⁸ Chapter 3, §5.

¹⁸⁹ *Ibid.*

¹⁹⁰ *Ibid.* p.31.

2.3 Justifying our actions

Finally, Hills argues that the fully virtuous agent must have a theoretical understanding of virtue because virtue often requires us to justify our actions to others.¹⁹¹ For example, recall the case of Nomy executing a will cited above. Whatever decision Nomy makes about the distribution of assets, someone might feel disgruntled. Hills suggests that if so, justice requires Nomy gives them an explanation of her decision. As Hills puts it:

‘The exchange of justifications (offering and accepting or rejecting them) is a way of relating to one another of fundamental moral importance. It is essential to people’s interests as rational agents with moral standing. Recognising and responding to this is part of having the right moral attitudes toward other people; giving people what you owe them, giving them what they deserve, is part of the virtue of justice.’¹⁹²

According to Hills, giving reasons-explanations is part of what we morally owe to each other as moral agents. More generally, we might think that because reason-giving practices are a central part of everyday life, if the fully virtuous person is an expert at living,¹⁹³ they will be skilled at explaining and justifying themselves. This seems highly plausible to me.¹⁹⁴ Since to give a reasons-explanation is to perform an action, if an agent cannot explain themselves when our normal everyday practices require them to (including our normal everyday practice of giving each other what we morally owe), they thereby fail to act well, and hence do not qualify as having full virtue. If so, there is reason to think that a fully virtuous person would have theoretical understanding of acting virtuously. Given this, it turns out that there is a reason to endorse an intellectual account of virtue.

3. A tension in Annas’ account

So we have a reason to advocate intellectualism in the case of virtue. Since reason-giving is an integral part of everyday life, the ideally virtuous person should be able

¹⁹¹ *Ibid.* pp. 27-34.

¹⁹² *Ibid.* p. 29.

¹⁹³ See Chapter 1, §2.

¹⁹⁴ For further support that reason-giving is essential to morality and virtue, see e.g., Falk (1953) and MacIntyre (1981, Ch.3).

to explain their actions, in virtue of theoretically understanding what they are doing. However, an important tension remains in Annas' account. As we saw in Chapter 2, Annas thinks that when acting expertly practical experts reason their way to the appropriate actions. Indeed, this is one of the key reasons Annas advocates a skill model of virtue, because 'exercising a virtue involves practical reasoning of a kind that can illuminatingly be compared to the kind of reasoning we find in someone exercising a practical skill' (p.1). In the case of practical expertise, Annas tells us, 'practical mastery is at the service of conscious thought, not at odds with it' (p.14). So, on Annas' account, practical experts think about what they are doing when acting expertly, which is not the case in the exercise of a routine habit.¹⁹⁵ However, as we saw in the last chapter, consciously thinking about what one is doing is inimical to acting-in-flow. When someone is acting-in-flow they are completely absorbed in the activity and do not deliberate about or reflect on what they are doing. As reported in the last chapter, when practical experts start to think about what they are doing they 'drop out of flow'.¹⁹⁶ Thus, there is a tension between the claim that the practical expert thinks about what they are doing and the claim that they act-in-flow. This tension is palpable in Annas' writings. For instance, she says:

'[A] 'flow' experience is best taken to pick out the direct engagement with the task characteristic of expertise, *unmediated by deliberation*, something active rather than passive, ... [by analogy] the person with more developed kindness or generosity will grasp what to do in a way that is *direct and unmediated by deliberation about what to do*'.¹⁹⁷

However, when considering the possible objection that the appeal to flow 'makes acting virtuously sound too easy' and ignores occasions when doing the right thing

¹⁹⁵Hills seems to suggest something similar when she says, 'it is characteristic of a fully virtuous person to be aware of her reasons for action and to articulate them when necessary [...] many different kinds of explanation, deliberation and advice are perfectly legitimate [...] But full moral virtue requires... a full responsiveness to moral reasons that goes beyond the sensitivity required to do the right action for the right reasons, and includes the ability to respond consciously, with awareness' (2015: 33-34). However, since Hills is not committed to the claim that the virtuous person acts-in-flow (or the claim that the skill model of virtue is correct), even if she does think that the virtuous person must be consciously aware of the reasons for her actions at the time of acting, it would not present the same tension.

¹⁹⁶ Chapter 5, §2.

¹⁹⁷ Annas (2011a: 71-74, my emphasis).

seems to require ‘breaking the ongoing rhythm’ and ‘stepping back to reflect’, Annas responds:

‘This is not an objection once we understand rightly what it is for a virtuous action to be easy and enjoyable, understood in terms of the psychology of flow. It is a matter of being *unimpeded in virtuous deliberation* and action, and this is something *which can well require detachment and reflection* if it is not to lapse into routine.’¹⁹⁸

Annas’ oscillation between the claim that acting-in-flow is unreflective absorption in the activity and the claim that acting-in-flow can involve detached reflection and deliberation is a symptom of the tension in her account. Annas is not unaware of this tension, and tries to relieve it by suggesting that:

‘The conscious thoughts [involved] seem to have disappeared; they are not taking up psychological room, or we should never see learners speed up as they become experts. The thoughts have, in a useful philosophical term, effaced themselves.’¹⁹⁹

Annas goes on to suggest that the thoughts involved are dispositional in the sense that although ‘effaced’ the agent will be able to ‘recover’ them if asked.²⁰⁰ However, it is difficult to understand what it means to say that the thoughts ‘are not taking up psychological room’ and, moreover, the evidence presented in Chapter 3 suggests that in the case of practical expertise, if there are such effaced thoughts the expert often *cannot* recover them.²⁰¹ Further, if my argument in the last chapter that Merleau-Ponty’s view of action best explains the practical expert’s ability to act-in-flow is correct, then Annas does not need to appeal to unrecoverable effaced (implicit) thoughts in order to account for the agent’s actions.

The source of the tension in Annas’ account is the force of the empirical evidence that acting-in-flow does not involve thinking about what to do, and the intuitiveness of the idea that full virtue requires one to explicitly reflect on the reasons to act in the situation. To remove the tension, it seems we are faced with a choice between rejecting either (i) the claim that when acting virtuously, the virtuous person thinks about what they are doing, or (ii) the claim that when acting

¹⁹⁸ *Ibid.* p.76, my emphasis.

¹⁹⁹ *Ibid.* p.29.

²⁰⁰ Annas (2011b).

²⁰¹ Cf. Hills (2015: 22).

virtuously the virtuous person acts-in-flow. The problem with rejecting (ii) is that since acting-in-flow is an empirically robust feature of practical expertise, this would be tantamount to rejecting the skill model of virtue. So, if we want to maintain that virtue is like a practical skill, then we should reject (i). It follows that if virtue is like a skill, then when acting in character the ideally virtuous person does not consciously think about what they are doing. To the extent that sometimes doing the right thing seems to require ‘breaking the ongoing rhythm’ and ‘stepping back to reflect’ we can say that this is because ‘we’ are all less-than-fully virtuous, e.g., most of us suffer from negative psychological influences, and lack the requisite experience of acting virtuously. From *our* perspective, there are certainly situations in which *we* should pause and think through what to do. But, on the account I am suggesting, this is not the case for the *ideally* virtuous agent. Rather, I propose that the ideally virtuous agent does not suffer from implicit biases and the like, and is so familiar with acting virtuously that, when acting in character, they are always drawn to do the right thing.

This raises a question of whether this way of resolving the tension is compatible with intellectualism. It might not be on Hills’ characterisation cited above—i.e., ‘It is characteristic of a fully morally virtuous person that she grasps explicitly why her action is right and can always explain the reasons why it is right’²⁰²—if ‘grasps explicitly why’ is taken to mean ‘consciously thinks about why’ and this is indexed to the time of acting. As I formulate intellectualism about virtue, however, it consists in two claims:

- 4) S has full virtue only if S can give reasons that explain why some particular action was, is, or would be the virtuous one in the circumstances.
- 5a) S has full virtue only if S has come to theoretically understand acting virtuously.

Neither 4) nor 5a) entail that the agent needs to be consciously considering reasons to act at the time they act. At most 4) holds that the fully virtuous person will be disposed to give a reasons-explanation, if asked, but it does not follow that they

²⁰² Hills (2015: 9).

were consciously thinking about what they were doing the moment before they were asked. Likewise, because of the dispositional nature of theoretical understanding,²⁰³ 5a) also does not entail that when acting in character the fully virtuous person must be consciously deliberating. So, at least as I formulate it, intellectualism looks compatible with the claim that the ideally virtuous person acts-in-flow.

One might worry at this point that something fishy is going on. Surely, we might think, the essence of intellectualism about virtue is that the virtuous person both grasps, and acts from, the right reasons. But if we endorse the claim that when acting-in-flow the virtuous person does not think about what they are doing, how do we know that the reasons they cite are the reasons that actually explain their actions? Moreover, does it even make sense in such a case to say that there are reasons that *explain* the virtuous person's actions?

I think this worry can be alleviated. Note that the first question only seems to make sense from within the framework of the standard view or the minority view of action.²⁰⁴ On the Merleau-Pontyan framework, when acting-in-flow the agent is totally absorbed in the task at hand and is not consciously thinking about anything. Which affordances solicit the agent to act depends partly on their practical skills *and* partly on their commitments, goals and so on. So, even though they are not thinking about anything when acting-in-flow, the agent's awareness of the situation is nonetheless informed by their commitments and goals because their commitments and goals affect the salience of relevant affordances.²⁰⁵ On this view, if the ideally virtuous person performs some virtuous deed in-flow and when asked to explain why they did it they give a reason, we need not think of this as them naming the causal antecedent of the action, or the thought consideration of which motivated them to act. Thus, the worry about whether the reason the virtuous person cites actually explains (read: 'caused' or 'motivated') their action or not, does not take hold. This leads onto the second question.

²⁰³ Chapter 4, §3.

²⁰⁴ Chapter 5, §2.

²⁰⁵ Chapter 5, §3.

If the fully virtuous person is not citing a cause or motive, how can we make sense of the idea that the reason the virtuous person gives *explains* their action? I can think of two ways of answering this question. First, we could say that often the reason cited explains the agent's action in the sense that it (implicitly) refers to the goal(s) of the action (recall Ted the baker, Chapter 4, §3.4). For instance, suppose that real-Trudy regularly gives money to charity collectors. Since real-Trudy is real she is still a learner in virtue, and sometimes she has to remind herself of the fact that giving to charity will help alleviate the suffering of others. Suppose that on one such occasion, real-Trudy's seven year old niece Becky asks real-Trudy why she gives money to charity. Real-Trudy could say 'because it will help alleviate the suffering of others' and this would be a legitimate explanation of her action. But she could also say 'because it's the generous thing to do' and this seems to be just as legitimate an explanation. The way both these utterances explain is by referring to a goal of the action viz., 'alleviating the suffering of others' or 'doing the generous thing'. In this way both utterances provide a (partial) teleological explanation of real-Trudy's action. Now, suppose a bearded wizard turned real-Trudy into ideally-virtuous-Trudy. On my account, when ideally-virtuous-Trudy gives money to charity she does so in-flow: her charitable action is directly solicited by her perception of an opportunity to help others, but she does not think about what she is doing. Now, suppose that one day ideally-virtuous-Trudy is out walking with her other niece Lola. As they pass a stand for Amnesty International, ideally-virtuous-Trudy puts some money in the collection bucket. Since ideally-virtuous-Trudy is acting in character, her action is done in-flow: seeing the charity stand immediately draws forth the action of putting some money in the bucket. Now, suppose that Lola asks her aunt why she gives money to charity. Just as before, it seems perfectly legitimate for ideally-virtuous-Trudy to say either 'because it will help alleviate the suffering of others' or 'because it is the generous thing to do'. If this is right, then this example suggests that if the fully virtuous person can cite a reason that (implicitly) refers to the goal(s) of the activity, this

would qualify as a legitimate explanation of their action, whether or not they were consciously considering the reason they cite.²⁰⁶

Another way we might try to account for how a reason that was not consciously considered can be cited by the fully virtuous person in a legitimate (non-causal) explanation of their good behaviour is by focussing on the pragmatic nature of explanations and the idea that full virtue requires being a teacher and role-model for others, mentioned in the last section. The pragmatic nature of explanations refers to the fact that which explanation is appropriate in the situation will depend on features of the context, not least the background beliefs of the audience and the purpose of giving the explanation. In action explanations this will normally include trying to make sense of the action to the audience. If we combine this with the claim that being fully virtuous requires being a good moral teacher and role-model, then we can say that another part of the purpose of giving reasons-explanations, for the fully virtuous person, is to aid the less-than-fully virtuous person (i.e., the audience) to become more virtuous. Given this, we might conclude that so long as the reason the fully virtuous person gives makes sense of the action to the audience and helps them become more virtuous, then the fully virtuous person has provided a legitimate pragmatic-ethical explanation of their actions.

Given these two possible ways in which the fully virtuous person giving a reason that they were not consciously considering might provide a legitimate (non-causal) explanation of their actions, I submit that, as I have formulated it, intellectualism is compatible with the claim that when acting in character the virtuous person acts-in-flow.

4. A commitment to goodness

It seems, then, that despite the fact that when they are acting virtuously the virtuous person does not think about what they are doing, this does not entail that

²⁰⁶ Cf. Alvarez (2010: Ch.6).

AR is false. Given that we also have an independent reason to think intellectualism is true, i.e., that our everyday reason-giving practices require it, the skill model suggests that virtue requires both theoretical understanding and practical understanding of acting virtuously. However, the account is not yet complete. The final requirement Annas mentions is that the fully virtuous person has ‘a commitment to goodness’.²⁰⁷ For Annas, this is not a commitment to a transcendental ideal of goodness like Plato’s Form of the Good, or God, that we ‘cannot aspire to achieve in anything like a complete or undistorted form’ in our everyday lives (p.114). Rather, the kind of goodness Annas delineates is the goodness ‘in the living of a human life’ taken as a whole (*Ibid.*). As Annas puts it, ‘The different kinds of goodness to which the exercise of the different virtues commits the virtuous person are unified, simply, in the person’s life being lived *well*. The life is thus unified in being lived virtuously’ (p.116).

One role of the commitment to goodness, suggested by this passage, is to unify the individual virtues.²⁰⁸ For Annas, the commitment to goodness also accounts for the special admiration we have for virtuous people as compared to the way we often admire experts at non-moral practical skills like surgery (pp.100-104).²⁰⁹ For instance, if an expert surgeon habitually sexually harasses his co-workers, then it seems that though most of us would acknowledge his surgery expertise, only the vicious would admire him. Since we are all learners in virtue, we can say that the

²⁰⁷ Annas (2011a: Ch.7).

²⁰⁸ Annas argues that the virtues are developmentally unified in (2011a, Ch.6). In the next chapter, I discuss further roles such a unifying commitment might play in relation to *wu-wei*.

²⁰⁹ In light of the four-stage model of skill cultivation I set out in Chapter 3—i.e., that the learner copies, repeats, perfects, and then comes to practically understand—the picture of virtue that emerges here bears a close resemblance to Linda Zagzebski’s (2010, 2017) neo-Aristotelian ‘exemplarist moral theory’. This is particularly salient because Zagzebski’s framework has provided insight into both Kongzi’s ethics (Olberding 2008, 2012, Zagzebski 2017: 84-90) and Zhuangzi’s ethics (Kidd 2018, forthcoming.) On Zagzebski’s view, ethical ideals are grounded by ‘direct reference’ to past, present and fictional people like the moral saint, the sage, and the hero, and the way we become more saintly, wise, or heroic is by emulation of these exemplars. For Zagzebski, emulation is ‘a form of imitation in which the emulated person is perceived as a model in some respect—a model of cooking, dancing, playing basketball, doing philosophy’ (2017: 131). Like Annas’ account, Zagzebski’s view is that what explains why a learner in virtue would want to emulate an exemplar is that she admires them, and that this admiration is tracking the exemplar’s goodness. Like the four-stage model of skill acquisition, the picture of moral education Zagzebski outlines focusses on imitation and emulation of moral role models. But because exemplars are few and far between, Zagzebski emphasises the importance of learning about the lives and deeds of past exemplars in other ways, e.g. by hearing or reading stories about them, and provides us with several such stories. Although Zagzebski thinks that her theory is compatible with Annas’ account of virtue (2017: 138, n. 12), she does not explicitly discuss or endorse a skill model of virtue.

commitment to goodness is a commitment to cultivate a (fully) virtuous character. In this way, the commitment to goodness can be seen as part of the ‘drive to aspire’ needed to cultivate a virtuous character.²¹⁰ Thus, a final requirement on the skill model of virtue is:

14) S has full virtue only if S is committed to goodness.

5. Conclusion

The modified skill model of virtue I propose holds:

4) S has full virtue only if S can give reasons that explain why some particular action was, is, or would be the virtuous one in the circumstances.

5a) S has full virtue only if S has come to theoretically understand acting virtuously.

5b) S has full virtue only if S has come to practically understand acting virtuously.

6) S has full virtue only if when S acts virtuously S acts-in-flow.

And,

14) S has full virtue only if S is committed to goodness.

However, because on the layered conception of awareness to act-in-flow is to attentively engage in what one is doing, it follows that 6) is equivalent to:

15) S has full virtue only if when S acts virtuously S is attentively engaged in what they are doing.

Moreover, because to exercise one’s ability to attentively engage in an activity is to manifest one’s practical understanding of the activity, it turns out that 15) is equivalent to:

²¹⁰ Chapter 2, §2.

16) S has full virtue only if when S acts virtuously S manifests their practical understanding of what they are doing.

Finally, because 16) entails 5b), and 5a) entails 4), we can reduce the account down to:

5a) S has full virtue only if S has come to theoretically understand acting virtuously.

16) S has full virtue only if when S acts virtuously S manifests their practical understanding of what they are doing.

And:

14) S has full virtue only if S is committed to goodness.

On my proposed account of practical expertise, the skill model of virtue does not ground intellectualism about virtue. However, just like Annas and Hills, I think that cultivating ethical expertise in the form of acquiring a fully virtuous character requires coming to theoretically understand the virtues in the sense that the *ideally virtuous person* would be able to give reasons that explain why one course of action is the virtuous one in the circumstances. I also agree with Annas and Hills insofar as they take the ideal of full virtue to be unachievable, yet worth pursuing nonetheless. However, unlike Annas and Hills (it seems), I also think that if when acting in character the ideally virtuous person acts-in-flow, then when acting in character the ideally virtuous person would not think about what to do, but would intuitively respond to the needs of the situation in virtue of having an embodied practical understanding of acting virtuously. Annas' appeal to the practical expert's ability to act-in-flow is thus more insightful when it comes to understanding virtue than she appears to realise. Moreover, it brings Aristotle's conception of ethical expertise into contact with Zhuangzi's notion of ethical expertise insofar as contemporary scholars of Zhuangzi's writings often draw strong analogies between *wu-wei* and the practical expert's ability to act-in-flow. The ability to act-in-flow, thus, provides a conceptual bridge between Annas' account of virtue and

contemporary accounts of *wu-wei* via the analogy with practical skills. In the next chapter, I discuss Zhuangzi's conception of *wu-wei*.

7. The skill model of *wu-wei*

1. Introduction

In this chapter, I apply the extended Merleau-Pontyan framework of layered awareness to the case of *wu-wei*.²¹¹ According to many ancient (pre-Qin) Chinese thinkers, the best kind of life is a life lived in line with the *Dao* (the ‘Way’)—a core feature of which is attaining the ideal of *wu-wei*. Although attaining *wu-wei* is plausibly at the core of much pre-Qin thought,²¹² my focus here is on the notion of *wu-wei* in early Daoism, as suggested by the writings of Laozi and Zhuangzi.²¹³ However, although it is common to take the writings of Laozi and Zhuangzi to be broadly continuous in such a way that they can be considered to form a school of thought (‘early Daoism’), the two texts are importantly different. For example, while Laozi provides many hints concerning the psychology and behaviour of the sage—putting a particular emphasis on *wu-wei* action—there is not much in terms of how we might achieve this ideal. Zhuangzi, on the other hand, gives a far richer picture of how we might cultivate *wu-wei* by drawing our attention to examples of practical expertise. Although Zhuangzi’s ‘skill stories’ do not explicitly associate practical expertise with *wu-wei*, insofar as they exalt the way practical experts behave, it is appropriate to interpret them as depicting *wu-wei* action. This has led many theorists to draw a direct analogy between Zhuangzi’s conception of *wu-wei* and ‘acting-in-flow.’²¹⁴

Given my arguments in previous chapters, several things follow. First, if the extended Merleau-Pontyan framework best accounts for the practical expert’s

²¹¹ A version of this chapter appears in Lee (forthcoming).

²¹² Slingerland (2003).

²¹³ *The Complete Works of Zhuangzi* (Watson 2013) and *The Daodejing of Laozi* (Ivanhoe 2002). I will refer to portions of each text by the initial letter of the eponymous author’s name and the number of the chapter the text appears in, e.g., Chapter 1 of the *Daodejing* would be (L1).

²¹⁴ Csikszentmihalyi (1990: 150-151), Velleman (2008), Barrett (2011), De Prycker (2011).

ability to act-in-flow (the conclusion of Chapter 5), then if *wu-wei* essentially requires acting-in-flow, the extended Merleau-Pontyan framework is also apt to illuminate the nature of *wu-wei*. Second, if the layered conception of awareness in action is correct, then we can say that the Daoist sage manifests their practical understanding of living in line with the *Dao* by attentively engaging in all their activities. My key suggestion, then, is that the way our awareness of what we are doing manifests our actions when we are acting-in-flow is the same way the Daoist sage's awareness manifests *wu-wei* action.

The chapter proceeds as follows. In §2, I briefly outline some characteristic features of the notion of *wu-wei* in early Daoism, and argue that the extended Merleau-Pontyan framework is apt to account for *wu-wei* as a way of acting. Although the analogy between *wu-wei* and acting-in-flow is a familiar one, some have found it wanting.²¹⁵ A central concern is that the intensity of concentration generally required for acting-in-flow is far too high for human beings to sustain throughout all their actions. If this is right, then if *wu-wei* requires acting-in-flow all the time, *wu-wei* would be beyond the reach of creatures like us. In §3, I address this worry and suggest that the Merleau-Pontyan framework can help diffuse it, by providing for a way that we can practise acting-in-flow. Finally, in §4, I discuss Zhuangzi's story of Cook Ding (Z3), drawing out an important difference between *wu-wei* and acting-in-flow. I conclude that *wu-wei* action is acting-in-flow that is guided by a commitment to live in line with the *Dao*.

2. *Wu-wei* as acting without intentions

A central theme in early Daoism is scepticism of the value of theoretical knowledge i.e., knowledge achieved by speech and reasoning (L3, 10, 20, 56, 65, 71, Z2, 7).²¹⁶ The literary style of early Daoist texts, which utilise aphorisms, poetry, metaphors, contradictions and anecdotes, as well as dialogues and stories,

²¹⁵ Barrett (2011).

²¹⁶ Eno (1996).

provides evidence of this scepticism.²¹⁷ For Zhuangzi, the dichotomies inherent in ordinary language and thought are a source of existential suffering and anxiety.²¹⁸

According to Laozi, the sage has ridden themselves of excessive desires (L1, 3, 12, 19, 20, 46, 50, 67) they don't act on ulterior motives (L38), and they don't expect anything in reward for their deeds (L10). For Laozi, the sage sees no value in accumulating material possessions (L12) or power.²¹⁹ Rather, the sage practises *wu-wei* ('non-action') (L2, 3, 10, 41, 46, 48, 66, 67, 86, 93). Since living in line with the *Dao* appears to require both scepticism of the value of theoretical knowledge, and extinguishing (certain) desires, motives and personal goals, *wu-wei* has been interpreted as a life free from intentional action.²²⁰ If we take this to mean acting without conceptually representing the performance of the action in thought, then this would give us an interpretation of 'practising non-action' (*wei-wu-wei*) as 'acting without conceptually representing the performance of the action in thought' or—acting without intentions.

If *wu-wei* is, or requires, acting-in-flow, and, if the extended Merleau-Pontyan framework best accounts for the ability to act-in-flow (the conclusion of Chapter 5), then the Merleau-Pontyan framework would be apt to make sense of *wu-wei* action. Recall that on the Merleau-Pontyan framework, intentional action does not require the presence of an intention that conceptually represents the performance of the action.²²¹ According to Merleau-Ponty, an agent's actions are drawn forth by perception. We perceive the affordances of our environment and immediately respond to those affordances that solicit us most strongly. For Merleau-Ponty, perception presents the world as offering us opportunities to act such that we literally see features of our environment in terms of the actions they afford us. For instance, in ordinary conditions, when I see a glass of water, my perception presents it to me as for-drinking. On this view, the more salient an

²¹⁷ Loy (1996).

²¹⁸ Slingerland (2003: 176-182).

²¹⁹ According to Ivanhoe, material wealth and power were considered unnatural values 'foisted upon [people] by social conditioning' (2002: XVIII).

²²⁰ Loy (1985: 79).

²²¹ See Chapter 5, §3.

opportunity is to an agent, the more urgently it solicits them to act—if I am very thirsty, seeing a glass of the water will draw me more urgently than if I'm not very thirsty. The salience of any particular opportunity to act can be affected by our thoughts, feelings, commitments, goals, etc. Thus, often perception and thought *jointly* bring about action. Nonetheless, on this view thoughts and intentions are neither necessary nor sufficient for action.

For Merleau-Ponty, the practical skills we have play an essential role in the way we perceive the world, and which affordances strike us as salient.²²² Becoming skilled at some activity involves becoming more and more familiar with that activity. The more familiar we are with an activity the more we see the world in terms of affordances relevant to its performance. For instance, a guitar affords a guitarist the opportunity to play, whereas it does not afford a non-guitarist the same opportunity. And if the guitarist is in the mood to play, then her perception of the guitar will solicit her to play it.²²³ Thus, the practical skills we have inform how we perceive the world, and in turn which features of the world solicit us to act.

Since Zhuangzi often depicts *wu-wei* action as exemplified by expertise in some practical activity—whether it is catching cicadas, carving bell-stands, or swimming in turbulent waters (Z19), the extended Merleau-Pontyan framework appears apt for making sense of *wu-wei*. On my extension of this framework, the practical expert's ability to act-in-flow is accounted for by their being able to attentively engage in the activity.²²⁴ Attentive engagement describes the way the agent acting-in-flow perceives the world wherein both their habitual awareness and their attentive awareness is directed at what they are doing, but they are not explicitly thinking about what they are doing.²²⁵ Given this, if *wu-wei* is analogous to acting-in-flow, then we can say that the Daoist sage attentively engages in all their activities.

²²² Merleau-Ponty (2013: 143-144), Romdenh-Romluc (2015).

²²³ Of course, there may be external conditions that prevent her from actually doing so.

²²⁴ Chapter 5, §5.

²²⁵ Chapter 5, §4.

3. A familiar analogy

The idea that *wu-wei* is analogous to acting-in-flow is a familiar one.²²⁶ However, it does seem implausible that *wu-wei* is *identical* to acting-in-flow. Nathaniel Barrett argues that *wu-wei* can't just be acting-in-flow all the time, as the cognitive load required to sustain such behaviour would be too much, given our limited capacities.²²⁷ This is an understandable worry because the most relatable examples of acting-in-flow (e.g., racing in Formula 1 and martial arts) are those that require an intensity of concentration that we could not sustain indefinitely. However, on the account being advanced here, there seem to be at least two ways of answering this worry, which are not mutually exclusive. The first is to say that living in line with the *Dao* should be considered a regulative ideal such that any increase in attentive engagement is considered an improvement and worthwhile.²²⁸ On this way of answering the worry, we might admit that it would take super-human abilities to achieve flow '24/7'²²⁹ but nonetheless still maintain that any increase in attentive engagement is an improvement and worthwhile. No doubt *wu-wei* in this 'flow-max' sense is out of the reach of many of us. But taken as a regulative ideal, attentively engaging with one's environment *as much as possible* is available to all.

The second way to answer the worry is simply to deny it. One way to do this is to point out that *wu-wei* need not involve super-human abilities because acting-in-flow need not be an overwhelming strain on cognitive resources. The examples used to illustrate acting-in-flow often involve intense concentration: competitive sports, dancing, playing a musical instrument, high-speed driving, playing chess, surfing, puzzle-solving, and so on. These are the kinds of activities many of us are familiar with and so provide relatable examples of acting-in-flow. And I agree that

²²⁶ Csikszentmihalyi (1990: 150-151), Velleman (2008), Barrett (2011), De Prycker (2011).

²²⁷ Barrett (2011).

²²⁸ Barrett considers but rejects this because the 'problem with treating spiritual ideals as regulative is that as such they fail to confront the problem of [human] limitations' (2011: 698, cf. 694-5). However, insofar as such limitations consist in 'our limited capacity of focused attention' (*Ibid.* p.684), this is not a problem that arises on my account. This leads me on to the second way to deal with the problem of cognitive overload in the main text.

²²⁹ Barrett (2011: 698).

maintaining concentration in the way these activities suggest would be mentally exhausting and is not something one could keep up all the time.

However, we might think that such ‘high-intensity’ flow experiences are but one way of attentively engaging in what we are doing, i.e., of acting-in-flow. If so, it need not be the case that acting-in-flow always requires such sustained concentration. There are plausibly many activities that promote acting-in-flow, but that don’t require this kind of intense focus, e.g., walking, flower-arranging, making-tea, even just breathing. Further, I suggest that we can actually *practise* acting-in-flow, and, thereby, habituate the appropriate skillset needed to act-in-flow. This is the same idea that underpins the Merleau-Pontyan account of action. The more we attentively engage in what we are doing, the more familiar we become with it, and the easier it becomes to act-in-flow. In this way, to practise acting-in-flow one must practise some skill, and it seems any activity that can be raised to the level of art, whether it is walking, making-tea, cooking, painting, or gardening, would be an appropriate vehicle for such practise, so long as one can attentively engage with that activity.²³⁰ My suggestion then is that attentive engagement can be sustained for longer and longer periods: one might cook-in-flow for hours, paint-in-flow for days, write-in-flow for weeks, or tend-a-garden-in-flow for years. One doesn’t need to take up a dangerous sport to practise acting-in-flow. If this right, then the concern that achieving ‘flow-max’ would be impossible for creatures like us seems misdirected.

One might worry here that this response is too quick, and that even with such mundane activities like walking, there will be circumstances in which maintaining one’s attentive engagement will require intense concentration, and that this will eventually lead to mental exhaustion. For instance, though normally walking-in-flow doesn’t require intense concentration, walking-in-flow on ice will require intense concentration. Thus, even for mundane skills, circumstances can be such

²³⁰ Cf. Eno (1996: 142) and Dreyfus (2005: 58). There is a caveat to this, given the non-intentional nature of attentive engagement, it seems plausible that intending to act-in-flow (under that description), may hinder the manifestation of it—in the moment, as it were. However, I don’t think this precludes indirectly aiming at acting-in-flow by consciously deciding to do something which one knows *promotes* acting-in-flow, i.e., some craft, art or other skilled-activity.

that attentively engaging with those activities will require intense concentration, which will eventually lead to exhaustion.

However, I think we can alleviate this worry by distinguishing activities in a more fine-grained way. Rather than simply taking, e.g., ‘walking’ as the activity, leaving the circumstances in which one is walking as external to it, we can instead incorporate the circumstances into the activity itself. Thus, we can distinguish, e.g., the activity of ‘walking-on-sand’ from the activity of ‘walking-on-ice’ from the activity of ‘walking-quickly-on-ice,’ and so on. By distinguishing activities at this grain, we allow for the possibility that the skills we need to habituate in relation to each of them is similarly fine-grained. We can then hold that just because someone can walk-on-sand-in-flow, this by itself does not imply that they can walk-quickly-on-ice-in-flow, since these two activities will require different (if overlapping) skillsets. To be able to walk-quickly-on-ice-in-flow we need to have practised *walking quickly on ice* to the point where we can do so in-flow, and *mutatis mutandis* for walking-on-sand-in-flow (and so on, for other specific activities). In this way, we can say that if someone needs to concentrate more intensely when walking-quickly-on-ice-in-flow than when walking-on-sand-in-flow, this only demonstrates that they are more skilled (practised) at the latter than the former. More generally, the suggestion is that even if there are circumstances in which acting-in-flow takes more concentration than in other circumstances, this only shows that the agent is not yet expert at that specific activity in those circumstances. Any exhaustion that results could then be attributed to this lack of skill.²³¹

It seems then that by holding either that *wu-wei* is best understood as a regulative ideal, or that acting-in-flow need not be mentally exhausting because it can be practised, or both, we can alleviate the worry that *wu-wei* can’t require acting-in-

²³¹ One might worry that this response will lead to an explosion of practical experts. By distinguishing practical activities at this grain, it turns out there are experts at tying-shoelaces, opening-doors, walking-up-stairs, and a myriad other everyday activities. However, so long as we bear in mind how difficult attentively engaging is, I am happy to embrace this consequence and to claim that the ascription of expertise in our everyday talk is limited by pragmatic factors like the rarity of the skill, and the scope of the activity. Expert mechanics are few and far between but most of us can tie our own shoelaces. Nonetheless, on this view, technically, every time you tie your shoelaces in-flow you exercise your expert skill at tying shoelaces, and, thereby, manifest your embodied understanding of tying shoelaces.

flow all the time. Yet, as mentioned above, it does seem implausible that *wu-wei* simply is acting-in-flow all the time, even if it may be necessary. This raises the question of what the relevant difference between them is. Barrett argues that the difference is that *wu-wei* also requires a kind of ‘personal transcendence’:

‘[While] flow experience is an important—perhaps even essential—engine of Zhuangzian spirituality, it cannot supply the spiritual equanimity of *wuwei* unless it is harnessed within a conceptual framework that interprets flow experience in ways that direct personal transformation and reorient engagement with the world. In light of the special characteristics of flow—especially unselfconsciousness—it seems likely that this personal transformation involves the interpretive reconfiguration of the subject’s experience of her own personal agency.’²³²

As far as it goes, this seems correct, and Barrett’s last sentence seems to nicely capture the spirit of my project. Moreover, the distinction in awareness I argued for in Chapter 5 might point to a nice explanation of ‘transcending one’s personal perspective,’ since attentively engaging with others plausibly encourages certain affections, e.g., empathy. Such an increase in empathy might well account for being able to take on the perspective of others and thus transcending one’s own.²³³ But even if this is possible, we don’t yet have any reason to think it guaranteed, or that attentively engaging with others wouldn’t, in some circumstances, lead to heightened contempt as likely as empathy.²³⁴ So, on this account too, there is also reason to think *wu-wei* is not identical to flow. However, as argued above, this is not because the limitations of human attention mean that acting-in-flow all the time is impossible, or that there isn’t still value in acting-in-flow as much as possible. Instead, if transcending one’s perspective in this sense is what separates acting-in-flow from *wu-wei*, then achieving the best kind of life requires something that directs our affections towards empathy rather than contempt. In the next

²³² Barrett (2011: 698-699).

²³³ This is perhaps another point where Barrett and I part ways, as he takes the kind of personal transcendence involved to require ‘the attainment of a radically different perspective’ that is ‘somehow discontinuous with normal life’ (2011: 699), so may not think even deep empathy can be an appropriate mechanism for the kind of personal transcendence necessary.

²³⁴ One might think that in some circumstances it would be appropriate for the sage to feel contempt. I’m not so sure of this. Moreover, it seems that in terms of the ideal ethical agent, if the sage did appropriately feel contempt towards someone, this would ultimately be grounded in their (deeper) empathy for some others (perhaps the contemptuous agent’s victims).

section, I suggest that the extra ingredient is a commitment to live in line with the *Dao*.

4. A commitment to *Dao*

The most famous exemplar of *wu-wei* in Daoism is found in Zhuangzi's story of Cook Ding (Z3). Let's consider one final case.

Case 7: The story of Cook Ding

'Cook Ding was cutting up an ox for Lord Wenhui. At every touch of his hand, every heave his shoulder, every move of his feet, every thrust of his knee— zip! zoop! He slithered the knife along with a zing, and all was in perfect rhythm, as though he were performing the dance of the Mulberry Grove or keeping time to the Jingshou music.

"Ah, this is marvellous!" said Lord Wenhui. "Imagine skill reaching such heights!"

Cook Ding laid down his knife and replied, "What I care about is the Way [*Dao*], which goes beyond skill. When I first began cutting up oxen, all I could see was the ox itself. After three years I no longer saw the whole ox. And now— now I go at it by spirit and don't look with my eyes. Perception and understanding have come to a stop, and spirit moves where it wants.²³⁵ I go along with the natural makeup, strike the big hollows, guide the knife through the big openings, and follow things as they are. So I never touch the smallest ligament or tendon, much less a main joint.

"A good cook changes his knife once a year— because he cuts. A mediocre cook changes his knife once a month— because he hacks. I've had this knife of mine for nineteen years and I've cut up thousands of oxen with it, and yet the blade is as good as though it had just come from the grindstone. There are spaces between the joints, and the blade of the knife has really no thickness. If you insert what has no thickness into such spaces, then there's plenty of room— more than enough for the blade to play about in. That's why after nineteen years, the blade of my knife is still as good as when it first came from the grindstone.

"However, whenever I come to complicated place, I size up the difficulties, tell myself to watch out and be careful, keep my eyes on what I'm doing, work very slowly, and move the knife with the greatest subtlety, until— flop! the whole thing comes apart like clod of earth crumbling to the ground. I stand there holding the knife and look all around

²³⁵ In light of my arguments in previous chapters, Cook Ding's/Zhuangzi's claim that 'Perception and understanding have come to a stop' might look awkward. It is hard to know what to make of this claim. Edward Slingerland translates this portion of the text: 'My sensory knowledge is restrained' (2003: 200) and Brook Ziporyn has: 'My understanding consciousness, beholden to its specific purposes, comes to a halt' (2009: 22). If the layered conception of awareness is the best framework for understanding *wu-wei*, then we might render Cook Ding's/Zhuangzi's thought as something like: 'My attentive awareness is clear of conceptual representations', which seems in line with Ziporyn's translation.

me, completely satisfied and reluctant to move on, and then I wipe off the knife and put it away.”

“Excellent!” said Lord Wenhui. “I have heard the words of Cook Ding and learned how to care for life!”

As this final comment makes clear, the skilful way Cook Ding carves oxen is supposed to be an example of how to live in line with the *Dao*, i.e., Cook Ding’s performance is an example of *wu-wei*. I suggest we can interpret Cook Ding as attentively engaged in what he is doing. In Zhuangzi’s story, Cook Ding’s butchery skill is likened to dancing or conducting an orchestra ‘in perfect rhythm’ and Ding is so absorbed in the task at hand that his ‘spirit moves where it wants’.²³⁶ On the account of practical expertise I’m suggesting, we can say that because Cook Ding has practised for many years, his actions are drawn forth spontaneously and effortlessly without him needing to think about what he is doing—his attentive awareness has merged with his habitual awareness. And by attentively engaging in the activity, Cook Ding manifests his practical understanding of carving oxen.

It is worth noting the part of the story where Cook Ding tells us that when he comes to a ‘complicated place’ he ‘sizes up the difficulties’. This is noteworthy for two reasons. The first is that it seems to fit nicely with the claim that a deficit in skill leads to the agent’s actions being drawn forth by attentive awareness of a conceptual representation of relevant affordances.²³⁷ The second reason leads on from the first in that if when things get complicated Cook Ding has to represent what to do next in thought, this seems to lend credence to the claim that *wu-wei* is best seen as a regulative ideal, since on this reading even the most famous exemplar of *wu-wei* does not act-in-flow all the time.

Even more noteworthy, for the task at hand, is that as part of his explanation to Lord Wenhui, Cook Ding avows a commitment to living life in line with *Dao*—what he cares about, he says, ‘is the Way.’ There are multiple roles a commitment to *Dao* might potentially play. The first concerns the fact that a characteristic of

²³⁶ Ziporyn translates this: ‘the promptings of the spirit begin to flow’ (2009: 22).

²³⁷ Chapter 5, §4.

acting-in-flow is that the activity be goal oriented.²³⁸ Importantly, this role seems to apply whether or not we take *wu-wei* to be a regulative ideal.

If, as I urge, we take *wu-wei* to be a regulative ideal, then it seems the person on the path to sage-hood, but not yet there, will still have many personal goals and desires which structure their daily activities. Although the path to sage-hood involves extinguishing most of these over time, until that is done, these goals and desires may well come into conflict. This would be inimical to acting-in-flow because it would require thinking about what to do in those circumstances in which they do come into conflict. In this way, the commitment to live in line with the *Dao* might play a *negentropic* role in that it could help ‘order’ the less-than-sage-like agent’s personal goals and desires, bringing them into harmony. As an analogy, we might think of the commitment as the captain that brings the team into line allowing them to work efficiently as one. Given this, we can see the commitment to live in line with the *Dao* as playing a kind of managerial role with respect to the regulative nature of the ideal.²³⁹ If, on the other hand, *wu-wei* is not taken as a regulative ideal in this way, then on my account achieving *wu-wei* would require acting-in-flow all the time. However, given the goal-oriented nature of acting-in-flow, and the fact that the sage has rid themselves of most desires and personal goals, it seems there would need to be some overarching ‘life-goal’ to structure their daily activities such that they can continuously act-in-flow. The commitment to live in line with the *Dao* would be apt to play this role.

Moreover, in either case, the commitment would seem to play an important role in terms of the actions performed, since on the Merleau-Pontyan inspired view having such a commitment would make certain affordances more salient than others, and thereby affect which affordances actually solicited the agent to act. If this is right, then for both the sage, and those on the path to sage-hood, attaining

²³⁸ Chapter 2, §4.

²³⁹ This is analogous to the role the commitment to goodness played of unifying the disparate virtues discussed in the last chapter. I take it that the rest of what I say about the role of the commitment to *Dao* also applies, *mutatis mutandis*, to the role of the commitment to goodness in the case of virtue.

wu-wei would seem to require that one has a fundamental commitment to live in line with the *Dao*.

Finally, the commitment to live in line with the *Dao* might play the role of directing us towards empathy rather than contempt, which would then provide for the ‘personal transcendence’ that differentiates *wu-wei* from acting-in-flow discussed in the last section. There are two ways this might happen. We might hold that the commitment to live in line with the *Dao* is itself a substantively moral commitment such that to have this commitment entails being committed to empathy.²⁴⁰ On the other hand, we might hold that the commitment to live in line with the *Dao* does not itself entail a commitment to empathy, but rather that people are by nature empathetic, and that being committed to the *Dao* allows this nature free reign by ‘masking’ certain tendencies that might serve to block its manifestation. This latter line of thought is congruent with the idea that the sage has extinguished excessive desires and personal goals, since such desires and personal goals might plausibly block any natural empathy we may have. Either way, my suggestion is that the commitment to live in line with the *Dao* is what directs us towards empathy rather than contempt, and so is what allows us to truly transcend our own personal perspective. Having such a commitment is, I think, what makes for the difference between acting-in-flow and *wu-wei*.

5. Conclusion

In this final chapter, I have argued that given Zhuangzi’s apparent exaltation of practical expertise, the extended Merleau-Pontyan framework is apt to make sense of the ideal of *wu-wei*. If *wu-wei* is like practical expertise, then on the account of practical expertise I have proposed, *wu-wei* requires one to attentively engage in what one is doing as much as possible, and to have a commitment to live life in line with the *Dao*. Moreover, if I am right, then by attentively engaging in all their activities the Daoist sage manifests their embodied practical understanding of

²⁴⁰ However, this line of thought is controversial, cf. Eno (1996: 142, 150, n. 49).

living in line with this commitment. I have addressed a key objection to this view, i.e., that acting-in-flow all the time would be beyond the capacities of creatures like us, and argued that by appealing to the Merleau-Pontyan framework, we can defuse this objection by claiming that the ability to act-in-flow is an ability that we can cultivate by practising some practical skill.

In the context of my overall project, we can say that there is a tighter analogy between *wu-wei* and practical expertise, than there is between virtue and practical expertise. This is because whilst cultivating practical expertise, virtue and *wu-wei* all require the learner to come to practically understand what they are doing, only cultivating virtue requires that the learner also comes to theoretically understand what they are doing.

8. Conclusion

My project in this dissertation has been to investigate the skill model of ethical expertise. The skill model of ethical expertise holds that becoming an ethical expert is (like) becoming an expert in a practical skill. The question I have been most interested in answering is: What is the nature of practical expertise? By appealing to the phenomenology of cultivating practical skills, and Merleau-Ponty's views on action and perception, I have argued that practical expertise essentially requires an inarticulate, embodied, practical kind of understanding of what one is doing. My key contribution has been extending Merleau-Ponty's account by developing a layered conception of awareness in action. On the assumption that awareness is what provides us epistemic access to the world, the layered conception of awareness in action allows us to say that exercising the ability to attentively engage in an activity itself manifests our embodied practical understanding of what we are doing. In which case, there is value in cultivating the ability to attentively engage in what you are doing. I then applied this account of practical expertise to both virtue and *wu-wei*. As I interpret Zhuangzi, *wu-wei* essentially involves being able to act-in-flow. On Annas' account, being virtuous also involves acting-in-flow. I have proposed that to act-in-flow requires one to attentively engage in what one is doing. If this is right, then the ability to attentively engage in what one's is doing is the common thread that links practical expertise, *wu-wei* and virtue, and contributes to the value of all three.

To end my investigation, I draw out some implications of my argument for our educational practices today. On my account, the primary way we cultivate the ability to attentively engage in what we are doing is by practising some physical activity that can be elevated to an elite level. Prominent activities apt to promote attentive engagement include playing sports, athletics, swimming, rock-climbing, dancing, martial arts, playing a musical instrument, painting, sculpting, pottery, and cooking. If we want our children to cultivate the ability to attentively engage in

what they are doing, then I think we should encourage these activities.²⁴¹ If we want our children to cultivate ethical expertise, then, for *wu-wei* we also need to encourage a commitment to live in line with the *Dao*, help them to rid themselves of excessive desires, and teach them not to value the accumulation of material wealth and political power. On the other hand, if we want our children to cultivate virtue, then we need to educate them in the principles of virtue, and encourage a commitment to the goodness of a life spent cultivating the various moral and intellectual virtues. In all three cases, we need to provide role-models for them to imitate. And since parents and teachers are the default role-models, if we want our children to cultivate ethical expertise, then parents and teachers will need to do the same.

²⁴¹ This advice contrasts starkly with the UK government's current educational focus on promoting science, technology, engineering and mathematics (STEM), rather than sport, music and art.

References

Allen, B. (2015). *Striking Beauty: A Philosophical Look at the Asian Martial Arts*. Columbia University Press.

Allen, K. (2019). "Merleau-Ponty and Naïve Realism." *Philosophers' Imprint* 19.

Alvarez, M. (2010). *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford University Press.

Angier, T. (2010). *Techne in Aristotle's Ethics: Crafting the Moral Life*. Continuum.

Annas, Julia (1993). *The Morality of Happiness*. Oxford University Press.

- (1995). "Virtue as a skill." *International Journal of Philosophical Studies* 3 (2):227-243.
- (2003). "The structure of virtue." In DePaul, M. & Zagzebski, L. (eds.) (2003), pp. 15-33.
- (2008). "The phenomenology of virtue." *Phenomenology and the Cognitive Sciences* 7 (1):21-34.
- (2011a). *Intelligent Virtue*. Oxford University Press.
- (2011b). "Practical Expertise." In Bengson, J. & Moffett, M. A. (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*. Oxford University Press, pp. 101-112.

Barrett, N. F. (2011). "Wuwei and flow: Comparative reflections on spirituality, transcendence, and skill in the Zhuangzi." *Philosophy East and West* 61 (4):679-706.

Barrington, K. (1968). *Playing It Straight*. London: Stanley Paul.

Benner, P. (1984). *From Novice to Expert: Excellence and Power in Clinical Nursing Practice*. Berkshire: Addison Wesley.

Berendzen, J. C. (2010). "Coping Without Foundations: On Dreyfus's Use of Merleau-Ponty." *International Journal of Philosophical Studies*, 18:5, 629-649

Bradford, G. (2015). *Achievement*. Oxford University Press.

Broadie, S. & Rowe, C. (eds.) (2002). *Aristotle: Nicomachean Ethics: Translation, Introduction, Commentary*. Oxford University Press.

- Chisholm, R. M. (1976). *Person and Object: A Metaphysical Study*. Open Court.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- Clark, A. & Chalmers, D. J. (1998). "The extended mind." *Analysis* 58 (1):7-19.
- Clarke, R. (2003). *Libertarian Accounts of Free Will*. Oxford University Press.
- Csikszentmihalyi, M. (1990) *Flow: The Psychology of Optimal Experience*. HarperCollins.
- (2014) *Flow and the Foundations of Positive Psychology: The Collected Works of Mihaly Csikszentmihalyi*. Springer.
- Davidson, D. (1963). "Actions, Reasons, and Causes." *Journal of Philosophy* 60 (23):685.
- (1973). "Freedom to act." In Honderich, T. (ed.), *Essays on Freedom of Action*. Routledge.
- De Prycker, V. (2011). "Unself-conscious control: Broadening the notion of control through experiences of flow and Wu-Wei." *Zygon* 46 (1): 5-25.
- DePaul, M. & Zagzebski, L. (eds.) (2003). *Intellectual Virtue: Perspectives From Ethics and Epistemology*. Oxford University Press.
- Dewey, J. (1922). *Human Nature and Conduct, An Introduction to Social Psychology*. Henry Holt.
- Doris, J. M. (1998). "Persons, situations, and virtue ethics." *Noûs* 32 (4):504-530.
- Dreyfus, H. L. (2000). "A Merleau-Pontyan Critique of Husserl's and Searle's Representationalist Accounts of Action." *Proceedings of the Aristotelian Society*, Vol. 100: 287-302.
- (2002). "Intelligence without representation – Merleau-Ponty's critique of mental representation the relevance of phenomenology to scientific explanation." *Phenomenology and the Cognitive Sciences* 1 (4):367-383.
 - (2005). "Overcoming the Myth of the Mental: How Philosophers Can Profit from the Phenomenology of Everyday Expertise." *Proceedings and Addresses of the American Philosophical Association* 79 (2):47-65.
- Dreyfus, H. L. & Dreyfus, S. E. (1986). *Mind over Matter*. New York Free Press.
- (1991) "Towards a Phenomenology of Ethical Expertise." *Human Studies* 14 (4), pp. 229-250.

- Driver, Julia (2013). "Moral expertise: Judgment, practice, and analysis." *Social Philosophy and Policy* 30 (1-2):280-296.
- Elgin, C. (2007). "Understanding and the facts." *Philosophical Studies* 132 (1):33 - 42.
- Eno, R. (1996) "Cook Ding's Dao and the Limits of Philosophy." In Kjellberg, P. & Ivanhoe, P. J. (eds.) (1996), pp. 127-151.
- Ericsson, K. A., Krampe, R. T. & Tesch-Römer, C. (1993). "The role of deliberate practice in the acquisition of expert performance." *Psychological Review* 100 (3):363-406.
- Falk, W. D. (1953). "Goading and Guiding." *Mind*, 62(246), 145-171.
- Gibson, J. J. (1977) "The Theory of Affordances." In Shaw, R. & Bransford, J. (eds.), *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*. Hillsdale, NJ: Erlbaum, pp. 67-82.
- Grimm, S. R. (2010). "The goal of explanation." *Studies in History and Philosophy of Science Part A* 41 (4):337-344.
- (2011). "Understanding." In Pritchard D. & Berneker S. (eds.), *The Routledge Companion to Epistemology*. Routledge.
 - (2012). "The Value of Understanding." *Philosophy Compass* 7 (2):103-117.
 - (ed.) (2017). *Making Sense of the World: New Essays on the Philosophy of Understanding*. Oxford University Press.
 - (2019). "Understanding as an Intellectual Virtue." In Battaly, H. (ed.), *Routledge Handbook to Virtue Epistemology*. Routledge, pp. 340-351.
- Hills, A. (2009). "Moral testimony and moral epistemology." *Ethics* 120 (1):94-127.
- (2015). "The Intellectuals and the Virtues." *Ethics* 126 (1):7-36.
 - (2016). "Understanding Why." *Noûs* 50 (4):661-688.
- Hursthouse, R. (1999). *On Virtue Ethics*. Oxford University Press.
- Ivanhoe, P. J. (2002). *The Daodejing of Laozi*. Hackett.
- Johnson, M. (2015). "Embodied understanding." *Frontiers in Psychology* 6.
- Kahneman, D. & Klein, G. (2009). "Conditions for Intuitive Expertise A Failure to Disagree." *The American psychologist*. (64): 515-26.
- Kelly, S. D. (2002). "Merleau-Ponty on the body." *Ratio* 15 (4):376-391.

- Kidd, I. J. (2018). "Spiritual exemplars." *International Journal of Philosophy and Theology* 79 (4):410-424.
- (forthcoming). "Following the Way of Heaven': Exemplarism, Emulation, and Daoism." *Journal of the American Philosophical Association*.
- Kitcher, P. (1981). "Explanatory unification." *Philosophy of Science* 48 (4):507-531.
- Kjellberg, P. & Ivanhoe, P. J. (eds.) (1996). *Essays on Skepticism, Relativism, and Ethics in the Zhuangzi*. Suny Press.
- Klein, G. (1998) *Sources of Power: How People Make Decisions*. MIT Press.
- (2009) *Streetlights and Shadows: Searching for the Keys to Adaptive Decision Making*. MIT Press.
- Krein, K. & Ilundáin, J. (2014) "Mushin and Flow: An East–West comparative analysis." In Priest & Young (eds.) (2014), pp. 139-164.
- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lakoff, G. & Johnson, M. (1980). *Metaphors We Live by*. University of Chicago Press.
- Lee, M. (forthcoming). "*Wu-wei*, Merleau-Ponty, and being aware of what we do." *Philosophy East and West*, 70:1, doi:10.1353/pew.0.0174.
- Leibowitz, U. D. (2011). "Scientific Explanation and Moral Explanation." *Noûs* 45 (3):472-503.
- Leibowitz, U. D. & Sinclair, N. (eds.) (2016). *Explanation in Ethics and Mathematics: Debunking and Dispensability*. Oxford University Press.
- Lipton, P. (2009). "Understanding without explanation." In de Regt, H. W., Leonelli, S. & Eigner, K. (eds.), *Scientific Understanding: Philosophical Perspectives*. University of Pittsburgh Press. pp. 43-63.
- Loy, D. (1985). "Wei-wu-Wei: Nondual action." *Philosophy East and West*, 35 (1):73-86.
- Lynch, M. (2017). "Understanding and Coming to Understand." In Grimm, S. (ed.) (2017), pp. 194-208.
- Machek, D. (2011). "The Doubleness of Craft: Motifs of Technical Action in Life Praxis according to Aristotle and Zhuangzi." *Dao: A Journal of Comparative Philosophy* 10 (4): 507-526.

- MacIntyre, A. (1981). *After Virtue*. Bloomsbury.
- McClelland, T. (forthcoming). "The mental affordance hypothesis." *Mind*, fzz036, <https://doi-org.ezproxy.nottingham.ac.uk/10.1093/mind/fzz036>
- (2015). "Affording introspection: an alternative model of inner awareness." *Philosophical Studies* 172 (9): 2469-2492.
- Mele, A. R. (1983). "Akrasia, reasons, and causes." *Philosophical Studies* 44 (3): 345-368.
- (1992). "Recent Work on Intentional Action." *American Philosophical Quarterly*. 29 (3): 199 - 217.
- Merleau-Ponty, M. (2013). *Phenomenology of Perception*. Translated from the French by D. Landes. Taylor & Francis.
- Milner, D. A. & Goodale, M.A. (1995). *The Visual Brain in Action*. Oxford University Press.
- Miyamoto Mushashi (2002). *The Book of Five Rings*. Translated from the Japanese by Wilson, S. Shambhala.
- Moeller, H-G. (2004). *Daoism Explained: From the Dream of the Butterfly to the Fishnet Allegory*. Open Court.
- Montero, B. G. (2015). "Thinking in the Zone: The Expert Mind in Action." *Southern Journal of Philosophy* 53 (S1): 126-140.
- Nagel, T. (1986). *The View From Nowhere*. Oxford University Press.
- Nakamura, J. & Csikszentmihalyi, M. (2009) "The concept of flow." In Sunder, C.R. & Lopez, S.J. (eds.) *Oxford Handbook of positive psychology*. Oxford University Press, pp. 89-105.
- Narvaez, D. & Lapsley, D. K. (2005). "The Psychological Foundations of Everyday Morality and Moral Expertise." In Lapsley, D. K. & Power, F. C. (eds.), *Character psychology and character education*. University of Notre Dame Press, pp. 140-165.
- O'Connor, T. (2000). *Persons and Causes: The Metaphysics of Free Will*. Oxford University Press.
- Olberding, A. (2008). "Dreaming of the Duke of Zhou: Exemplarism and the analects." *Journal of Chinese Philosophy* 35 (4): 625-639.
- (2011). *Moral Exemplars in the Analects: The Good Person is That*. Routledge.

Priest, G. & Young, D. (eds.) (2014) *Philosophy and the Martial Arts: Engagement*. Routledge.

Pritchard, D. (2009). "Knowledge, Understanding and Epistemic Value." *Royal Institute of Philosophy Supplement* 64: 19-43.

Riggs, W. D. (2003). "Understanding 'Virtue' and the Virtue of Understanding." In DePaul, M. & Zagzebski, L. (eds.) (2003), pp. 203-227.

- (2009). "Understanding, Knowledge, and the Meno Requirement." In Haddock, A., Millar, A. & Pritchard, D. (eds.), *Epistemic Value*. Oxford University Press.

Romdenh-Romluc, K. (2011a). *Routledge GuideBook to Merleau-Ponty and the Phenomenology of Perception*. Routledge.

- (2011b). "Agency and embodied cognition." Proceedings of the Aristotelian Society, Vol. CXI, Part I, pp. 79-95.
- (2012). "Thought in action." In Zahavi, D. (ed.) *Oxford Handbook of Phenomenology*. Oxford University Press, pp. 198-215.
- (2014). "Habit and attention." In Moran, D. & Jensen, R. (eds.) *The Phenomenology of Embodied Subjectivity*. Springer, pp. 3-19.
- (2015). "Merleau-Ponty: actions, habits, and skilled expertise." In Dahlstrom, D., Elpidorou, A. & Hopp, W. (eds.) *Philosophy of Mind and Phenomenology: Conceptual and Empirical Approaches*. Routledge, pp. 98-116.

Rubén, D-H. (ed.) (1993). *Explanation*. Oxford University Press.

Ryle, Gilbert (1949). *The Concept of Mind*. Hutchinson & Co.

Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.

Slingerland, E. (2003). *Effortless Action: Wu-Wei as Conceptual Metaphor and Spiritual Ideal in Early China*. Oxford University Press.

Sosa, E. (1991). *Knowledge in Perspective*. Cambridge University Press.

Sterelny, K. (2010). "Minds: extended or scaffolded?" *Phenomenology and the Cognitive Sciences* 9 (4): 465-481.

Stichter, M. (2007) "Ethical Expertise: The Skill Model of Virtue." *Ethical Theory & Moral Practice*, 10: 183-194

Strevens, M. (2013). "No understanding without explanation." *Studies in History and Philosophy of Science Part A* 44 (3):510-515.

- Swartwood, J. D. (2013). "Wisdom as an Expert Skill." *Ethical Theory and Moral Practice* 16 (3): 511-528.
- Taylor, R. (1963). *Metaphysics*. Prentice-Hall.
- Varela, F., Thompson, E. & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.
- Velleman, D. J. (1992). "What Happens When Someone Acts?" *Mind* 101 (403): 461-481.
- (2008). "The Way of the Wanton." In Adkins, K. & MacKenzie, C. (eds.) *Practical Identity and Narrative Agency*. Routledge, pp. 162-92.
- Watson, B. (ed.) (2013). *The Complete Works of Zhuangzi*. Columbia University Press.
- Webber, J. (2002). "Doing Without Representation: Coping with Dreyfus." *Philosophical Explorations*, 5:1, 82-88.
- Wiland, E. (2012). *Reasons*. Continuum.
- Wilkenfeld, D. A. (2013). "Understanding as representation manipulability." *Synthese* 190: 997–1016.
- Wilkenfeld, D. A., Plunkett, D. & Lombrozo, T. (2016). "Depth and deference: When and why we attribute understanding." *Philosophical Studies* 173: 373–393
- Yearley, L. (1996) "Zhuangzi's Understanding of Skillfulness and the Ultimate Spiritual State." In Kjellberg, P. & Ivanhoe, P. J. (eds.) (1996), pp. 152-182.
- Zagzebski, L. T. (1996). *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.
- (2001). "Recovering Understanding." In Steup M. (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.
 - (2010). "Exemplarist virtue theory." *Metaphilosophy* 41 (1-2): 41-57.
 - (2017). *Exemplarist Moral Theory*. Oxford University Press.
- Ziporyn, B. (2009). *Zhuangzi: The essential Writings with Selections from Traditional Commentaries*. Hackett.