KNOWING-HOW: A SUBSET OF KNOWLEDGE AND OF ABILITY

Thesis submitted by Robert Segall for the degree of Master of Arts by dissertation in the Department of Philosophy, Faculty of Humanities, University of Cape Town

August 2003

This work has not been previously submitted in whole, or in part, for any award of any degree. It is my own work. Each significant contribution to, and quotation in, this dissertation from the work, or works, of other people has been attributed, and has been cited and referenced.

Signature Date 26/11/03

Signed by candidate

Acknowledgements

I would like to thank my supervisor, Dr Jeremy Wanderer. His thoughtful, penetrating and constructive criticisms of earlier drafts have been both an essential part of writing this thesis and an enjoyable education for me.

I would especially like to thank my wife, Nan Yeld, for her unfailing warmth and kindness.

KNOWING-HOW: A SUBSET OF KNOWLEDGE AND OF ABILITY

ABSTRACT

The concept of knowing-how as a form of knowledge distinct from propositional knowledge (or knowing-that) has been invoked widely in the half century since Ryle introduced it. It has figured in the philosophy of science, ethics, epistemology and philosophy of mind.

I begin by looking at Ryle's concept as it appears in the twenty first century, without his philosophical behaviorism and especially in the light of the influence of the computational theory of mind. This last has profoundly influenced our worldview, although few believe that it adequately models the mind/brain.

I characterize knowing-how in detail and propose a definition - something that surprisingly has been neglected in the epistemological literature. For A to know-how to φ (ride a bicycle, play chess, speak Basque)

1. A can entertain φing for a purpose.

2. A is acquainted with a set of practical procedures for successful φing.

3. A exhibits or has exhibited reasonable success at φing.

4. A complete propositional account of φing is not available to A.

5. A must have learned to φ via a process of instruction which may include self-instruction (and may be an unconscious process).

From the consequences of this definition, in particular its implications for knowledge of other minds, I conclude that knowing-how is a distinct subset of ability. I then argue that this subset of ability is also a subset of knowledge.

Given the denial by some philosophers of the very existence of knowing-how, I seek some paradigmatic exemplars.

I examine grammatical knowledge on the basis that in broad terms Chomsky's account of language is a scientific theory on a par with most of contemporary science. On this basis, but *contra* Chomsky, I argue that our grammatical competence with our own idiolect is knowing-how. Thus knowing-how is indeed a subset of knowledge as well as a subset of ability.

I then consider evidence from the immense recent progress in neuroscience. The neurobiological studies of vision provide compelling evidence for a form of knowledge that is not propositional knowledge. Essential to our visuomotor action is knowledge that is necessarily embodied knowledge, and again it is argued that in language use we have a clear exemplar of knowing-how.

Next, I turn to the area of consciousness and propose that the best account of our apprehension of some qualia - those whose apprehension is enabled or enhanced by learning - constitutes another persuasive example of knowing-how.

As a last example, I examine whether virtue can be seen as knowing-how.

Finally, I counter the objections to Rylean knowing-how raised in a substantial recent paper by Stanley and Williamson.

Knowing-how, a form of ability that is simultaneously a form of knowledge, is a categorization that enables a better view of the concept of knowledge. That there should be a various forms of knowledge – conscious, unconscious, *a priori*, necessary, empirical and practical is unsurprising given the complexity of the human brain. Characterizing each is an essential element in arriving at a comprehensive theory of knowledge.

CHAPTER 1. RYLE'S REGRESS ARGUMENT AND ITS CONTEMPORARY REFUTATION

> *There is (I continue to claim) a characteristic cluster of properties that are typical examples of local mental processes reliably shared with one another but not with typical instances of global ones. Three of these features are most pertinent to our purposes: local mental processes appear to accommodate pretty well to Turing's theory that thinking is computation; they appear to be largely modular; and much of their architecture, and of what they know about the proprietary domains of application, appears to be innately specified*
>
> *By contrast, what we found out about global cognition is mainly that it is different from the local kind in all three of these respects; and that, because it is, we deeply do not understand it.*
>
> Jerry Fodor[1]

1.1 Ryle's Analysis

Ryle, in his book "The Concept of Mind"[2], introduced the concept of knowing-how as something distinct from propositional knowledge, or knowing-that. He was intent on debunking a conception of mental activity that he referred to as the "intellectualist legend", which he took to be the generally accepted view of mind in the first half of the twentieth century. He saw it as a vestigial residue of Cartesian dualism, which had seduced and confused thinkers, most of whom were not in fact themselves dualists. Essentially, he regarded the view he was criticizing as a category mistake – the mind was wrongly seen as something that did things that enabled human performances like playing

---

1 J.A. Fodor, "the mind doesn't work that way", Cambridge: MIT Press, 2000, p 5

[2] G. Ryle, "The Concept of Mind", London: Hutchinson, 1949

chess and playing soccer intelligently. In fact, the mind is manifest in the intelligent playing of chess or soccer. As he says,

> "The crucial objection to the intellectualist legend is this. The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle." (p 30)

It is difficult to expound this view further, better than Ryle does himself,

> "To put it quite generally, the absurd assumption made by the intellectualist legend is this, that a performance of any sort inherits all its title to intelligence from some anterior internal operation of planning what to do. Now very often we do go through such a process of planning what to do, and, if we are silly, our planning is silly, if shrewd, our planning is shrewd. It is also notoriously possible for us to plan shrewdly and perform stupidly, i.e. to flout our precepts in our practice. By the original argument, therefore, our intellectual planning process must inherit its title to shrewdness from yet another interior process of planning to plan, and this process could in its turn be silly or shrewd. The regress is infinite. And this reduces to absurdity the theory that for an operation to be intelligent it must be steered by a prior intellectual operation. What distinguishes sensible from silly operations is not their parentage but their procedure, and this holds no less for intellectual than for practical performances. "Intelligent" cannot be defined in terms of "intellectual" or "knowing how" in terms of "knowing that";" (p 31)

Ryle gives numerous examples, but his essential point is straightforward, that there must be a knowing-how prior to knowing that. Our mental activities are manifested in being a good marksman or being a good centre forward or cyclist, and are not confined to making inferences from propositions, although that too is the mind in action.

> "Overt intelligent performances are not clues to the workings of minds; they are those workings." (P 58)

He would add, of course, that stupid performances are likewise the workings of minds.

In his positive analysis, as distinct from his criticism of the "intellectualist legend", Ryle has had numerous critics. The chief purpose of this chapter is to begin to argue that Ryle's knowing-how is a vital contribution to our analysis and characterization of knowledge, despite the fact that I defend the restoration (in its modern form) of the "intellectualist legend" and do not defend Ryle's philosophical behaviourism, which is now often rejected[3]. I consequently separate the arguments against Ryle's general stance from his specific advocacy of knowing-how as something distinct from knowing-that. This last, which is the present concern, does not fall with the failure of his philosophical behaviourism.

1.2 The Resolution of the Regress Problem

1.2.1 Introduction

In discussing Ryle's argument and its attack on the so-called intellectualist account of the workings of the human mind, we face a problem. The current counter to Ryle's argument, or to put it another way, the modern form of the intellectualist legend, rests on a theory of mind that draws, at least to some extent, on a loose, implicit form of a computer model of some, but only some, aspects of the mind. This is so, in the vague general sense that, in accounting for some mental phenomena, it draws on the notion of a primitive form of representation that is, very approximately, analogous to that in the computer.

---

[3] As Goldman puts it, "Clearly, the states that mental predicates denote interact in systematic ways with states denoted by other mental predicates, and this interaction is important to their meaning." (A.I. Goldman, "Philosophical Applications of Cognitive Science", Boulder: Westview Press, 1993, p 72) Other objections, less telling against Ryle's sophisticated form of behaviourism, include the neglect of the "inner" aspect of mental properties, and Putnam's objection, involving the super Spartans who manifest no pain behaviour.

Such an argument is presented in this chapter.

However, as the second paragraph of the above quotation from Fodor implies, we have no clear idea how to begin an argument on any other basis, and, thus, the whole of this chapter is open to challenge. This is because the argument I am mounting is proposed as being quite general. That is, it is not confined to local mental processes such as those, for example, within the ambit of the putative (and plausible) language module model. On the other hand, I would claim that within any ultimate theory of mind there simply must be some elementary representations of the kind envisaged. Proceeding as Ryle's critics do is at least in accord with the norms of conventional science, namely making use of the best available picture.

1.2.2 Fodor's Account

Fodor[4] says we do things without knowing how we do them, like judging depth. We can show experimentally that we do this last by *inter alia* measuring texture gradients. This is achieved by the central nervous system. Thus we have mental competencies about which we know nothing. At this point, Fodor considers Ryle's objection I have quoted above, "The crucial objection to the intellectualist.... ", and says that this is perfectly correct. He suggests, however, that one must make the essential distinction between what he calls mental adverbials, which relate to traits (like being clever) and mental competencies like speaking Latin. Note that speaking Latin well would

---

[4] J.A. Fodor, "The Appeal to Tacit Knowledge in Psychological Explanation", Journal of Philosophy, 65, 627 – 640 (1968)

be a mental adverbial since it refers to the manner of speaking Latin. Thus Ryle would be right were it not for the fact that mental competencies can be progressively simpler and end in what Fodor refers to as an elementary operation. Ryle assumes that each prior operation contains the ingredient of the same trait, namely being intelligent, but being intelligent is not doing something; it is doing something in a certain way. Ryle fails to see that this trait (of intelligence) can (but may not) arise out of an ultimate simple competence.

Fodor puts it graphically by envisaging a series of homunculi with instruction manuals. The manual the little man consults for speaking Latin had better be the same as that for speaking Latin well, he says, or we would indeed be heading for a Rylean regress, since there would be a manual on speaking (Latin well) well. In other words speaking Latin is a competence. Being intelligent is a trait; you can't be intelligent well (or badly). There is nothing wrong with seeing the brain as a huge team of homunculi, provided that it is recognized that the stupidest of these must be extremely stupid. Thus at bottom there is, as Ryle asserts, knowing-how, but the intellectualist legend survives in the modern form. That is, a model (at least for the modular parts of the brains functioning) along computational lines is defensible.

1.2.3 Dennett's Account

Dennett[5], like Fodor, acknowledges that Ryle is correct when he asserts that the regress of representers has to stop somewhere, with systems having

---

[5] D.C. Dennett, "Styles of Mental Representation", Proc Aristotelian Soc 83, 213 – 226 (1983)

merely tacit knowledge. However, while Dennett and Fodor are in essential agreement, one needs to be careful in distinguishing the way they use the term tacit knowledge. In Dennett's terminology, something is explicitly represented if and only if there exists in a functionally relevant place in the system a formula or tokening for which there is semantics or interpretation and a mechanism for reading the formula. Chomsky, in this terminology, would say there was an explicit (albeit unconscious) representation of one's grammar in humans (Fodor would refer to this as tacit knowledge, as would many other writers[6]). Implicit, for Dennett, is defined as what is implied logically by the explicit. Thus Euclid's theorems are implicit in Euclid's axioms. Tacit for Dennett is what underlies explicit. That is what is meant in the phrase "merely tacit knowledge" in the first sentence of this paragraph. It is a mere minimal capacity to represent; it is what underlies things like our depth perception, as discussed above in describing Fodor's position[7]. Not enough is known to say what simple processes embodying Dennett's tacit knowledge are in humans, but the argument is that something that loosely parallels the

---

[6] This major issue of the nature of linguistic knowledge is taken up in Chapter 5. Polanyi, in his extensive writings, does not allow for the clear separation of what he treats as tacit knowledge in the way that Dennett uses the term. For Polanyi, the ultimate basis of all knowledge resides in embodied activity, and tacit knowledge is an essential element in all knowing (M. Polanyi, "Personal Knowledge", New York: Harper and Row, 1964). This issue is tangentially discussed in Appendix B within the framework of Wittgenstein's rule following argument. Polanyi's approach is not directly discussed further, as the contemporary arguments about the legitimacy of "knowing-how" are not easily put into the philosophical framework employed by Polanyi. This is not to deny the importance of Polanyi's early insights, but if the arguments of this thesis can be sustained they are certainly not negated by Polanyi's approach. Polanyi's central idea of subsidiary awareness is implicit in the discussion of grammatical knowledge in Chapter 5, and in the work of various neuroscientists taken up in that chapter.

[7] Thus, for example, it is implicit for Dennett that using Arabic numerals we can multiply as we normally do. This is not implicit in Roman numerals. How do you multiply XIX by LIX within the system of Roman numerals? Many writers would say that for most people the information about multiplication in the first two sentences of this footnote is tacit knowledge.

simple representations via on-off switches in a computer is present in us. As Fodor famously puts it in a later paper than the one cited in 3.1[8], "no computation without representation", but of course, as Fodor acknowledges, not all mental activity is computation.

Dennett states it most clearly when he says,

> "we want our storers and manipulators to be stupider than our understander (of which they are proper parts); otherwise will get into a Rylean regress."

Dennett goes on to describe how a sophisticated system can be built up from the hypothesized tacit knowledge. A simple example is the pocket calculator. There is no explicit representation of the rules of arithmetic. To multiply 6 by 7, the calculator acts on the input instruction and stores 6 seven times in a buffer and stops. Certainly the designer knew the rules of arithmetic but the calculator only has to perform as though it knows the rules.

I hereinafter refer to what Dennett calls tacit knowledge as Basic Tacit and use the term tacit knowledge as it is generally used in the literature, i.e. in the way Fodor uses it, which I described above.

The distinction is important, particularly in the later discussion of the knowing-how of language use (Chapter 5). To spell it out more explicitly:

(i) Basic Tacit knowledge is the "knowledge" embedded in the most elementary unconscious operations of the human mind that enables things like depth perception or olfactory stimuli or the maintenance of blood

---

[8] J.A. Fodor, "The Language of Thought", Hassocks, Harvester Press, 1975

pressure. Basic tacit knowledge is in limited operation even in the clinically brain dead. It does not fall under any usual definition of knowledge.

(ii) Tacit knowledge is not conscious, nor is it merely unconscious; commonly it cannot be brought to consciousness. Fodor puts it this way,

> "If an agent regularly employs rules in the integration of behaviour, then if the agent is unable to report these rules, then it is necessarily true that the agent has tacit knowledge of them."[9]

This allows tacit knowledge to include unconscious knowledge in the Freudian sense. There are differences between this last and what is involved in most knowing-how. Freudian unconscious thoughts can in principle become ordinary propositional knowledge, perhaps through psychoanalysis, i.e., Electra could come to realize she wanted to sleep with her father. Moreover, Electra's unconscious wish to sleep with Agamemnon is integrated with her desires and beliefs; indeed it is the cause of her passion to avenge her father's death. These differences will be discussed further in Chapter 5. There I consider in detail the process of speaking one's native language and claim that it is a prime example of knowing-how.

---

[9] J.A. Fodor, "The Appeal to Tacit Knowledge in Psychological Explanation", Journal of Philosophy, 65, 627 - 640 (1968), p 634; care is necessary here so as not to pre-empt the issue of whether tacit knowledge is knowledge in some defined sense (say as justified true belief). There is no universality of usage but I am explicitly not using tacit knowledge in the way Harman uses it in the following passage (and nor are Dennett or Fodor). "Competence is knowledge in the sense of knowing how to do something; it is ability. It is not the sort of knowledge that can properly be described as "tacit". Tacit knowledge is knowledge that something is the case." (G.H. Harman, "Psychological aspects of the theory of syntax", The Journal of Philosophy, 64. 75 - 87 (1967), p 81). The important point is that the term "tacit knowledge" should not in my usage be thought of as the noun knowledge qualified by the adjective tacit. "Tacit knowledge" is a single expression that I have attempted to explicate in this chapter.

1.3 Rebuttal of the Rebuttal

There is a way of defending Ryle's refusal to accept the intellectualist legend and thus rejecting the Fodor/Dennett position outlined above. One can maintain that Ryle must be right because if, at the bottom of the processes of thought, there are simply ignorant homunculi that distinguish between 1 and 0, that can build "or" and "and" gates and so on, ultimately to produce seemingly intelligent behaviour, something must be missing, at least in the area of beliefs and desires and other conscious thoughts. The fact that a computer can beat Kasparov at chess seems like intelligent behaviour, but these critics would say (i) it tells us nothing about how humans play chess and (ii) it is not intelligent behaviour any more than a computer based expert system is behaving intelligently when it is able to answer recondite medical questions correctly. The argument runs that it is just a huge database, commonly also with some built in heuristics.

In other words, these critics argue that Ryle is right in thinking that the intellectualist model leads to an infinite regress because the Dennett/Fodor type of answer cannot generate human creativity. They agree with Ryle's critique without adopting his logical behaviourism, and the basis of agreeing with Ryle is not that they necessarily share any of his views but that they reject any computational theory of mind. Such a system must follow rules and hence cannot explain consciousness. Thus at the lowest level, there has to be "something" which does contain whatever it is that our high level thinking

contains[10]. Whatever that "something" is cannot be reduced to competencies[11]. Alternatively, and in the same spirit, these critics reject this whole picture of "levels" and, for example, see consciousness and creativity as an unexplained "emergent" property at some sufficient stage of animal complexity.

It is not easy to completely refute this position. As the initial quotation from Fodor says, we know we simply do not understand many aspects of mentality. Consequently, those who claim that there is a something quite essential we don't know about may be right.

However, there is the traditional counter argument of normal science, namely that one works with the current theory with maximal explanatory force. The various forms of the notion of the possibility of accounting for complex behaviour via the manipulation of elementary representations have had some success. Nineteenth century chemistry made huge progress despite the complete ignorance of the fundamental physics that explains the Periodic

[10] See for example, T. Nagel, "Searle: Why We Are Not Computers", in "Other Minds", Oxford: Oxford University Press, 1995, or, in a quite different and extreme form, Whitehead, who says, "Thus an electron within a living body is different from an electron outside it, by reason of the plan of the body" (A.N. Whitehead, "Science and the Modern World", New York: New American Library, 1925, p 76).

[11] A somewhat similar argument is sometimes put as follows: (i) A procedure for achieving some specified result is known as an algorithm if (a) every step of the procedure is moronic; (b) at the end of each step it is moronically clear what is to be done next; (c) the procedure is guaranteed to lead to the specified result in a finite number of steps. (ii) The algorithmic assumption states (a) There exists, or can exist, some algorithmically calculable system exhibiting general intelligence (b) There is a practicable algorithm (known or unknown) for calculating the behaviour of the system (iii) General intelligence in humans includes second order quantifier logic (iv) There are nonsemidecidable sets in second order quantifier logic (v) The algorithmic assumption therefore cannot be accepted. I think it is an open question whether the algorithmic assumption is true, i.e. (v) may not be entailed by (iv). Godel's second incompleteness theorem is sometimes invoked for the same ends. If accepted, these reasonable arguments show only that the computational theory of mind cannot be the full story. They do not show that it is not an important part of the story.

Table, namely quantum mechanics. Indeed, that same chemistry gave us the Periodic Table, the subsequent search for missing elements and interest in optical spectroscopy, and hence ultimately contributed to the discovery of quantum mechanics.

Note that to mount this counter argument one need not disagree with those who say Big Blue is not being intelligent when it beats Kasparov at chess. I entirely agree with this last, and believe that the intelligence resides in the designers and programmers. However, this does not mean that tacit knowledge in Dennett's sense is not the underlying basis for human intelligence and creativity. It means only that the seemingly intelligent performance of computers does not demonstrate intelligence. It is equally true that the occasional seemingly intelligent performance of some university professors does not demonstrate much intelligence; it may merely signify a good education. The role of computers in all of this discussion has been to alter the zeitgeist, and make it apparent that some clever things can be done via simple means and some millions of years of trial and error, together with a feedback system.

1.4 Summary

In essence I have argued that we need not concern ourselves with the precise form of one of Ryle's own concerns in advancing the distinction between knowing-how and knowing-that. Ryle's argument, that there must be knowing-how prior to knowing-that, is essentially correct. However, a contemporary view like that of Dennett or Fodor would not see this as defeating what Ryle called the intellectualist legend in its modern form; a form that postulates that

the lowest form, or more accurately the ultimate basis of, "knowing" is a homunculus of great stupidity.

This last does not detract in any way from Ryle's extraordinary insight. He used a regress argument, which I claim can now be countered by a view arising from the computational theory of mind. This view holds that underlying all our mental functioning, from reflex responses to propositional attitudes via unconscious thoughts, beliefs and desires, there is what I have called "Basic Tacit" knowledge. Further, Ryle advocated philosophical behaviourism, which, as I indicated, I am rejecting. Nevertheless, I argue in the following chapters, that Ryle was importantly and fundamentally right in distinguishing knowing-how from knowing-that, and in seeing it as akin to ability.

CHAPTER 2. THE NATURE OF KNOWING-HOW

> *If I know how to do x, it does not follow that I know how to explain how x is done. There is at least that much distinction between knowing how and knowing that.*
>
> Jerry Fodor[12]

2.1 Introduction

Stanley and Williamson, in a recent article[13], have proposed that knowing-how is a subset of knowing-that, contrary to the view advanced by Gilbert Ryle more than 50 years ago in "The Concept of Mind"[14]. Before putting forward the argument that Ryle was essentially correct, and, moreover, that he was also right in not making too much of the distinction between knowing-how and ability, I will try to isolate as precisely as possible what is being considered as knowing-how. However, even prior to this, it is necessary to make some preliminary remarks about knowing-that.

Knowing-that is propositional knowledge. The typical examples of knowing-that are the capacity to make statements like,

(1) $\sqrt{2}$ is irrational.

(2) Men are on average taller than women.

---

[12] J.A. Fodor, "The Appeal to Tacit Knowledge in Psychological Explanation", Journal of Philosophy, 65, (1968) p 634

[13] J. Stanley and T. Williamson, "Knowing How", Journal of Philosophy, 98, 411 - 444 (2001). However, see also the rebuttal, J. Koethe, "Stanley and Williamson on Knowing-how", *ibid*, 99, 325 – 328 (2002). These papers are discussed in Chapter 8.

[14] G. Ryle, "The Concept of Mind", *ibid*

These are simple statements of fact, the former readily provable, and the latter empirically verifiable. A widely used but contested account of knowledge is that it can be characterized as justified true belief. I have implicitly used such a view in defending (1) and (2) above, as knowing-that[15].

However, there is a similar type of statement that might commonly be expressed in ordinary English as,

(3) Gloria knows how to multiply.

We are then attributing to Gloria some knowledge, which would be categorized as knowing-that. In other words, it is knowledge of the same type as the first two examples – ordinary propositional knowledge. Gloria knows the precise procedure for multiplying, so that given two numbers she can generate the product, and, moreover, she could give an account of this procedure. Thus there are some simple physical feats and some simple or complex mental feats that can be easily described in propositional terms. These, like (3) above, are nevertheless commonly described by the expression knowing how to φ as in (4) (5) or (6)

(4) The children know how to recite the first five letters of the alphabet.

(5) I know how to solve quadratic equations.

(6) Samantha knows how to open the door.

---

[15] See footnote 30 for a brief further remark on the limited use made of the contested JTB theory.

These are examples of propositional knowledge or knowing-that. Samantha, if pressed, could give a precise account of how she opens the door.

The foregoing account needs some amplification. In the simplest cases (1) and (2), it would have been consistent with (4) (5) and (6) to write something like "Patricia knows that $\sqrt{2}$ is irrational" and "most people know that men are on average taller than women". In the case of (3), the real difference is that to express the knowing-that, would require a lengthy formulation describing the propositional knowledge. It might read; Gloria knows that to multiply two natural numbers, one procedure is to write the numbers below one another

abcd

efg

and knowing the decimal system (i.e. that 'a' here represents an integer between one and nine inclusive, multiplied by 1000, 'b' an integer between zero and nine inclusive, multiplied by 100 and so on), the significance of zero, and using the elementary multiplication tables for the integers from one to nine, etc. The long windedness of this formulation is, of course, a partial explanation of why the expression "knows how" is often used to describe instances of knowing-that.

We often use knowing-how when describing some procedure followed by other people as in (7) and (8):

(7) George knows how they serve breakfast at the Mount Nelson Hotel.

(8) Nicola knows how to use a cell phone.

Here, George could give a non-indexical account of the practice and Nicola knows what is in the cell phone instruction manual. In other words George knows that they serve breakfast in a certain defined fashion and Nicola is in possession of a set of facts. Again, despite the use of the term, this is not Ryle's knowing-how. Of course, "George knows that they serve breakfast in the Mount Nelson Hotel", means something quite different from (7). This difference in construction from the manner in which one would rewrite the other examples given above is not uninteresting. In the other cases, with finite rather than tensed clauses, one does not get a well-formed sentence; "Nicola knows that to use a cell phone" requires something further.

To take a slightly more complex case, there are many people who do not know how to solve cubic equations but do know how to solve quadratic equations. The procedure for solving cubic equations is markedly more elaborate than that for quadratic equations, although fully describable. Consider, for example, a school class of 12 year-old students who can solve quadratic equations and this is as far as their studies in algebra have gone, i.e. they have been told no facts about cubic equations. We present them with a suitably simple cubic equation, say $x^3 + x^2 - 10x + 8 = 0$ and ask them to find the roots. The bright children will rapidly find the answer either by substituting a few small positive and negative integers or by drawing a careful graph from say $-5 < x < 5$. If asked how they did it, these children would say (i) we knew it was worth a bit of trial and error because you gave us what we assumed was a problem we could solve and we have learned what the roots of an equation are; or (ii) we knew that we could solve many quadratics graphically and it seemed clear that the methods would carry over to cubic

equations, or something along these lines. The other children would say we did not know how to do it.

For the bright children, we can accurately say they knew that to solve the given equation one followed a certain procedure. For the others, we can say they did not know how to solve the equation because they were unable to make the necessary inferences from the same knowledge base. One again we seem to have a clear case of knowing-that – of knowledge plus inferences from knowledge. Again, note that we can tell little from the use of the terminology "knowing how", but knowing-that contains a great deal more than factual knowledge; it really includes "able to deduce that".

Despite what has just been said, there is the beginning of an element of knowing-how in this last case, which has been masked by the use of the term "bright" in referring to some of the children. There is a certain facility for doing mathematics that some of the children have and others do not, which does not correlate with their general intelligence. Some of those who failed this test would, on other criteria, be brighter than those who passed it and *vice versa.* In some humans there is a knowing-how in general terms about mathematics that others have to a lesser degree, and which is susceptible to training. Moreover, it is also susceptible to psychological and attitudinal factors outside of general intelligence – witness the empirical evidence that girls on average perform less well than boys in mathematics.

Consider an even more complex situation. Centuries ago making a sword was an extremely important skill. If the swords were too brittle they would shatter and if they were too ductile they would bend. There was considerable effort,

much folklore and many legends, like that of Excalibur. There were also individuals whom the kings or emperors recognized as expert sword-makers. By contrast, in the 21st century the control of the elemental composition (carbon content and so forth) and heat treatment is thoroughly understood, so that producing steel with properties thus and so, for a particular purpose, is a question of knowing-that. (This possibility of something that was knowing-how but subsequently becomes knowing-that, is discussed further below.)

Thus it could be said

(9) A top sword-maker in the 14th century knew how to make a sword.

This is an example of knowing-how, which I propose to discuss in this chapter. Some contemporary examples are as follows.

(10) Hanna knows how to ride a bicycle.

(11) Max knows how to do approximate arithmetical sums.

(12) The better clergy know how to interpret the bible.

(13) Children generally know how to speak by the time they are four years old.

2. 2 Identifying Ability

There are a number of distinctive features of knowing-how that might be common cause with Stanley and Williamson and it is useful to enumerate them. However, first it is necessary to say something about "ability" which is commonly contrasted with knowing-how. Indeed, I initially draw this contrast myself, although a more complex position is finally adopted.

There is a sense of "can do" that means no more than physically possible. Nearly anyone in the public house is physically able to hit the bulls-eye on the dartboard and may do so by chance. To demonstrate the ability to throw a dart accurately (I am not talking about the ability to win in a darts championship, for reasons that will become clear later) would require some consistency of performance. Ability includes a simple acquired physical skill. It also includes some things commonly thought of as natural abilities – to walk and run and the like. (We speak of physical disability where these cannot be performed but I need to remark that I am here, at least loosely, defining ability, not describing usage. Many feats characterized as knowing-how are commonly termed ability and vice versa.)

Among the things we are physically able to do are many partially or totally involuntary actions. We breathe, cry, digest, vomit, perspire, maintain more or less constant blood pressure, become sexually aroused (with various physical manifestations) and so on. Perhaps the most extreme of these is the maintenance of blood pressure. This is something we can do when pronounced clinically brain-dead[16], so there is nothing involved that could be described as involving knowledge or even belief. There are things within our physical capacity that are best described as tropistic. Tropisms are simple mechanical or chemical feedback processes, or combinations of such processes, that may sometimes look like organized motivated behaviour. Sunflowers turn towards the sun and many simple organisms execute

---

[16] When surgeons remove the liver from a brain-dead human for the purpose of an organ transplant, the blood pressure and heart rate of the "dead" person rise.

complicated manoeuvres by being sensitive to light and to symmetry. For example, a creature might seem to be "heading towards a light source" if its leg movements slowed down on the side where one of its two eyes received more light.

In the sense in which knowing-how is used in the literature, none of these involuntary actions is considered as knowing-how; this is clearly correct. Since one concern is the distinction between knowing-how and ability, it is simple and proper to remove involuntary or tropistic actions from the regime of ability as they are far removed from the borderline with knowing-how.

One could provisionally define ability as the capacity to undertake an act (or not to do so) as a result of an intention to do so (or not). However, we should not exclude the common possibility of undertaking (or not undertaking) acts, which we could intend, without in fact having any intention. We treat breathing as an ability – normally we do not intend to breathe but we can and do on occasion hold our breath or breathe more rapidly and so on. Similarly, walking and drinking and raising our arms are abilities but digesting is not. Controlling our heart rate is not ability for most of us, but some people have developed this ability.

In the previous paragraph I have provisionally defined ability in terms of capacity, intending to narrow down from the simplest form of "can do" to something narrower. That is, to distinguish between the fact that almost everyone can by good fortune, given sufficient attempts, cause a basketball to pass through a hoop, and the further fact that some individuals can do this with remarkable consistency. The latter I call ability and define it, to put it

loosely, as being able to do something involving a degree of complexity or difficulty when intended or desired, whether intended or desired or not[17].

However, there is a further intimately connected question which is important for a discussion of ability, and, as it will emerge, for knowing-how. This question can be put as follows. What justifies us in saying that someone has the ability to throw a dart accurately? One can define the ability operationally and say that Cassandra can throw a dart accurately when she can, from four metres, place a dart within a circle 10 cm in diameter not less than 80 times in 100 attempts. It does not matter for my purposes that this is clearly quite arbitrary.

The question is, then, how I can justify saying Cassandra can throw a dart accurately, i.e. has the ability. Normally, one says this when Cassandra has given a demonstration, i.e. she has before our eyes hit the aforementioned target, say 80 times from 85 tries, or performed some comparable feat. If we are of a sceptical turn of mind, we might ask Cassandra to repeat the feat on a number of occasions, separated by some days or weeks, just to ensure that there had not been some abnormality. What licenses us to say that Cassandra can throw a dart accurately? There is a difficulty here, as has been pointed out by Thomas[18]: there is a category difference between the episodic

---

[17] It is clear that there will, of necessity, be a statistical element in this. One can have ability without being able to successfully exercise it every time we try to do so.

[18] G. B. Thomas, "Knowing How: A Study of the Concept of a Human Ability", PhD Thesis, Harvard University 1963. I am indebted to this work by Thomas for drawing my attention in great detail to the relevant portions of the work of Wittgenstein on this topic. Thomas develops these ideas of Wittgenstein in an interesting and extensive manner, but to a conclusion that I believe is false. Nevertheless, my own brief remarks here owe a great deal to the careful work of Thomas.

performances we have witnessed and the time independent assertion that there is an ability. We say Cassandra can throw darts accurately, or has the ability, or let us say, "C can t", between performances. There are, of course, circumstances that disallow C can t; she may be injured or be flying to London and have no darts and so on. However, in the absence of such constraints, what is the general justification? Is it simply induction? It is induction if we predict, after 85 on-target attempts out of 90, that the next attempt will probably succeed (if and only if Cassandra is still trying). However, when we say, "C can t", some years after last seeing Cassandra, we are saying, "If she wants to, C can t, because I have seen her t-ing." Or, alternatively, it has been reliably reported that she did demonstrate t-ing on some occasion in the past.

The first of these is an inductive inference about an event, namely, the next throw by Cassandra. That is, the inference was made from 85 events about another event. However, to say "C can t" is not to predict an event by induction, it is to ascribe something to C.

A parallel might be observing many throws of a dice and then asserting it has one chance in six of the appearance of say a three. In other words, one is ascribing to the dice, whether rolled or not, the "ability" to lie with three uppermost with the probability of one in six.

What is therefore required is some argument for a comparable regularity between the dice and a human agent. Clearly, this extreme is not attainable, nor does it exist. However, Wittgenstein has argued that our use of terms makes use of behaviours, and this "form of life" is a necessary part of our ascription of ability to people from observations of their performance.

Wittgenstein does not discuss this topic in a direct way, but I propose to now argue that it is a reasonable inference from some reasonable assertions in the "Philosophical Investigations".

The point that Wittgenstein makes is that although no finite number of past observations that C can t proves that C can now t, this does not undermine justification. As he puts it[19], "Justification by experience comes to an end. If it did not it would not be justification" or[20], "What people accept as justification - is shewn by how they think and live." In other words, we are entitled to say that "C can t" from the evidence of past performances because that is what we do. The ascription of ability has a place in activities, customs institutions. Wittgenstein makes the point most clearly in an account he gives in the Brown Book[21] of a hypothetical tribal activity. The members of the tribe place bets by laying a piece of gold under the picture of the competitor they believe will win. They have their stake doubled if he does win and lose it otherwise. The members of the tribe need have no terms in their language for ability. However, the punters will favour a competitor who has shown greater strength in the past.

In what follows I therefore treat as unproblematic the ascription of ability and knowing-how on the basis of past performance (in those circumstances where

---

[19] L. Wittgenstein, "Philosophical Investigations" 485.

[20] L. Wittgenstein, "Philosophical Investigations" 325.

[21] L. Wittgenstein, "The Blue and Brown Book", Oxford: Blackwell, 1958, 58

there is not any special reason to suppose otherwise)[22]. Ryle himself had no qualms on this score; he took the matter for granted. In his characteristic style he says[23]

> "There is no one signal of a man's knowing how to shoot, but a modest assemblage of heterogeneous performances generally suffices to put it beyond reasonable doubt whether he knows how to shoot or not."

2.3 Characterizing Knowing-how

Returning now to the consideration of knowing-how, I want first to enumerate some *prima facie* relevant characteristics of knowing-how. The aim here is then to arrive at a definition of knowing-how. I do this in section 2.4.

It is worth making one further preliminary remark. I use "knowing-how" throughout for a major subset of what is commonly called tacit knowledge and it is this subset only that concerns me here. Many writers[24] include as tacit knowledge the fact that we can give answers to questions we have never previously thought of and these are answers that have never been conscious. For example, we know that the answer to the question, "Do elephants in the wild wear ear muffs?" is "No". It would seem better to refer to this tacit propositional knowledge as latent knowledge. This latent knowledge will not be considered further.

The more obvious characteristics of knowing-how, then, are as follows.

---

[22] It is perhaps too quick to regard this as unproblematic. There may be approaches to this issue via probability theory.

[23] G. Ryle, "The Concept of Mind", *ibid*, p 46.

[24] See for example the textbook M. Devitt and K. Sterelny, "Language and Reality", Cambridge: MIT Press, 1999, p 175.

(a) Knowing-how normally requires ability, but it is not synonymous with ability.

One can know how to play the piano and become crippled with arthritis and lose the ability - such a person still knows-how to play the piano.

(b) Knowing–how is intentional

Being able to recognize that one is in pain, to walk or to eat food is ability, but none of these is knowing-how. These are merely examples of events that can be caused by an agent. The case of pain here is somewhat different from the other two examples in that one is conscious of being in pain, but normally one does not bring this about. One can know-that one is in pain just as one can know-that one is digesting; the former is conscious and the latter is a learned thing; a pig would not know that it was digesting. However, to detect that one is in pain is simply an ability that we share with pigs.

I have used "intentional" here in its lay sense of it being the case, where knowing-how is claimed, that the knower plans or desires or wishes that what she does should happen. Of course this is commonly true of many abilities, like walking; the claim here is that it is a necessary feature of knowing-how.

(c) Knowing–how is learned

Walking is ability, but it would develop in a normal human without the parental encouragement commonly devoted to it. To be able to walk is not an instance of knowing–how. Of course the skill of the winner of an Olympic walking race is quite different and is indeed knowing-how. If someone without training won an Olympic walking race by virtue of natural aptitude, i.e., just by walking as

rapidly as she could, this also would simply be ability (although of a high order). Thus training and practice and often also coaching commonly enhance knowing-how[25]. Of course, coaching or teaching can develop knowing-that, but the subtle difference between coaching and teaching is relevant here. Usually one coaches students to pass an exam in quantum theory, whereas one teaches them quantum theory. Some of the former is knowing-how and the latter is knowing-that, although applying it to new problems involves knowing-how as well[26].

I have inescapably used learned somewhat loosely here. One does learn to walk, and one learns many things in life without explicit study or teaching, often merely via social interaction. My intention is to indicate the contrast between two extremes: (i) acquiring knowing-that by simply being told, for example, that π is the ratio of the circumference of a circle to its diameter, and (ii) acquiring the knowing-how of skiing by the painful process of instruction interspersed with numerous falls.

(d) Demonstrating knowing-how is necessarily an extended process

Knowing-how involves being able to execute a process or series of processes. It would not demonstrate that one knew how to do approximate

---

[25] Nothing that has been said implies that coaching is all that is required to develop high order skills. A good violinist needs good hearing and the ability to discern small differences in pitch.

[26] I have said in (13) above that speaking one's native tongue is knowing-how. This is probably the best-known example of tacit knowledge or knowing-how. Chomsky denies that this is knowing-how, but does vigorously argue that we have tacit knowledge of grammar and we learn the labels in our native tongue for various basic (pre-scientific) concepts. It is difficult to claim that a four year old has propositional knowledge of the rules of grammar that she so conspicuously uses (and occasionally flouts) in her daily speech (This major and controversial topic is taken up in detail in Chapter 5.)

arithmetical sums with one instance of such ability – such as 13 x 14 < 300. In contrast, to say "√2 is irrational" is to show that one knows this fact. Of course, there are many examples of knowing-that where repeated performance is also required to show the knowledge. To show that one knows the multiplication table up to x12 would require a number of examples of correct performance. The point is that this "process nature" is a necessary feature of exhibiting knowing-how.

(e) Knowing-how can be a matter of degree, like knowing-that and ability. However, establishing an empirical ranking is often problematic in the case of knowing-how, unlike ability and knowing-that.

Consider the proposition, "I know how to play tennis", uttered by various players. It may be the case that Leyton Hewitt has less propositional knowledge about tennis than Pete Sampras or Rod Laver. He currently has more ability. What about his knowing-how? It is clear we could empirically rank these three individuals in regard to their propositional tennis knowledge (via a questionnaire), and their ability could readily be ranked on the court. In contrast, we cannot rank their knowing-how by any empirical test[27]. The reason is that Rod Laver could have more tennis knowing-how than the others but age prevents him from exploiting this. What I mean by knowing-how, in addition to being able to serve, volley and so on, are things like court craft,

---

[27] I will later discuss some cases where the degree of knowing-how can in fact be ranked. One example is the use of one's native language, where experiments have been conducted on speakers who were not given an opportunity to learn to speak in the so-called sensitive period of childhood. Marked differences in knowing-how to speak are found in matters like verb agreement and inflections for aspect and number.

temperament, capacity to sense opponents' weakness, willingness to change strategy at the right moment, capacity for intense concentration and so on. These last are all means of knowing what to do rather than knowing what to think, and are the essence of knowing-how[28].

If one is to regard tennis playing as a single activity rather than as an amalgam of serving, volleying, forehand etc., it might be best to speak of winning at tennis as the activity, i.e., to give a teleological account. This seems not to alter the argument just made. The knowing-that of "winning tennis" is the subject of numerous books. Attentive study of such works has

---

[28] Although Ryle takes knowing-how to be intelligent performance, there is some ambiguity in his position as has been pointed out by Scheffler (I. Scheffler, "Conditions of Knowledge", Chicago: University of Chicago Press, 1965, p 100). Ryle clearly states that knowing the rules of chess is not sufficient to know-how to play. But then Ryle says in his book, "The Concept of Mind", *ibid*, p 41, even if a person cannot cite the rules that person does know-how if, "he normally does make the permitted moves, avoids the forbidden ones and protests if his opponent makes forbidden moves."

Even if this player cannot cite the rules, the mere fact that he confines the black bishop to black squares and so on, is not the knowing-how of chess play, which has to do with strategy and other high order skills of chess playing. This is not a simple slip on Ryle's part. In his other major writing on this topic ("Knowing How and Knowing That", Collected Papers, Volume 2, London: Hutchinson, 1971) Ryle says on page 217, "Knowing a rule is knowing how. It is realized in performances which conform to the rule, not in theoretical citations of it."

I think Ryle is roughly right to take intelligent performance, whether of skiing or cycling or chess playing, as knowing-how. The case of following simple rules is probably best excluded from knowing-how. Thus, if I can tie a simple knot, having received instructions, but cannot give the instructions, it seems to me nevertheless a case of knowing-that; I place end A over end B and then over A and so on. Similarly, if I drive on the left-hand side of the road (in England or South Africa) I do so because I know that it is legal. Knowing-how is not manifested in this aspect of my driving, but simply knowing that.

I pursue a precise definition of knowing-how in section 2.4, which follows. However, I think Ryle is confusing the issue here, in his eagerness to argue that all performance is knowing-how, to make his case against the intellectualist legend. This is most easily seen by imagining a very boring game consisting solely of two-players, alternately tossing a coin, and the only further rule is that the first to score heads four times in succession is the winner. The game is devoid of knowing-how and the performance of the game requires only the knowing-that of the two rules.

failed to give most people either the ability or the knowing-how; they still double fault commonly (lack of ability) and even more commonly at crucial points (lack of knowing-how).

(f) Knowing-how cannot be cast into propositional form

The medieval sword-maker discussed above (p 24-25) has a store of experience and thoughtful observation that enables her more reliably to deliver the goods, albeit not infallibly. But she does not know what it is that leads to a good sword. We now know she did not know[29]. She can, of course, give a general account of her practice. However, she varies it in various ways from considerations she is unable to articulate, but which on each occasion she nevertheless feels is a correct modification.

(g) Knowing-how does not essentially involve belief

In this section, where I characterize knowing-how as a precursor to giving a definition, I am partly implicitly contrasting it with propositional knowing. In doing this I am using, as a basis, a justified true belief account of knowledge. This account is contested and widely believed to be incomplete[30]. However, most, but not all, of the theories of knowledge claim that propositional

---

[29] I have thus acknowledged that what was once knowing-how can become knowing-that. Hence, my saying, knowing-how cannot be cast into propositional form, means only that if someone is performing by knowing-how, she cannot herself currently cast what she does in propositional form.

[30] The JTB theory has been challenged by the Gettier problem (E. Gettier, "Is Justified True Belief Knowledge?" Analysis, 25, 121 – 123 (1963)) and various solutions have been proposed, such as Nozick's truth tracking (R. Nozick, "Philosophical Explanations", Oxford: Oxford University Press, 1981) and various reliabilist approaches. The resolution of these differences does not impinge directly on my account of knowing-how.

knowledge (knowing-that) does require belief. In other words, most theories of knowing-that require that it involves taking a distinctive attitude – a belief – towards a proposition[31].

We can believe that we know how to play tennis. This is a belief in a proposition. Knowing-how does not consist of propositions. In this respect it is like ability. There is no belief involved in my ability to walk. Of course, things are a little more complicated because knowing-how normally involves a modicum of knowing-that. I must know the rules of chess as part of knowing-how to play and good chess players invariably know a good deal more[32]. To this extent, there are beliefs. However, there is no proper analogue in knowing-how of the false belief that $\sqrt{2}$ is rational, which is not knowledge. Someone who rides a bicycle in a dangerous fashion, and falls off occasionally and is somewhat inept, is commonly said not to know how to ride. But this is not a denial of "she knows how to ride a bicycle", in the Rylean sense. If someone knows how to φ, even if not especially well, they know how. Ignorance of the knowing-that component of tennis might disqualify one from knowing-how (if one thought the aim was to hit the ball into the net) but

---

[31] Welbourne (M. Welbourne, "Knowledge", Bucks: Acumen Publishing Limited, 2001, p 110 advances an example of a theory that does not do so. He says, "Knowledge is just our word for what a speaker communicates and a hearer accepts in episodes of telling [in favourable circumstances]." I mention this dubious theory only to make the point that knowing-how is distinct from knowing-that, even in the case of theories of knowledge that do not take knowledge as a form of belief. Knowing-how is not knowledge at all on Welbourne's theory.

[32] It is usual among good chess players to have an encyclopaedic knowledge of the classic openings and end games and often the whole of famous matches. This vast store of knowing-that is no substitute for knowing-how, though it is of great value. The knowing-how of chess is complex. It has been shown that grandmasters and masters can recall positions of some two dozen pieces presented for a few seconds almost perfectly, while weaker players can do this for only about six pieces (A.D. de Groot, "Thought and Choice in Chess", The Hague: Mouton Publishers, 1978). Note that this is a skill and has nothing to do with the role of memory in an occurrent game of chess.

knowing-how comes only in degrees, from Serena Williams to the local schoolboys to zero.

Scheffler[33] has pointed out that there are circumstances where use of the terms knowing and believing part company, even in the case of propositional knowledge. Many fundamentalist students know the theory of evolution because they are taught the theory. However, they do not believe it. Of course, these students do not themselves believe that chimpanzees and humans have a common ancestor, and would themselves say they did not know this, since it is not a justified true belief (or does not conform to whatever theory of knowledge they happen to hold). I might say they know but do not believe; they would say they neither know nor believe, but have learned that others believe falsely. While Scheffler's point is quite correct, this usage of know does not alter the general contrast made above between knowing-how and knowing-that.

2.4 Defining Knowing-how

Carr[34] has given a careful account of knowing-how and has proposed that the following conditions must be met:

A knows how to φ only if:

(1) A can entertain φing as a purpose[35]

---

[33] I. Scheffler, "Conditions of Knowledge", Chicago: University of Chicago Press, 1965, p 15.

[34] D. Carr, "Knowledge in Practice", American Philosophical Quarterly, 18, 53 – 57 (1981), p 57

[35] We can speak of knowing how to avoid life's pitfalls or knowing how to behave decently or knowing how to act at a vice regal reception or knowing how to tie a reef knot. We have here a series of uses of the term knowing-how, which are successively more specific. The last I

(2) A is acquainted with a set of practical procedures for successful φing

(3) A exhibits reasonable success at φing[36]

This definition would require slight modification (and it is clearly stated in Carr's article that he is well aware of this) to incorporate the arthritic pianist. Perhaps "has exhibited" could be substituted, in (3), for "exhibits". However, from the present point of view, a fourth clause, which is entirely in the spirit of Carr's analysis, is required:

(4) A complete propositional account of φing is not available to A[37].

This last is necessary to exclude feats like knowing how to recite the first four letters of the alphabet.

Finally, in relation to defining "knowing-how", I would add a further condition:

(5) A must have learned to φ via a process of instruction which may include self-instruction (and may be an unconscious process)[38].

---

would classify as knowing-that (left over right then right over left). The first is so diffuse as to defy characterization in much less than a full-length novel. Edith Wharton's heroine, Lily Bart, in "The House of Mirth", being trusting, foolish and yet principled, clearly does not know how to avoid life's pitfalls. I want to encompass by, "A can entertain φing as a purpose", some definable act or set of acts like riding a bicycle, playing tennis or having a conversation in Italian, and want to exclude global aims like avoiding life's pitfalls. The line here is however fuzzy, in Chapter 8 I briefly discuss knowing-how to keep a secret.

[36] A minor caveat might be that this is not absolutely necessary. We can imagine a skier who has lived entirely in extremely mountainous terrain and has never skied on a gentle slope. We would commonly acknowledge that she knew how to ski on gentle slopes.

[37] A partial propositional account can be given of, say, skiing; no account would enable someone to ski immediately. It is difficult to give a general statement of what constitutes a "complete account".

[38] The point of this elaborate restriction of the nature of learning is to exclude the cases where Gloria has learned to multiply 12 times 13 by being told that the answer is 156, or where she learned to recognise blue by being shown a piece of blue paper and told it is blue. Knowing-how is not acquired in these instant ways. However, it is not easy to adequately define a "process" of instruction.

Carr might well argue that his (1) (2) and (3) imply (5). However, while he explicitly excludes opening one's mouth as knowing-how, his three minimal conditions do not. It seems that condition (5) excludes nothing that would fall under a current understanding of knowing-how.

It should be noted that this definition of mine is not a behavioural definition and hence might not accord with Ryle's philosophical behaviourism. However, as I remarked earlier, Ryle's insight about knowing-how is not beholden to his philosophical behaviourism.

It is useful at this point to consider in general terms the adequacy of the definition I have proposed. There are several possibilities.

(a) The definition fails by including too much

(b) The definition fails by excluding too much

(c) The definition is sound

(d) The definition is incoherent, as the concept of knowing-how is incoherent.

In succeeding chapters I consider various candidates for knowing-how and examine whether they meet the definition. In other words, I try to establish that my defined set is not the null set and hence does not fail totally as a result of (b) above.

With regard to the possibility of including too much – no instances of knowing-that can fall under the terms of the definition since they are excluded via clause 4[39]. However, Hawley[40], who has a definition of knowing-how in terms of a counterfactual success condition plus warrant, has an interesting Gettier-like[41] case for knowing-how. She says,

> "First, consider Sally, out for an ill-advised winter walk in the hills. She has no idea what to do in the event of an avalanche. When an avalanche occurs, she mistakes the snow for water, makes swimming motions, and - luckily - escapes the avalanche, since in fact the way to escape an avalanche is to make swimming motions. Sally satisfies the counterfactual success condition for escaping avalanches - if she were to try, she would succeed, for she is prone to mistaking snow for water. Yet it seems that her success is really a matter of luck: Her two failings (not knowing-how to escape from an avalanche, mistaking snow for water) cancel each other out. Despite her success, she does not know how to escape avalanches."

This sort of example is ruled out by my definition via clause 5; Sally has not learned how to escape avalanches. Indeed this seems to be a nice legitimation of clause 5. Imagine I have learned and know-how to play the flute. Imagine also that there is some activity, among some isolated islanders, whose purpose is rain-making, and this is done silently with a bamboo tube with markings spaced as on a flute and with motions identical to flute playing. It would seem counter intuitive to say of me, who had never even heard of the activity, that I knew how to do it. The argument, however, is inconclusive. Hawley might argue that this is nevertheless Gettier-like, since I do in fact

---

[39] This point is discussed further in Chapter 8, when considering the paper, already mentioned, by Stanley and Williamson.

[40] K. Hawley, "Success and Knowledge-How", American Philosophical Quarterly, 40, 19 – 31 (2003), p 27

[41] E. Gettier, "Is Justified True Belief Knowledge?" Analysis, 25, 121 – 123 (1963)

know how to instantiate the islander's practice. I merely do not know that I know.

I have not included ability in the definition at this stage, because I believe it is implicit in the definition as given. My position is that *prima facie* it is implicit as a result of clause 4 in the definition - that a complete propositional account of φing is not available to the person who knows-how to φ – while meeting the other clauses which establish the activity as do-able and purposive. For example, the able cyclist cannot give such an account of cycling. The issue of whether knowing-how is a subset of ability is discussed in detail in Chapter 3.

However, there are ways of failing to be ability while meeting clause 4. There are activities for which no full propositional account can be given, and I would claim they are not abilities. Telepathy and clairvoyance are undertaken as purposive activities and their proponents do not claim to give a full propositional account of how they are performed. However, I would argue that these do not meet the other clauses of the definition. If they were shown to be real abilities (i.e. they were performed successfully), it seems they would have to be accepted as examples of knowing-how (unless there was a remarkable parallel advance in our knowledge of the underlying science enabling these arts to be fully described). Another example is given later in Chapter 7, that being kind, an example of a virtue, arguably fails to be ability because there is no such settled state. Individuals are kind in a set of circumstances conducive to their being kind; they do not know-how to be kind in general. In contrast, being able to play chess or ski or cycle are settled human configurations. This may, however, only suggest some minor modification of clause 3, perhaps

along the lines of requiring consistent performance over a period of time and in diverse conditions.

I have said nothing about the definition including too little. That is, are there clear examples of knowing-how that do not fall under the definition? The purpose of the present work is to establish whether Ryle's knowing-how is a useful concept, and I believe an instance of this last sort is highly unlikely, although it would be welcome in this context. In any case, I have been unable to invent such an example and I do not pursue this point further.

The coherence of the concept of knowing-how, an issue raised under (d) above, is, of course, the topic of the remainder of this study.

2.5 Consequences of the Definition

A feature of the definition (1) to (5) above is that it excludes from knowing-how a situation where (i) A has never been able to φ or (ii) is unable to φ except where prevented from φing by a physical or mental disability incurred after acquiring the relevant knowing-how. Stanley and Williamson write[42] of a ski instructor who knows how to perform a complex manoeuvre but cannot do it. I do not believe this is knowing-how (or at least it should not be assimilated with what Ryle wrote about, i.e., it should not be included in a discussion of knowing-how in the sense of someone who is able to execute the manoeuvre) All that this erudite skiing instructor can do is to describe in propositional terms what the manoeuvre consists in and make some helpful suggestions.

---

[42] J. Stanley and T. Williamson, *ibid*, p 416.

She clearly doesn't know-how because she can't, although physically fit and *compos mentis*. If someone suggests that she knows how to play tennis although she has never played, she can only mean (i) she knows the rules and (ii) she knows from observation that it is good to get to the net, to get ones serve deep, to be swift about the court and so on. She could write a programme for a tennis video game but that is not knowing-how. Many of the Koreans who packed the terraces at the 2002 World Cup do not know how to play soccer in the sense that Ryle used the term: they do all, however, believe they know that the way X should have scored is thus and so. This point is important, and I believe it is the source of much confusion. It is quite true, as I have already mentioned, that the expression knowing-how is commonly used for describing a situation of knowing-that. For example, like Stanley and Williamson, Alter says[43], "A guitar teacher may know how to play well, but lack the ability to play well." What is meant here is that the guitar teacher (i) can recognize good guitar playing; (ii) can successfully instruct others, with the talent she herself lacks, to play well; and (iii) could give a propositional account of how one (if one has talent) could learn to play well. She does not have what Ryle means by knowing-how to play the guitar. We could correctly say, in Ryle's sense, that she does, for example, know-how to recognize good guitar playing.

On one point my definition distinguishes my position from that outlined by Carr. He reserves knowing-how for physical attributes such as bicycle riding.

---

[43] T. Alter, "Know-How, Ability, and the Ability Hypothesis", Theoria, 67, 229 – 239 (2001), p 232.

He argues that ability and knowing-how coalesce in the case of mental ability. He rightly says that the mental ability to multiply and what is called "knowing how to multiply" are one and the same. However, the ability to multiply is knowing-that and this is the reason why one would not in normal circumstances say, "She knows how to multiply but cannot multiply". In contrast, it is perfectly proper to say, "She knows how to ride a bicycle but cannot because she broke her leg". However, it is also quite proper to say, "She knows how to speak Chinese but cannot since her terrible accident". Indeed it is well established that cognitive abilities of this sort (including, of course, the example above of multiplication) can not only be lost due, say, to a head injury but can return without any training or learning. Any skilled physical activity like skiing clearly involves the central nervous system. The false distinction Carr makes arises from an essentially trivial source. Namely, the commonly visible known causes of physical inability, like a skier with a broken leg, contrasted with our substantial ignorance of the details of brain damage phenomena.

Thus, *contra* Carr, I am saying that both mental (chess playing) and physical (cycling) abilities can fall under knowing-how[44]. Further, both of these, and also knowing-that, can be present and yet cannot be manifested. For example, being sufficiently inebriated would incapacitate one from all of chess playing, cycling and elementary multiplication. For this reason, his example of

---

[44] I do not intend to imply that these are sharply distinct categories, fencing or boxing requires thought and dexterity.

a particular mental ability where it seems strange to say one has the ability but cannot manifest it is not compelling.

Some philosophers might object at this point that if one includes mental ability such as chess playing within knowing-how, which indeed I do, there is a major difficulty. Namely, if Deep Blue has been successfully programmed to beat Kasparov at chess, does this not prove that chess playing, like elementary multiplication, is knowing-that? The crucial point here is that Deep Blue does not play chess the way Kasparov plays chess. Deep Blue is full of knowing-that, including some heuristics[45]. In speed of elementary computations Kasparov is totally unable to compete with Deep Blue and does not do so. We do not know the details of how able human chess players play[46], any more than the medieval sword-maker knew the details of how to heat-treat messy

---

[45] To labour the point made earlier, it would be usual for someone to say, "Deep Blue certainly knows how to play chess – it can beat Kasparov". This, however, is not to say that Deep Blue has knowing-how in Ryle's sense.

[46] It is true that there is considerable overlap between some of the strategies built into Deep Blue and the way humans play chess. It could hardly be otherwise. For example, the pruning αβ algorithm whereby Deep Blue makes most use, in surveying future moves, of responses to the opponent's best next move as computed, is something that competent chess players will naturally do.

Kasparov in 2003 played and drew a $1 million match series against Deep Junior, an Israeli computer. He has looked at the software and remarked, "Junior is very humanlike. It's a computer version of me. It plays forcefully, imaginatively and takes risks." (New York Times January 21 2003) Although Deep Junior has more strategic programming and less speed than Deep Blue, it can nevertheless examine more than 100 000 times as many possible moves per second as can Kasparov. After the first game, which Kasparov won, one of the two Israeli programmers who designed Deep Junior, Shay Bushinsky, said the machine realised it was in a difficult position when it took 25 minutes to make its ninth move. None of these anthropomorphic remarks alters the fact that the computer is entirely reliant on knowing-that.

The very poor performances of computer efforts to play the ancient Asian game of Go are interesting in the context of knowing-how. An ordinary human player easily beats the best programmes. Top Go players can readily recognise a good move but have great difficulty in saying why it is a good move. Along with intuition, pattern recognition is a major part of the game.

dilute iron alloys. Certainly, an intelligent adult can play noughts and crosses (tic-tac-toe) like a computer would, because she can readily fully analyse the game. This is not so with chess.

This last is relevant to the question raised earlier of knowing-how becoming knowing-that, as understanding progresses in a particular area. In most theories of knowledge, it is accepted that empirical knowledge of a particular time, such as the knowledge that the earth is flat, or that space is Euclidean, may cease to be knowledge. And the same fate will doubtless befall some parts of what is currently accepted as knowledge. In the same way, we may come to know precisely in propositional terms how to perform what was once an art. I have cited sword making but there are numerous other examples. The manufacture of medium quality white wines (post the actual grape production) can now be entrusted to a capable chemist, rather than a grizzled, fourth generation, French vintner. Antibiotics have made routine the treatment of a range of bacterial diseases that were once best handled by wise and experienced doctors, who brought their know-how of long practice to the attempted cure. It is important to distinguish here the Deep Blue case, where the "knowing-that" computer may be the best chess player in the world, but the human champions still use knowing-how since they are unable to emulate the computer, from cases like the sword manufacture, where a propositional account can now be given and used directly by a human agent.

This argument might raise the question of where to draw the line between knowing-how and knowing-that. Suppose a medieval sword maker had been extremely successful throughout her life. Would one say she knew-that? On

my present definition the answer is no. She does not have a propositional account of her art. Consider a chess player who not only became world champion, but also absolutely dominated the game, and won every match in fewer than 40 moves, with only twenty minutes on her clock. I think we would ask her if she had some new knowledge. We might ask, "Have you completely analysed the game? Have you some entirely new strategy? Why do you think Kasparov was beaten by Deep Blue in those three games and what would you have done at White move 16, from which point Kasparov seemed to face inevitable defeat?" Unless we received some entirely unexpected answers one would just say she was a knowing-how prodigy.

Consider another case to examine the boundary between knowing-that and knowing-how. Take the series 1, 5, 11, 19, 29. ... One might say that for someone to say "the next term is 41" is a case of knowing-that. The person either knows that $a(n) = n^2 + n - 1$, or she knows that the successive increments are 4, 6, 8, 10, 12 and so on, or she knows both facts, or some other facts. Suppose she is asked for the next term when presented with the series 1, 9, 55, 193, 501... .She must know the convention that the next term is to be derived by a procedure that yielded the previous terms from their predecessors. That is, I am excluding, by a well established mathematical convention for series, the perfectly logical possibility that this series is, in fact, 1, 9, 55, 193, 501, 1, 9, 55, 193, 501 and so on. This last convention is another piece of knowing-that. But how does she proceed? She might look at the difference between the terms, obtaining the unhelpful-looking 8, 46, 138, 308. She might then say, looking again at the original numbers, "let me subtract one from all of them", obtaining 0, 8, 54, 192, 500. Now she thinks,

"let me divide the second by 2, the third by 3 and so on", obtaining 4, 18, 48, 100. "Now I can go on!" she says[47], we have 1 x2^2, 2 x 3^2, 3 x 4^2, 4 x 5^2, so the next number is (6 x 5 x 6^2) + 1 = 1081. And so she deduces $a(n) = n^4 - n^3 + 1$. This last process for the more complicated series, I would argue, is in part knowing-how, and the border between knowing-that and knowing-how in this case lies somewhere between the two examples[48]. The point is that, in the more complex example, one could not give an account in propositional terms, prior to doing the task. All one has is the notion of looking for regularities, and, in lay terms, some mathematical nous (or in philosopher's terms knowing-how).

My account of knowing-how is based on the observation that humans can acquire skills, without a full understanding of their achievement. Trial and error allows us to improve our performance in many domains without conscious thought. Is knowing-how commonly and often, although certainly not always, a mundane form of human creativity? Or, to approach the question from an entirely different angle, do chimpanzees exhibit knowing-how when they successfully use an instrument to collect ants from a nest or learn to ride a bicycle, as they readily do in captivity? There are complex questions here because some would argue that at least the first was ordinary knowledge and others would deny knowledge to creatures without language. The chimpanzees certainly meet my criteria (2) to (5). If non-human animals

---

[47] L. Wittgenstein, "Philosophical Investigations", Oxford: Blackwell, 1953, 151.

[48] It is not necessary for our argument that the border lies in the same place for different individuals. There are mathematically minded individuals for whom the second series is as obvious as the first.

display intentionality then they can acquire knowing-how. One might be reluctant to ascribe creativity to a chimpanzee, but one's intuitions are not of much value. On the other hand, it is usual to ascribe knowing-how to an outstanding opera singer.

The chimpanzee's performances are variously regarded in the literature. Once again we are led back to the complexity of the performance; opening a door is knowing-that, brilliantly dancing the principal role in Swan Lake is an exhibition of knowing-how. How is the line to be drawn? The heart of the matter in my account arises from criterion (4), that a complete propositional account of φing is not available to the φer. Clearly, animals are included by this criterion (although excluded where there is failure of intention). Small children playing noughts and crosses are included as knowing-how, whereas in their noughts and crosses playing, intelligent adults know-that, and are excluded.

On the definition (1) to (5), whether animals exhibit knowing-how becomes the usual issue of questioning animal intentionality. The remaining difficulty is then only to give as precise an account as possible of the divide between full propositional knowledge of a process and its absence. I have suggested, in talking about the hypothetical wizard at chess, that it is not a question of extremely impressive knowing-how. The boundary will, as in many other problems, be fuzzy. At one extreme (confining ourselves to processes, rather than obvious propositional knowing-that of the "√2 is irrational" type) are things like the ability that most of us have to multiply the natural numbers up to 12. Many adults can do this because they have memorized the relevant

multiplication tables – a clear case of knowing-that. At the other extreme is the old chestnut – knowing-how to ride a bicycle. What about a beginning bridge player? She has learned the rules and gradually learns how to play. According to the present argument, she knows how to play bridge when her experience has given her skills that transcend the advice she might get from a book on "How to Play Bridge". The book would say things like, "third hand plays high, declarer must consider entry problems initially, against no-trumps the opening lead should be" etc. She knows-how to play bridge when some of these generalisations are not treated as rigid rules, when she allows a host of subtle factors to determine her play, which she could not give a weighting to even if she could enumerate many of them.

I have so far made no distinction, but there is a difference, between playing bridge or chess and riding a bicycle. This difference should be considered. The former two are done within the framework of rules, whereas the latter is determined essentially by function (we can consider riding a bicycle somewhere other than on a public highway, where there are some rules but this is a triviality). As Wittgenstein famously says[49],

> "To obey a rule, to make a report, to give an order, to play a game of chess, are customs (uses, institutions)."

Thus there is an inescapable element of knowing-that in chess or bridge. We could simply give a child or a chimpanzee a bicycle, and they would by trial and error learn to ride and become progressively more efficient, either in play

---

[49] L. Wittgenstein, "Philosophical Investigations", Oxford: Blackwell, 1953, 199.

or because they had a need or desire to cover distance more readily than on foot[50]. Giving either character a pack of cards or a chess set would not result in bridge or chess. Does this difference make any difference to the conception of knowing-how? It seems not, as, given the rules of the games, these rules determine functionality, so that for the knower there is no real difference. Conversely, society can impose customs on a functional activity, as in the different styles of horse riding once imposed on "ladies" and "gentlemen".

2.6 Two Sorts of Knowing-how

Both the description and the definition of knowing-how have left open two very different understandings of the phenomenon. In the example of the 14th century sword manufacture I have said that there is a way of making good swords such that a propositional account could be given. In the past this was not known and hence knowing-how was required. In the case of speaking one's native tongue there may be rules, which we do not know, but it is also possible that the most expert linguists' collective understanding is an expression, in approximate rules, of a process that cannot truly be described by rules. The language faculty may have evolved by evolutionary processes

---

[50] Dummett attempts to make a sharp distinction in the following way. Suppose you have never ridden a bicycle and you are asked whether you can ride one. An intelligible response would be to say, "I don't know, I have never tried". If however you were asked if you could speak Spanish, never having any sort of exposure to Spanish, this would not seem a proper response. It is not the sort of thing one could try without learning (M. Dummett cited in M. Luntley, "Contemporary Philosophy of Thought", Oxford: Blackwell, 1999, p 91). However, if one were asked about whether one can play billiards, never having been exposed to the game, the response, "I don't know, I have never tried" does not seem entirely foolish (one might be a snooker player). My point is that there is no knowing-how without some knowing-that, and (i) this is a matter of degree, and (ii) knowing-how covers an immense range, as will emerge throughout this thesis.

so complex that no set of consistent rules can properly describe it. An analogy might be the "design" of an aeroplane in a wind tunnel by trial and error so as best to accommodate, for optimal performance, various conflicting rough "rules" (really empirical generalizations). There then may be no rules to be known. Or, to put it more cautiously, the rules may be so complex and context sensitive that we can only ever aspire to approximate rules.

It is possible and plausible that humans could have evolved skills that they can now deploy in a trial and error process, to learn, say, how to ride a bicycle. In this case, there is no account in propositional terms of bicycle riding, and it can neither be the same as, nor can it be a subset of, knowing-that. Stanley and Williamson seem to be aware of this, and they make a proposal to get over the difficulty that this entails for their position, that bicycle riding is knowing-that. They introduce the notion of "practical guises of propositions"[51] which they acknowledge is obscure. It seems to involve the notion that one could come to know-how by imitating someone who knows-how. This idea is certainly outside the ordinary conception and cannot be right, since while this might occasionally be true, it cannot be a general truth[52]. It also incorporates the objectionable notion I referred to above, that one can know-how without being able, despite the absence of any physical impediment.

---

[51] J. Stanley and T. Williamson, *ibid*, p 429.

[52] This is not to deny that one can profit from the examples of others in learning to know-how to do something. Indeed it is a characteristic of learning in this domain that one is commonly shown how rather than told how. Examples are being shown a video of one's own performance by the coach, while she draws attention to its shortcomings ostensively, or the piano teacher playing the piece herself for your instruction.

CHAPTER 3. KNOWING-HOW AS A SUBSET OF ABILITY

3.1 Knowledge of Other Minds

There is a field of human activity that I have as yet said nothing about that is viewed by some philosophers as a prime example of knowing-how. This is the area that is part of what is often referred to as "folk psychology" and is to do with how we interpret and often predict the behaviour of other humans. It sometimes also includes, in the opinion of some, the explanation and prediction of the behaviour of certain non-human animals.

As we shall see, consideration of folk psychology forces a reconsideration of the knowing-how versus ability dichotomy that I have maintained up to this point. Indeed, I conclude that the dichotomy is false. Knowing-how is a subset of ability and it is a proper subset; not all ability is knowing-how, just as not all cruelty is torture, although all torture is cruelty.

The behaviour of humans is extremely complex and poorly understood and often the behaviour of others seems inexplicable, even that of close friends. On the other hand, there is a vast area of predictable behaviour. We attend a lecture theatre at the appointed time and the lecturer appears, as we expect, and talks for 50 minutes and he is as boring but well informed as we expected, as usual he makes no eye contact with the class, he rubs out everything he writes on the blackboard promptly after he has written it and he paces back and forth constantly. Without much effort, we deduce certain things from this behaviour. He turns up because he does not think to do otherwise as that is part of his job, which he accepts, so we conclude that he is ordinarily dutiful; he is boring because he is not very extrovert or original or

maybe he does not care to be interesting; he is well informed because he is interested in philosophy, or was once. We guess maybe he is insecure with this eager erasing of what he has written and his failure to make eye contact and his frenzied pacing about.

This analysis of the lecturer's behaviour may be wholly or partly wrong but the point is that we analyse in this way and we are often at least approximately right. Indeed if we were not, we could not conduct ordinary affairs in society.

How is this achieved? The ordinary deductions that are the common stuff of our daily lives are part of what is known as folk psychology, and one view of this is that it constitutes a theory of human behaviour that we share. Clearly, some have a more developed version than others and more often make analyses and predictions that are subsequently shown to be correct. Nevertheless, we all have the ability to make these generalizations to some degree. This view is commonly dubbed the Theory Theory. To understand it we need to look at the claim that we are indeed dealing with a theory in the usual sense of the word. In a well-established discipline like physics, to take one example, we explain the fact that all freely falling bodies in a vacuum fall with constant acceleration, by asserting there is a law of gravitation. That is, we assert that there is regularity in nature that we have discovered and we can make generalisations about all falling bodies by extrapolating from our finite number of observations.

The problem with the Theory Theory is that there is no evidence that there are in fact such generalisations about human behaviour. Indeed, there seems little in common between folk psychology and even a fairly recent body of theory

like that in, say, Biochemistry. There is a counter argument that says there are indeed generalizations about human behaviour, and those who see a chaotic profusion of behaviours are failing to see the wood for the trees. We all fear danger and avoid pain; there are many such global generalizations. Theory Theory advocates commonly suggest[53] that we are born with a theory-of-mind module and while there are wide cultural variations in the language of the Theory Theory, the innate language independent theory-of-mind module (largely inaccessible to consciousness) grinds out our intuitions about the behaviour of our conspecifics - *homo sapiens sapiens*.

An alternative to Theory Theory is Simulation. Here it is argued that we do not have a known body of theory[54] but instead, being human ourselves, we imaginatively and empathetically put ourselves in the position of the person whose behaviour we are trying to understand. A common analogy for the contrast between these two approaches is to think of designing, say, the hull of a boat. Analogous to the Theory Theory approach is to begin from the basis of known fluid dynamics. The analogy with simulation is to imagine designing by testing a series of models in a water tank. We might try to model a friend's behaviour via our own imagined behaviour (since we too are human) in the same circumstances as our friend, making an allowance for our known

---

[53] See for example G. Botterill and P. Carruthers, "The Philosophy of Psychology", Cambridge: Cambridge University Press, 1999, Chapter 4.

[54] It does seem that a pure simulation theory is hard to maintain, and many who accept a simulation account actually allow of a mixed account. In particular, a pure account faces some difficulty with the self-ascription of mental states.

differences in temperament[55]. It can be argued that here also one is assuming underlying regularities in nature, which explain the utility of the models in the tank, and are in turn a prerequisite for the simulation model of analysis of other minds. However, defenders of the simulation approach argue that it is not appealing to regularities in human behaviour analogous to events in the world of things. It is claimed that a human agent and only a human agent can appreciate the force of reasons, and that there is no parallel in the behaviour of falling bodies[56].

Whatever the best account of the capacities involved in folk psychological explanations, our capacity to predict and explain the behaviour of others, or to deduce the emotional state of others, raises some basic questions about knowing-how. When we see, or better, sense, that an acquaintance is likely to behave in a certain way or is in pain or is anxious, or afraid, or excited, or in any one of the many possible emotional states, we have acquired a belief about them. As I have described, there are crudely two theories of how we do this. One is that we mentally put ourselves in their situation and see how we ourselves feel, i.e., we attempt a simulation. The other view is that we observe

---

[55] I am not suggesting, either in the Theory Theory or the Simulation case, that unconscious irrational elements may not play a role in the analysis. We may well know, of a friend, that she had a violent authoritarian father, and that this plays a part, little understood by her, in her responses to certain situations.

[56] One difficulty with the simulation idea is that there is good experimental evidence that we are often wrong in our predictions of how people will behave in well-defined experimental situations. If 95% of us are wrong in our predictions of about how 95% of people will behave in a particular situation, this is no difficulty for Theory Theory, we just do not have a very good theory in relation to these conditions. But it is a problem for the simulation theory because on this hypothesis we ought to have been able to simulate how we would behave, i.e. the putative analogy with the models in the water tank breaks down. There is some evidence that a high percentage of people wrongly predict what a substantial majority will do in some well defined circumstances (S. Stich and S. Nichols, "Folk Psychology: Simulation or Tacit Theory", in "Folk Psychology", Ed M. Davies and T. Stone, Oxford: Blackwell, 1995.)

their behaviour closely and apply a theory of the mind's working. This ability, however it is explained, is universal among humans - no esoteric skill is learned. There are indeed great variations in the ability to make sound judgements about the future behaviour or the emotional state of others and it is also true that the skills develop as we mature. Experiments have shown that four year olds are able to attribute false beliefs to others but a four year old is often deceived as to the emotional states of others that might be obvious to an adult.

Is it merely ability (as I called it in the preceding paragraph), to understand other minds or is it knowing-how? Is it like walking or like talking? It meets all the criteria for knowing-how to φ, *viz.*:

1. A can entertain φing for a purpose.
2. A is acquainted with a set of practical procedures for successful φing.
3. A exhibits or has exhibited reasonable success at φing.
4. A complete propositional account of φing is not available to A.
5. A must have learned to φ via a process of instruction which may include self-instruction (and may be an unconscious process).

On the other hand, walking, which I have characterized as ability, meets all of these except the last, and it could well be argued that one does learn to walk, even if, unlike talking, one does not need to be taught.

I did describe a feature that distinguishes ability from knowing-how, namely that one can lose the former while retaining the latter. It seems that one could

not lose the ability to make judgements about other's minds while retaining the knowing-how. The cases considered previously were of two types: first, loss through physical disability, like the arthritic pianist; second, the loss, as a result of brain damage, of the capacity to speak a language or do mental arithmetic.

The ability to read other minds is a significant challenge to the position largely adopted so far. There is not in this domain some local physical feature (arthritis like) or local mental feature (particular part of the central nervous system) that could be knocked out, leaving us knowing but unable. Even if it were firmly established that we had some "theory of mind" module that enabled our folk psychological performance and it was defective (as is often proposed as an account of autism), there is no analogy with the crippled pianist. In the case of autism it is the ability that is missing, not its manifestation.

There is another feature of knowing-how that I included in the description but not in the definition; namely, the fact that one cannot determine a ranking of knowing-how between individuals, whereas one can do this with regard to ability. In the "knowledge of other minds" case we again find it falling on the ability side of the divide, rather than the knowing-how side (if that is how it is appropriate to make the division). It is a commonplace that there are individuals who are more "in tune" with the moods and feelings of others and who are more often able to reliably predict the behaviour of others. Indeed, it is well within the bounds of contemporary psychological research to devise an

effective approximate measure of this skill[57]. I assume here that it would be agreed between those who regard knowing other minds as simulation and those who regard it as a theory of mind plus observation, that there is an element of skill in the process, albeit this element is perhaps less present in the latter view[58]. It is a skill to manage to work successfully with any complex theory. The knowing-that of quantum mechanics is not the same as the knowing-how of solving new problems in quantum mechanics.

It is counter intuitive to demote the sophisticated skill of reading other minds to mere ability while categorizing bicycle riding as knowing–how. This stems from the intuition that knowing-how, as I have defined it (and as it is generally perceived), is more dependent on rationality in some form than many other abilities. Opening one's mouth hardly seems rational at all.

It is claimed that severely autistic people have little or no ability to read other minds, but it makes little sense to suggest they know-how but cannot.

One line of argument is to say, "Well so much the worse for our intuitions of a hierarchy of cognitive capacities". There is certainly ample evidence in other domains that intuition is a poor guide. It is difficult to conjure up anything more

---

[57] See, for example, the Reading the Mind in the Eyes Test, S. Baron-Cohen, T. Joliffe, C. Mortimore and M. Robertson, "Another advanced test of theory of mind: evidence from very high functioning adults with autism or Asberger Syndrome", Journal of Child Psychology and Psychiatry, 38, 813 – 822 (1997). Typically such tests ask the subject to choose between different emotions, say, happy, anxious, fearful and sad, which are expressed in a series of different photographs showing only the eyes. The tests are controlled by separate tests ensuring that the subject is competent at face recognition and, say, gender recognition.

[58] In early versions of Theory Theory it was taken to be a body of doctrine, known to the layperson, and hence one might argue its application would be knowing-that. Later versions accept that the knowledge is tacit to a significant extent.

counter-intuitive than a particle without mass or a function that is everywhere continuous and nowhere differentiable. However, the problem is more deep-seated. We seem not to have a satisfactory way of deciding whether "reading other minds" is ability or knowing-how. Walking may properly be regarded as ability since it meets the criteria for ability and is akin to other abilities. Reading other minds seems to be in the heartland of the examples of knowing-how that require mental activity; in addition it clearly meets our earlier criteria of ability.

Thus the dilemma of deciding whether knowing-how is ability or knowing-how is distinguishable from ability disappears. It is both. Knowing-how is a subset of ability.

3.2 Summary

In Chapter 2, I distinguished tropistic responses from ability, and both of these from propositional knowledge (knowing-that), taking the last, in the interests of simplicity, to involve justified true belief. In this chapter, in considering folk psychology I have argued that our capacities in folk psychology constitute ability and that they constitute knowing-how, in terms of my explicit definition.

This conclusion, from the case of folk psychology, seems readily to generalize to the more familiar examples of knowing-how, like bicycle riding. That is, all knowing-how is ability (not all ability is knowing-how). The fact is that there is neither adequate knowledge of, nor a meaningful theory of, general intelligence. It is easy to understand a pianist with arthritis and not too difficult to comprehend the effects of, say, aphasia. The genuine distinctions between ability which is not knowing-how, and knowing-how (which is also ability), in

terms of the clear possibility in many cases of retaining the latter while losing the former or being able to rank the former and often not the latter, is essentially trivial.

Knowing-how has become part of the *modus operandi* of our minds. The border between simple ability, and ability which is knowing-how, is not clear-cut. The important point is that knowing-how is a subset of ability. At one extreme we have capacities, which I have declared do not even rank as ability, such as maintaining one's blood pressure. Then there are abilities like being able to walk. There are more complex learned abilities like cycling and skilful chess playing, which are knowing-how, and even more sophisticated learned abilities such as being able to know about other minds[59], which is also knowing-how. All of these are part of the way the mind operates. The fact that one can "lose" an ability, as in the case of the arthritic pianist, simply illustrates that in this case the knowing-how cannot be manifested, and the distinction is not a real difference between knowing-how and ability, but between having an ability and being able to manifest that ability[60]. (Of course,

---

[59] Assessing the degree of sophistication of chess playing versus making deductions or predictions about other people's behaviour is probably a futile exercise. Many would claim that the latter skill is at least partly innate, so one is not comparing like with like.

[60] Thus, I would claim that knowing-how to speak and understand a language is ability. Chomsky is clearly the most distinguished champion of the view that this is not so. He says (N. Chomsky, "Rules and Representations", Oxford: Blackwell, 1980, p 51), "Imagine a person who knows English and suffers cerebral damage that does not affect the language centres at all but prevents their use in speech, comprehension, or let us suppose, even in thought. Suppose that the effects of the injury recede and with no further experience or exposure the person recovers the original capacity to use the language. In the intervening period, he had no capacity to speak or understand English, even in thought, though the mental (ultimately physical) structures that underlie that capacity were undamaged. Did the person know English during the intervening period?" Chomsky's answer is clearly "yes" and hence he says this knowing is not ability. As I have already indicated, I think this is strictly true but in a trivial sense and this knowing-how is best seen as ability. Chomsky's indignant response to this line of argument is that this is "inventing a new term "ability" that is used in

if ability cannot be fully manifested, as in the case of the ageing tennis player's knowing-how, it cannot readily be measured.)

It might be argued that I have merely stipulated that the failure to be able to manifest ability (through amnesia or injury, say) does not disqualify knowing-how from being properly characterized as ability. There are two points to be made. First, it is emphatically not being claimed that ability and knowing-how are synonymous; the claim is that the latter is a proper subset of the former, as euthanasia and murder are proper subsets of killing. Second, we are interested in whether knowing-how is distinct from knowing-that and is akin to ability.

On the present view there is a marked contrast between knowing-how and knowing-that: the latter consists of the manifold simple and sophisticated examples of the mind's capacity to store and manipulate information in dealing with propositional attitudes.

---

the sense of "knowledge" and is radically different from "ability" in its normal sense" (N. Chomsky, "Language and problems of knowledge", Chomsky Symposium, Madrid, 1986, cited in ""A Review of D'Agostino, "Chomsky's System of Ideas", "Noam Chomsky Critical Assessments", Ed C.P. Otero, London: Routledge and Kegan Paul, 1994, p 480.) However, the fact is that Chomsky is citing a case where ability cannot be manifested. It does not alter the fact that aluminium oxidizes rapidly when exposed to air, that oxidation proceeds extremely slowly because of the formation of an oxide layer, which is highly impermeable to oxygen. There are conditions for a disposition to be manifested, in this case a clean aluminium surface.

In the same discussion as that cited above from "Rules and Representations", Chomsky raises the issue of the person who is in a dreamless sleep, and asks rhetorically does he know his name. Thus, we might agree that someone in a dreamless sleep does know his name (propositional knowledge), does know-how to ski (knowing-how) and can sign his name (ability of a general kind and not knowing-how), so the possibility of manifesting any of these at a particular instant is not constitutive of any of them. (The large subject of what it is to know a language and its relation to knowing-how is taken up in Chapter 5)

Finally, knowing–how is, I think, a major part of our mind's functioning, at least in comparison with knowing–that. How significant either component is, is a major current empirical question. Chomsky remarked in 1975[61],

> "conscious beliefs will form a scattered and probably uninteresting subpart of [our] cognitive structure"

The experimental evidence gathered since seems to confirm this view[62].

Of course, what is interesting depends on one's perspective, and to a moralist the subtleties of, say, our auditory system, are uninteresting. Aristotle suggested that we become virtuous by practicing virtue, just as we become skilled at playing the lyre by practice. Is behaving morally knowing-how? This complex and questionable form of knowing-how is discussed in Chapter 7.

---

[61] N. Chomsky, "Reflections on Language", New York: Pantheon, 1975, p 163.

[62] See for example, S. Dehaene, "The Cognitive Neuroscience of Consciousness", Amsterdam: Elsevier, 2001.

CHAPTER 4. KNOWING-HOW AS KNOWLEDGE

"Nature, by an absolute and uncontrollable necessity has determined us to judge as well as to breathe and feel"

David Hume[63]

4.1 Propositional Attitudes and Subdoxastic States

A necessary basis for the claim that the concept of knowing-how is a form of knowledge, distinct from knowing-that, must be to establish that some form of knowledge other than propositional knowledge is possible.

In turn, to establish this last, we must characterize propositional knowledge by some specific feature. This will allow us to put forward a candidate for knowledge that lacks this feature but is nevertheless knowledge.

To anticipate the argument, my claim is as follows.

(i) Inferential integration is the specific hallmark of propositional knowledge.

(ii) There are states - subdoxastic states - that are not inferentially integrated.

(iii) Subdoxastic states are integral to knowing-how, and this is knowledge because it involves sequences of information processing via subdoxastic states[64].

---

[63] D. Hume, "A Treatise of Human Nature", 2nd edn, L.A. Selby-Bigge and P.H. Nidditch (eds), P.H. Nidditch, (rev), 1978, I, iv, i, 183.

Propositional knowing consists of a particular attitude towards a content. The proposition might be "π is a transcendental number", and the attitude knowing (since the proposition is true). I have already remarked that there is a major philosophical debate around a knowledge claim of this sort[65], but it serves for my purpose of discussing subdoxastic states.

I follow the account of subdoxastic states given in the classic 1978 paper by Stich[66]. In this paper Stich establishes a distinction between beliefs and subdoxastic states - certain other psychological states that

> "play a role in the proximate causal history of beliefs, though they are not beliefs themselves.[67]"

Thus, for Stich, both beliefs and subdoxastic states have semantic content and figure in causal explanations.

Stich is concerned to show that the subdoxastic states can be distinguished in a meaningful and important way from beliefs. He argues that while they are not beliefs, these states are instrumental in much of our belief formation. Of

---

[64] Thus for me, as for Ryle, knowing-how is manifested at the personal level, being a good marksman, for example. However, it is characterized at the sub-personal level, and this was not an issue for Ryle.

[65] One complication is that, under other circumstances I might not know the fact about π, but I might know that a particular one of the various infinite series that converge to π gives a transcendental number, but not know that this series converges to π. This raises the important issue of sense and reference; the issue of the semantic relevance of the mode of presentation that arises in Chapter 8, where I discuss the paper by Stanley and Williamson, already cited.

[66] S.P. Stich, *ibid*

[67] S.P. Stich, *ibid*, p 499

course, not all belief formation relies directly on subdoxastic states. Many of our beliefs are simply inferred from other beliefs.

Consider the information we obtain via visual perception. There is now a huge body of evidence on the complexity of the processes that operate as we look at a scene. For example, an increase in the magnitude of the retinal image is generally associated with approach of the object; however, if this is brought about by means of lenses the object still seems to approach. This seeming approach is the result of an unconscious process. That is, an object is taken as being of a known size by a normal observer.

To take another example, we are able to focus both eyes on a single object and hence are capable of stereoscopic vision, which is a great aid to depth perception. The principle of stereoscopic vision can be described in terms of the vision process involved in the use of a stereoscope, which presents an image from two slightly different angles, so that the eyes can merge them into a single image in three dimensions[68]. We can, however, make judgements of depth with one eye closed, or about objects in a painting, which, of course, is two-dimensional. Moreover, we receive an impression of depth from two random dot patterns, one of which is slightly displaced relative to the other. Some dots simply appear to be in front of others, although we cannot say why or how we have obtained the belief that this is so. As Stich puts it[69],

---

[68] This is a somewhat crude account, the merging is done via the visual system as a whole, and the "single image" is more of a metaphor than the word "image" would suggest.

[69] S.P. Stich, *ibid*, p 503

> "underlying the belief there is, no doubt, a complex psychological mechanism which serves to measure the degree of binocular disparity and to use this information in the production of an appropriate belief about apparent relative depth. If, as seems inevitable, the process leading from retinal stimulation to belief involves various psychological states which represent features of the retinal images, then the states are subdoxastic. Though they are causal antecedents of belief, there is a strong intuitive inclination to insist that they are not beliefs themselves."

In Chapter 5, I discuss some aspects of the visual system in detail. For the present, it is sufficient to notice that our brain/mind analyses incoming information in ways that are outside our conceptual capacities. The contents of such states are formulated in terms of concepts that are clearly not possessed by the average perceiver. One prime example is our use of a 2nd order difference operator (gradient of the gradient) to detect edges or gradual peaks and troughs in the gradient of intensity in the visual field. It is the states underlying such capacities that are the subdoxastic states.

Stich gives another, slightly different example, where male subjects found more sexually attractive, photographs of women in which the pupil in the photograph had been enlarged by the experimenter, compared with an otherwise identical set of photographs. The men formed the belief that the women with the enlarged pupils were more attractive via psychological states which served to record the ("unknown") information about pupil size. These states are again subdoxastic but contribute to belief formation. The subdoxastic states allow inference to belief but in a very restricted fashion. The subdoxastic states related to the pupil size would not for example lead to the belief that the pairs of photographs of identical women (apart from pupil size) presented at different times, had different pupil sizes. This is despite the

fact that this is clearly a belief closely related to the belief which was generated – namely that the women with larger pupils were more attractive.

Of course, it could be argued in this case that the men had an unconscious belief that the relevant women had larger pupil sizes. There is no evidence for this from the experiments, and it is not plausible that a similar objection along these lines can be raised against subdoxastic states in general.

These and other examples provide overwhelming evidence that our mind/brain processes information in ways that cannot themselves be described as belief formation but which lead to belief formation. The evidence is detailed and the result of extensive experimental studies using a variety of techniques. Some of this work is described further in Chapter 5. For the present, it is sufficient to have established the reality of subdoxastic states.

It can be seen from this account that the mental state of having an item of propositional knowledge, involving an attitude towards a content, is a different entity from a subdoxastic state.

Evans has characterized propositional knowledge by means of the so-called Generality Constraint, which relates closely to the property of inferential integration mentioned in (i) above. I need now to discuss this, to pursue my argument that we can characterize propositional knowledge and then distinguish (via the subdoxastic states) a form of knowledge that is other than propositional knowledge. My claim, of course, is that this is knowing-how.

4.2 The Generality Constraint[70]

To avoid the complexity of indexical expressions and of conversational implicature, Evans[71] considers demonstrative thoughts of the form Fa, which consist of ascribing the property F to an object a. Evans says that to do this we require a concept of an object and a concept of a and also a concept of a property and a concept of F. One needs to appreciate that there are objects and they can be distinguished from one another, i.e. they can be individuated. Objects, he says, can be individuated by their spatio-temporal location. We must be able to appreciate that there are objects b, c, d etcetera, about which it may be possible to say truly or falsely, Fb, Fc, Fd etcetera. Further, if we are able to have the demonstrative thought Fa, we must be able to have the thought of other predicates applying to a, i.e. we must be able to entertain the thoughts, true or false, Ga, Ha and so on.

This, the so-called Generality Constraint, can be expressed as follows. If one can grasp the meaning/entertain the belief that x is F and one can grasp the meaning/entertain the belief that y is G, then one must be able to understand what is meant by x is G and y is F.

One might believe that a platypus is a monotreme and that a donkey is a mammal. One can then also understand the belief that a donkey is a

---

[70] What follows leans heavily on two papers by Martin Davies, "Tacit knowledge and the structure of thought and language", in "Meaning and Interpretation" Ed C. Travis, Oxford: Blackwell, 1986, Chapter 5, and "Tacit Knowledge and Subdoxastic States", in "Reflections on Chomsky", Ed A. George, Oxford: Blackwell, Chapter 7.

[71] G. Evans, "The Varieties of Reference", Ed J. McDowell, Oxford: Oxford University Press, 1982.

monotreme and a platypus is a mammal. It happens that the former beliefs are true and the latter respectively false and true, but the holder of the former correct beliefs could in general understand the latter (and presumably in general be aware of their truth status[72]).

The important point for my purposes is that the Generality Constraint is constitutive of propositional attitudes. An illustration will clarify why this is the case. A parrot might say, "Fido is a mammal" and moreover might say this only and always when it sees Fido, the dog. One would not as a result say that the parrot had the concept mammal. On the other hand, if a parade of creatures were presented successively to the parrot and it said of the cat, the giraffe, the bat and the whale that they were mammals and denied that the shark, the crocodile and the frog were mammals, we would begin to think this parrot had the concept, especially if the parrot added that the whale suckled its young while the shark did not.

Evans's claim then is that the propositional attitude involving the proposition Fa is making a judgement. Fa is either true or false, and what is essential about a propositional attitude is that a judgement is being made. A thermometer reading may be true or false, but a thermometer is not making a judgement. Content in a person's demonstrative thoughts, or propositional

---

[72] I say "in general" because, for example, someone might believe that monotreme and mammal were exclusive categories, and not be aware that the monotremes are a subset of the mammals. They would, nevertheless, at least be aware that monotremes were egg laying if they could genuinely entertain the belief that a platypus was a monotreme. Of course, someone might simply have been told that a platypus was a monotreme and know no more than this. It seems doubtful whether such a person has a grasp of the meaning. Putnam claims to understand the terms elm and beech, despite being unable to distinguish an elm from a beech. This seems to me to be a dubious claim. Despite qualms of this sort I believe the argument that follows goes through.

attitude statements expressing them, is normative[73]. There are normative patterns, and to have a propositional attitude is to understand a concept and hence to be able to conform to a pattern of use. Wittgenstein illustrates this with the example of continuing a simple series[74]. To write down 14 as the next number in the series 2 4 6 8 10 12 is mandated by the pattern of living in which we have become acculturated, with the way we have lived.

> "It is not possible to obey a rule "privately": otherwise thinking one was obeying a rule would be the same thing as obeying it[75]."

The difference between a human being saying "Fido is a mammal" and the case of the parrot is that the person is making a judgement. The judgement may be false; what our sentient human friend Gloria took to be Fido may be a stuffed replica of Fido and hence not a mammal at all. Nevertheless, it is constitutive of a propositional attitude that it conforms to the Generality Constraint. In having the propositional attitude to Fa one is making a judgement. To say, "snow is white" is true iff snow is white is non-trivial. For the parrot to say "snow is white" commits it to no opinion, not merely on the question of the truth or falsity of the utterance, but to nothing at all about the very possibility of truth or falsity.

---

[73] There are complex debates around this approach, which I simply accept here. See for example W.S. Sellars, "Empiricism and the Philosophy of Mind", in "Science Perception and Reality", London: Routledge and Kegan Paul, 1963 and S. Kripke, "Wittgenstein on Rules and Private Language", Oxford: Blackwell, 1982.

[74] This conformity to a pattern of use has been discussed in Appendix B, where Wittgenstein's rule following is considered.

[75] L. Wittgenstein, "Philosophical Investigations", 202

We can look at this characterization of propositional attitudes via the concept, elaborated by Stich[76], of inferential integration of beliefs. This analysis is closely related to that given by Evans. From ordinary beliefs, whether true or false, we can generate many further beliefs. If I believe it is raining (to use a favourite philosophers' example) and I believe I do not want to get wet, I infer that I would be wise to take my umbrella if I am going outside and so on. This phenomenon, of the elaborate interconnected network of beliefs via inference, is referred to by Stich as inferential integration of beliefs. We do not infer all the beliefs that are logically consistent with our corpus of beliefs. Moreover, we do sometimes wrongly infer false beliefs from true beliefs, and much less frequently vice versa, but nevertheless there is a network of beliefs capable of indefinite expansion in every direction. As Stich puts it, beliefs are inferentially promiscuous. In addition, we normally endeavour to adjust our beliefs on acquiring a new belief in order to maintain consistency among our beliefs.

Superficially, it seems that Stich's inferential integration is entailed by Evans's Generality Constraint and thus adds nothing new. However, it does add an additional useful perspective to the picture of the systematic semantic analysis of propositional attitudes.

A formal proof that the Generality Constraint entails inferential integration looks to be a daunting task because of the difficulty of giving a formal specification of the domain of demonstrative thoughts within a natural

---

[76] S.P. Stich, "Beliefs and Subdoxastic States", Philosophy of Science, 45, 499 – 518 (1978)

language. However, it can be seen that it is plausible. One might believe that chalk is white and snow is white and snow is cold. One might never have thought about the proposition that chalk is cold. Given the Generality Constraint, this must be a proposition one could grasp. Consequently, it can readily be added to one's stock of beliefs either that chalk is cold or that chalk is not cold or that chalk can be either cold or not cold, since all are propositions that can be grasped. This grasping of new beliefs is true from pairs or indeed various n-tuples of all of our beliefs. It is this possibility of adding new beliefs from existing beliefs, and thus the necessary interdependence of beliefs, that entails inferential integration.

Now if we accept the Evans account involving the Generality Constraint, or alternatively, Stich's account via inferential integration, we have a clear distinction between belief or knowledge as mental states, and the subdoxastic states identified by Stich.

As Evans[77] puts it,

> "It is one of the fundamental differences between human thought and the information-processing that takes place in our brains that the Generality Constraint applies to the former but not the latter."

What then is the relation of subdoxastic states to knowing-how? Riding a bicycle does not involve belief in any proposition. So, since it is not justified true belief, what makes knowing-how knowledge? This is taken up in the next section.

---

[77] G. Evans, *ibid*, p103, fn. 22

4.3 What Makes Knowing-how Knowledge?

Pre-lingual infants are aware of the continued existence of objects that are out of sight and can tell the difference between 2 objects and 3 objects and between 2 or 3 and larger numbers, although not between say 5 and larger numbers. These are beliefs, and they are not justified in the sense that my belief that $\sqrt{2}$ is irrational is justified. Similarly, when a small child utters a sentence, it has the belief that it has indeed done just that, and not merely emitted noises. Knowledge, as Ryle[78] put it, is "not an elite suburb of belief". These feats just described are performed on the basis of subdoxastic states. The feats are instances of knowing-how: information is being processed. The same is true of cycling.

The capacities of pre-lingual infants that I have just described are innate. How then are they justified? We could perhaps accept direct justification in a case like knowing that a rose we see is red. But a belief that "all natural languages have a grammar which has a phrase-structure part" cannot claim such a direct justification. Such knowledge, which is tacit knowledge in the case of the ordinary unsophisticated language user, is not justified true belief. I use justified true belief here in the sense of epistemically responsible belief[79]. But for the knowing-how of language and aspects of the knowledge of pre-lingual infants, there is nevertheless a case to be made that we are in fact dealing with knowledge.

---

[78] G. Ryle, "On Thinking", Oxford: Blackwell, 1979, p 95.

[79] I distinguish this from reliably grounded belief, which I return to below. Some writers refer to this last as justified true belief and others seem to leave this distinction aside.

Justified true belief in the above sense is a useful concept. Nevertheless, *inter alia*, the complexity of the brain suggests that it does not exhaust the concept of knowledge. The argument that it does so rests on presuppositions about mental functioning that are questionable. For example, Gallistel doubts that there is anything that can usefully be called general intelligence, something, that those who argue in favour of knowledge as exhaustively described as justified true belief, implicitly attribute to a person. He says[80]

> "If problem-specific learning mechanisms are required to explain everything from the learning of the solar ephemeris in bees to song learning in birds and language learning in humans, should we none the less continue to imagine that there exists a general-purpose learning mechanism in addition to all these problem-specific learning mechanisms? What structure could it have? It is like trying to imagine the structure of a general-purpose organ, the organ that takes care of the problems not taken care of by adaptively specialised organs like the liver, the kidney, the heart, and the lungs."

As he remarks earlier in the same article,

> "Adaptive specialisation of mechanism is so ubiquitous and so obvious in biology, at every level of analysis, and for every kind of function, that no one thinks it necessary to call attention to it as a general principle about biological mechanisms.[81]"

Of immediate relevance to the claim that knowing-how is knowledge, is the problem of giving a satisfactory theory of knowledge in animals other than humans. Kornblith[82] describes the complex cooperative behaviour of ravens in

---

[80] C.R. Gallistel, "The Replacement of General Purpose Learning Modules with Adaptively Specialized Learning Modules" in "The New Cognitive Neurosciences", M.S Gazzaniga (Ed), 2nd Edition, Cambridge: MIT Press, 2000, p 1183

[81] C.R. Gallistel, *ibid*, p 1179)

[82] H. Kornblith, "Knowledge in Humans and Other Animals" in "Philosophical Perspectives 13 Epistemology", Ed J. E. Tomberlin, Oxford: Blackwell, 1999, pp 327 – 346.

distracting hawks and osprey to steal the prey from these birds. He also describes careful experimental studies of the elaborate behaviour of the piping plover in protecting its young by distracting predators, using an elaborate broken wing display, making a convincing appearance of being injured. Kornblith then says the following.

> "Even if we grant that the sophistication of the plover's behaviour requires the use of intentional idioms, what is it that licenses talk of genuine knowledge here, rather than the more prosaic belief?"

The answer is an evolutionary one, and the claim that we are dealing with knowledge is given in terms of the requirement for a species to behave successfully in an environment, i.e. we have reliably produced true belief as knowledge. On this basis, the instances given above of the performance of pre-lingual infants are knowledge.

There are various objections to the notion that animals other than humans can have knowledge at all, or knowledge that is relevant to philosophical study. I think that animals do have knowledge comparable with at least part of what we call knowledge in humans. Indeed it seems intuitive to think that a chimpanzee in captivity knows that it is normally fed at a certain time each day and we know that these chimpanzees can learn to ride a bicycle. The former is, of course, propositional knowledge, if it is in fact knowledge, and the latter, knowing-how.

There is an argument that the relevant sense of knowledge in philosophical study is only what can be predicated of humans who have reached "the age of reason". That is, the information processing that I have described in pre-lingual infants is not to be counted as knowledge. There is a sense in which a

thermometer knows what the temperature is, or a supermarket door knows when someone is approaching. Some philosophers argue that we must have reasons for beliefs for these to count as knowledge in the relevant philosophical sense. I have argued, following Evans, that it is sufficient if we are making a judgement, provided that the proposition believed is true. One can have propositional knowledge without reasons. For example, I can say that I know it is cold today without having reasons in the ordinary sense. Moreover, if we can know without knowing that we know, it cannot be essential to have reasons. Imagine a painter who is asked to paint an imaginary scene and conscientiously does so, say of a farm building. Her parents see that the scene is precisely one the painter saw as a child[83]. The painter knows, prior to the parents saying so, that the farm building has, say, four front windows, but she can hardly be said to have believed it. In general, in cases of knowing without belief, one does not have reasons.

There is, in the area of knowing-how, no reason to suppose that some animals other than humans have a different kind of knowledge from that of humans: indeed, my argument does imply that if my riding a bicycle is knowing-how, so is the same feat performed by a chimpanzee. As Kornblith says

> "there is a long history of underestimating the cognitive sophistication of animals, and, in particular the extent of continuity between animal and human cognition, which makes it appear unlikely to me that there is a distinction in kind between animal and human cognition, rather than a distinction in degree of sophistication."

---

[83] I have taken this example from C. B. Martin and M. Deutscher, "Remembering", The Philosophical Review, 75, 161 - 196 (1966) p167, who use it to illustrate a different point.

Either in humans alone, or somewhere between amoebae and humans, knowledge manifests itself. It seems acceptable that quite "lowly" creatures have beliefs as part of their repertoire of mental states. Often these beliefs will be true. However, to ask whether an individual plover has a justified true belief seems to be an inappropriate question.

I have instanced unconscious knowledge as an example of knowledge where we have knowledge without reasons. I have also, in defending the constitutive nature of the Generality Constraint in propositional attitudes, claimed that having a propositional attitude involves judgement. However, the converse is not true, namely that the presence of judgement, in the absence of a propositional attitude, requires the Generality Constraint.

Returning to animals, we know that birds learn songs as nestlings[84]. We know this is complex; not all birds learn songs; some are endowed with certain fixed songs. Some learn fixed songs from their parents or older peers. Some learn to innovate and learn to sing unique songs. This last seems to be the case where individuals as well as merely conspecifics are recognized by means of their song. All of this learned singing would be knowing-how by my definition were it not for clause 1, "A can entertain φing as a purpose". We do not think of birds entertaining singing for a purpose. As we humans see it, they do sing for a purpose – to warn of predators, to call to prospective mates and so on. However, we do not envision them as having second order thoughts such as

---

[84] C.R. Gallistel, A.L. Brown, S. Carey, R. Gelman and F.C. Keil, "Lessons From Animal Learning for the Study of Cognitive Development", in "The Epigenesis of Mind", S. Carey and R. Gelman (Eds), Hillsdale, NJ: Lawrence Erlbaum Associates, 1991, Chapter 1.

entertaining the idea of singing for the purpose of warning their family. Our picture is that the bird perceives danger and singing results. In contrast, we see ourselves as being able to think of crying "danger" in cases of danger, but also in the absence of any danger, when we want to deceive others.

The plover, in attempting to deceive a predator in its elaborate pretence and decoy behaviour, is, Kornblith argues, demonstrating knowledge, and in my view knowing-how, because it does meet clause 1 in this instance. An exercise of judgment is called for, and the relevant subdoxastic states result in the judgement that the behaviour is called for. The behaviour continues while the danger persists and ceases when the danger ceases. On the other hand, birds or monkeys often utter warning cries as a response to certain general types of observation. For example, they will utter these cries on seeing a high-flying, circling bird although often there is no danger[85].

The dividing line between ability that is knowing-how and ability that is not, is fuzzy, and so necessarily is that between what is knowledge and what is not. Nevertheless, I believe my definition of knowing-how is a definition of something properly called knowledge, although it is not a definition of something that could be instanced as justified true belief[86].

---

[85] This is an example of the obvious truth that evolution does not always lead to true belief. For survival, it is better to respond to some cases where there is no danger than to fail to respond to real danger. However, Kornblith's position requires only generally reliably formed belief, not certainty or infallibility.

[86] There are those who claim that knowledge is not an appropriate concept for detailed analysis. Marvin Minsky for example says, "though prescientific idea germs like "believe," "know," and "mean" are useful in daily life, they seem technically too coarse to support powerful theories, we need to supplant, rather than to support and explicate them. Real as "self" or "understand" may seem to us today, they are not (like milk and sugar) objective

There is now ample experimental evidence, some more of which I present in Chapter 5, that some human performances can be explained via information-bearing subdoxastic states. Other human performance requires an explanation in terms of belief / desire psychology. I decrease my pupil size via the former and explain why I drink red wine in preference to white via the latter. The claim that has been made in this chapter is that, for example, one normally correctly identifies sentences as ungrammatical, and hence knows they are ungrammatical, via subdoxastic states. This information processing, which enables knowing-how as I have defined it, and cannot be expressed in propositional form, is nevertheless knowledge. It is grounded in our evolutionary history but it is clearly not justified true belief. Indeed it is not propositional knowledge within the scope of any of the current accounts.

Thus knowing-how reliably generates correct beliefs in the case of language. Of course knowing-how to ride a bicycle generates no beliefs: it generates reliable cycling, reliable action. Both are examples of Ryle's knowing-how, and it is fair to say that we have little fundamental understanding of either.

---

things our theories must explain, they are only first steps to better concepts." (M. Minsky, Behavioural and Brain Sciences, 3, 439 – 440 (1980)). The fault with this line of reasoning is that Minsky has no serious alternative (he has some ideas of knowledge as a graded concept), and we must work with the best we have to refine our concepts like knowledge and belief. This is not to argue that the concept of knowledge as we have it at present is sacrosanct, but only that it has to be replaced, if it is to be replaced, with a better concept. The history of the concept of temperature is that it was at one time a confused muddle of what we would now think of as heat and temperature. Indeed heat and temperature are entirely different types of physical property; heat is extensive whereas temperature is intensive. Once the kinetic theory of gases was developed, the notion of mean molecular kinetic energy as temperature fully clarified the concept. Thus, in this case, the terms were indeed explicated rather than simply supplanted.

CHAPTER 5. KNOWING-HOW AND LANGUAGE

We are not simply a Sartrean néant, but an embodied percipient whose shadow falls on what we take to be there....

Henry Pietersma[87]

In the Cartesian framework, the basic stance of mind toward the world is one of representing and thinking about it, with occasional, peripheral, causal interaction via perception and action. It has been known since Bishop Berkeley that this framework has fundamental epistemological problems. It has been a more recent achievement to show that escaping these epistemological problems means reconceiving a human agent as essentially embedded in, and skilfully coping with, a changing world; and that representing and thinking about the world is secondary to and dependent upon such embeddedness.

Tim Van Gelder[88]

5.1 Introduction

In this chapter I apply what has been said thus far to substantiate my claim that speaking one's native tongue is knowing-how. Moreover, the example of language illustrates that knowing-how is knowledge. In section 5.2, I argue this on the basis that it meets my definition and that it involves the operation of subdoxastic states in engendering belief.

In sections 5.3 and 5.4, I turn to the recent extensive work in neuroscience to clarify the central importance of mental functions not under our conscious control, in activities where we have the firm subjective impression that they

---

[87] H. Pietersma, "Merleau-Ponty's Theory of Knowledge" in "Merleau-Ponty: Critical Essays", Ed H. Pietersma, Lanham MD: University Press of America, 1990, p 107

[88] T. Van Gelder, "What Might Cognition Be, If Not Computation", in "Mind and Cognition", W.G. Lycan (Ed), Oxford: Blackwell, 1999, p 188

are indeed consciously controlled. I am not here concerned with the Freudian unconscious. The point of the account of this recent experimental work is to show that a major part of our complex mental activity is not conducted via our propositional knowledge. It can then be seen that language is part of a continuum of forms of knowing-how that draw on our embodied knowledge. It is this embodied knowledge that is the basis of knowing-how, from bicycle riding to language use. It is thus clear both why knowing-how is a subset of ability and why it cannot be propositional knowledge.

In Chapter 2, I developed a definition of knowing-how to φ (cycle, play chess, speak Urdu) and I claimed, as a consequence of this definition, that knowing-how was ability. However, by no means are all abilities knowing-how. In Chapter 4, I argued that knowing-how is a form of knowledge. It is now necessary to consider real examples which instantiate these claims

In addressing language, which is widely claimed to be knowing-how, I emphasize again that a common expression like "I know how to speak English", tells us nothing in either direction. It tells us no more than to say that to know how to open a safe all one needs is to know the combination – a clear case of knowing-that[89].

We know much more than we can tell, as Polanyi puts it, and we know things we do not know we know: both of these are exemplified in our ability to speak grammatically. On the other hand, as I have already said, there are those who

---

[89] Of course, the safe door may be stiff or the safe may be awkwardly placed at the back of a dark cupboard, but the point is clear - that English usage is not a proper guide to the use of Ryle's knowing-how.

say that this is not ability, but unconscious ordinary propositional knowledge. Notable among these last is Noam Chomsky[90].

I attempt to establish, *contra* Chomsky, that using language is a paradigmatic example of knowing-how and finally summarize this view in section 5.5.

I have, here and elsewhere, taken for granted a broadly Chomskyian picture of our language competence. I discuss this briefly in Appendix A.

5.2 Language, Subdoxastic States and Knowing-how

Chomsky refers to our language as our own idiolect, referring not to some (in his view non-existent) abstract entity, say English, but to the way an individual speaks her language. The individual might be the monolingual offspring of monolingual parents in a country where the language she learns is called English. I will refer to this national language as English[91].

Consider our ability to distinguish between grammatical and ungrammatical sentences. While ignorant of formal grammar, most people are able to agree with a high degree of consistency on which of a set of sentences is grammatical. Nevertheless, the same people can give no coherent account of how they have formed the belief that causes them to say a particular sentence is grammatical or not.

---

[90] N. Chomsky, "Knowing Language: Its Nature Origin and Use", New York: Praeger, 1986, p 269

[91] Chomsky is of course right in implying that this includes a vast diversity of ways of speaking, and he has pointed out that many Dutch citizens near the border with Germany can better understand what the local Germans say than what their own more distant compatriots say. As someone remarked, a language is a dialect with an army and a navy (the remark is often attributed to Max Weinreich but there is no reference cited).

To achieve this, one uses a set of states that store information on English grammar. As I have discussed in Chapter 4, Stich postulates that these last states are subdoxastic and are not beliefs. He gives two reasons. First, he says we normally have access to the content of our beliefs and in this case we do not have such access. Second, he says that these states are not inferentially integrated.

Taking first the issue of access: an immediate rejoinder might be that we do not have access to many of our beliefs, as evidenced by the now extensive studies in psychology on the nature of the unconscious. However, Stich argues plausibly that while this is true, it is true under particular circumstances. The cases where we do not have access to genuine beliefs can commonly be seen to be cases where there is some psychological mechanism at work blocking access. There is a sharp contrast between my understandable lack of access to my belief that I would like to marry my mother (given Freud's account of the workings of the unconscious and its relation to various powerful social conventions and deep-seated taboos), and to an imagined general community-wide psychological barrier that prevents us having access to our information-bearing psychological states relating to English grammar.

The argument for the role of subdoxastic states in language use is that we have many information-bearing states that are inaccessible to us and for which there is no reason to suppose a psychological blocking mechanism. Under normal circumstances we have access to the content of our beliefs. We do not have access to the content of our subdoxastic states.

Moreover, and this is vital, these linguistic subdoxastic states are not inferentially integrated or as Evans might put it, they are not subject to the Generality Constraint. As Davies and Stone[92] say,

> "one of the marks of tacit knowledge is that its content does not have to be conceptualized by the subject. Tacit knowledge of principles about theta roles, or case assignment, or c-command does not require the subject who tacitly knows the principles to possess the concepts of theta role, or case, or c-command."

Note that this is a different matter from the lack of access, which is shared by our unconscious states. Stich makes the same point by comparing an explicit belief in a grammatical rule by a linguist, with a rule simply stored in a subdoxastic state. In the latter case one could not be caused to infer, for example, that Chomsky was in error in some portion of his theory, whereas in the former this is exactly what might happen.

It is important for my argument that there is a sharp difference between the unconscious/conscious distinction and the belief/subdoxastic state distinction. In fact, our unconscious states often are inferentially integrated. The mother's refusal to believe that her only son committed a series of crimes for which there is overwhelming evidence is a direct result of her unconscious wish to believe no ill of her beloved offspring. Likewise, the conviction of parents that their child, missing in action and presumed dead, years after the war's end, is alive and well, is a consequence of their unconscious wishes. Many such instances of the effects of unconscious beliefs in the formation of a network of conscious beliefs are documented in the psychological literature, and are

---

[92] M. Davies and T. Stone, "Folk Psychology", Oxford: Blackwell, 1995 p 12

indeed a commonplace of ordinary discourse – "she won't believe it because she doesn't want to know".

It seems clear that to speak Spanish or English or whatever requires extensive manipulation of information and is not achieved via exclusively using propositional knowledge. We simply do not have access to the necessary knowledge in propositional form. Thus, accepting Evans's argument, subdoxastic states are part of our language endowment and hence we speak our native tongue not by virtue simply of propositional knowledge but by knowing-how in precisely the sense Ryle used the term in accord with the definition I have given[93], namely:

1. A can entertain speaking English (or another natural language) for a purpose.

2. A is acquainted with a set of practical procedures for successful speaking English (or another natural language).

3. A exhibits or has exhibited reasonable success at speaking English (or another natural language).

4. A complete propositional account of speaking English (or another natural language) is not available to A.

---

[93] This is not to deny Evans's statement, "Tacit knowledge of the syntactic and semantic rules of the language are not states of the same kinds as the states we identify in our ordinary use of the terms "belief" and "knowledge." (G. Evans, "Semantic Theory and Tacit Knowledge" in "Collected Papers", Oxford: Oxford University Press, 1996, p 338 (paperback edition)) On the contrary, it is precisely because knowing-how involves tacit knowledge that it is distinct from knowing-that and is nevertheless knowledge (as well as ability). This point is taken up further below.

5. A must have learned to speak English (or another natural language) via a process of instruction which may include self-instruction (and may be an unconscious process).

So knowing-how can be seen in this case to be not only a subset of ability but also a subset of knowledge.

As I have stated several times, this claim does not mean that there is not a vast amount of propositional knowledge involved in the ordinary use of English. Neither Chomsky nor anyone else would claim that my saying, say, that there are infinitely many transcendental numbers, involves only knowing-how.

There are two possible challenges to my claim about language as knowing-how – that it is both ability and knowledge, and that, although it involves knowing-that, it cannot be characterized as knowing-that. The first, already mentioned, is made by Chomsky, who appears to argue that knowing-how to use one's idiolect, including the tacit knowledge of grammar, is ordinary propositional knowledge. I have dealt with this in discussing the views of Stich and Evans above[94] and discuss it further in section 5.5 below.

---

[94] Knowles, in a recent paper (J. Knowles, "Knowledge of grammar as a propositional attitude", Philosophical Psychology, 13, 325 – 353 (2000)), has argued, in support of Chomsky, for the position that the ability to speak grammatically is ordinary propositional knowledge. The argument advanced by Knowles goes as follows. If one accepts, as I do, a distinct grammatical competence, this, he claims, completely undermines a levels analysis in the spirit of Marr (D. Marr, "Vision", San Francisco: Freeman, 1982). Therefore one is left only with the position that we have full propositional knowledge in linguistics, since this position does not require the sort of levels analysis given by Marr for visual perception where a component of tacit knowledge is accepted.

The second challenge is to say that what I am calling knowing-how is merely ability. (I am not now considering the universal claim by Stanley and Williamson that knowing-how can be subsumed under knowing-that. I deal with this in Chapter 8.) The present challenge is that speaking a language

---

Marr's theory of vision postulates it as proceeding by the computation of symbolic descriptions of the image. There are three stages of this computational psychological process. The first level states which function is to be computed (a parallel in mathematics might be say to compute ($x^2 - 1$)). The second level states which algorithm is to be used to compute the function. In our example one might compute x^2 and then subtract 1, or one might subtract 1 from x and store this, add 1 to x and store this, and then multiply these two together (($x^2 - 1$) = (x-1)(x+1)). The third level then states how the chosen algorithm is to be realized in hardware (in our mathematical case, say a simple four function calculator).

The cognitivist claim, as Knowles calls it, is that given this distinct unconscious knowledge of grammar, the levels analysis fails. The reason is that one cannot find the necessary role for each level of analysis in the speaker's competence. Finding such a role is a requirement of Marr's account.

In discussions of this sort it is usual to consider Peacocke's modified version of the Marr account (C. Peacocke, "Explanation in computational psychology: language, perception and level 1.5", Mind and Language, 1, 101 –123 (1986)). This seems to be sufficiently general for the arguments for and against the exclusion of knowing-how to be adequately considered. Consider then the argument against knowing-how based on denying the Peacocke account because no role can be found for all of the Peacocke/Marr levels.

Peacocke introduces a level 1.5 between levels 1 and 2 (as the name implies). This level states the information on which the algorithm draws, i.e. at level 1.5 there will be a description that constrains the algorithms that might be used. Knowles proceeds to a very simple argument. Level 1.5 has, he says, adequately explained the competence; there is no role now for a level 2 or level 3. Proposed rules about algorithms, Knowles says, concern performance mechanisms, not grammatical competence. However, this seems to misunderstand Peacocke. The Marr analysis on Peacocke's account applies to the whole language system, just as in Marr's original discussions it applied to the whole visual system. It is as though Knowles was making a claim in the visual case that the optical system of lens plus retina did not require a three level analysis, and therefore Marr's analysis is wrong. It seems intrinsic to the Chomsky position, that is being accepted here, and of course by Knowles, that the linguistic competence model, considered independently, does not require analysis like that provided by Marr. The analysis applies to the whole language system. This point is made in detail by Rattan in a recent paper (G. Rattan, "Tacit knowledge of grammar: a reply to Knowles", Philosophical Psychology, 15, 135 – 154 (2002)).

In any case, the Knowles argument is dependent on a very precise commitment to Marr's model (as modified by Peacocke), a model that is certain to be a simplification. This is not a criticism of the model; successful scientific theories are normally an abstraction from and simplification of reality. This is because aspects of the situation being studied that are tentatively considered to be of little or no importance are put aside in the legitimate interests of simplicity. In the complex and ill understood area of human physiology it seems unwarranted to place such literal faith in the Marr model, in a domain outside its original intended application, and thus to draw a far-reaching philosophical conclusion.

involves ability and propositional knowledge. It might be said, "I can call it knowing-how, but labeling it thus does not make it a form of knowledge distinct from propositional knowledge". This claim would go as follows;

(i) There is clearly propositional knowledge involved in speaking one's idiolect. (This is beyond dispute.)

(ii) The component that does not involve propositional knowledge - the subdoxastic states - can be properly classified as ability and not as knowledge.

It is true that, for example, we alter our pupil diameter according to a number of environmental stimuli, including light intensity. Moreover, this is done via subdoxastic states that cause appropriate bodily responses of which we are unconscious. This is clearly not knowing-how; indeed it is not even ability. It probably occurs, like maintaining our blood pressure, even when we are brain dead. Walking is ability, not knowing-how. What is the difference between walking and speaking Urdu, that entitles me to say that the latter is knowing-how, a subset of ability, whereas the former is merely ability and not knowing-how?

My answer has been that the defining (fuzzy) line between the subset of ability that is knowing-how and the subset that is not, is given by my definition, and that the former subset is akin to propositional knowledge, but nevertheless distinct. The ability say to clench one's fist is very different from the ability to speak Urdu. The difference lies in the complex interaction between the information bearing subdoxastic states in the language case, and

the fact that this results in reliable beliefs. An example of this would be the belief that one utterance is a well-formed sentence of the language and another is not.

5.3 Information Without Awareness

As I have mentioned a number of times, we all continuously "do" things of which we are unaware. I have cited maintaining our blood pressure: something we do when we are declared sufficiently dead for our vital organs to be removed for the benefit of others. We breathe without intending to do so, both when asleep and awake and heavily anaesthetized. These automatic functions are primarily controlled by the Medulla Oblongata – the lowest portion of the brainstem (at the top of the spinal cord).

Automatic functions may nevertheless be complicated. Consider someone who has extensive damage to the primary visual cortex (area V1). If such a person cannot report whether or not they see a light, they are said to be cortically blind. Yet even a cortically blind person's pupil-diameter will change with light variations in a way that a truly blind person's, say with a severed optical nerve, will not. This can be dismissed as a simple reflex, but experiments have shown that pupil diameter in such patients is responsive to far more than light intensity. The determinants of pupil diameter include stimulus motion and spatial frequency[95]. The point of importance for my purposes is that there is a continuum of complex processes going on in the

---

[95] L. Weiskrantz, "Outlooks for blindsight: explicit methodologies for implicit processes", Proc Roy Soc, B239, 247 – 278 (1990)

mind/brain in mammals in their activities ranging from automatic functions, to deliberate acts like choosing to stand up, to telling deliberate lies, to realizing one has just taken the last step in making a significant contribution to relativistic quantum mechanics.

It turns out that contemporary research has shown that some of our deliberate acts are not the result of processes under our conscious control. This bears directly on the claim for knowing-how. Indeed, it is quite arguably conclusive evidence. Some of the modern evidence is in the area of knowing how to shoot, one of Gilbert Ryle's insightful examples of more than half a century ago, cited in section 2.2. The evidence is primarily in the area of vision, which for a variety of theoretical and experimental reasons is our best-understood sensory faculty.

It is necessary to give some detail of the experimental evidence as part of an evaluation of the significance of what Merleau-Ponty called our "motor intentional" understanding. We are crucially embodied minds and no full comprehension of the complete range of intentional states can be had without recognition of this. As Charles Taylor puts it[96],

> "The human subject is an agent, engaged in activity, and engaged in the world. He is an embodied subject."

He goes on to say that this is both a crashing truism and often overlooked.

---

[96] C. Taylor, "Embodied Agency" in "Merleau-Ponty: Critical Essays", Ed H. Pietersma, Lanham MD: University Press of America, 1990, p 1

One reason for the neglect of our embodiment has been the huge success of what since Descartes have been the paradigm sciences - mathematics and the physical sciences - that can and do steadfastly ignore our embodied nature. Contrariwise, the main reason for its contemporary resurrection in science is the huge advances in experimental techniques, developed in these same physical sciences, which allow extraordinarily sophisticated experiments on animal and human subjects. Twentieth century developments like Nuclear Magnetic Resonance Spectroscopy, Positron Emission Tomography, High Speed Photography, Electron Microscopy, Electroencephalography, Magnetoencephalography, and many others have transformed experimental psychology and neurophysiology. Moreover, the addition of the immense current computing power at the back end of many of these devices has enabled or further enhanced the capacity to experimentally probe the finest details of animal, including human, responses to a huge array of inputs.

In the next section I give an account of the experimental evidence on the mammalian visual system that establishes the centrality of our embodied nature in guiding our actions and clarifies how this is distinct from our conscious intentional stances. In the final section of this chapter, I discuss how this bears on knowing-how as I have defined it.

5.4 The Visual Brain in Action

This section heading is the title of a detailed volume in the Oxford Psychology Series[97], which is a comprehensive survey of the relevant literature. The proposal made by the authors on the interpretation of this huge body of evidence is the basis of the argument made here for knowing-how. The authors themselves make no reference to knowing-how or to Ryle, Wittgenstein or Merleau-Ponty; their interest is in the experimental data on mammalian vision and their implications for a theoretical account of the visual[98] faculty in humans.

Our sight is our main tool in interacting with the world. The blind face two problems resulting immediately from their disability. The first is the difficulty in navigating the physical environment and it is immediately significant that a solution for this is to use a sighted animal, a guide dog. The second is the difficulty they have in representing the world, in the lack of a major part of the human capacity to enable the representation of the world, to identify and re-identify faces and places and objects and to store these in memory. This is mitigated by the acute use of the other senses and the human capacity to learn of the sighted world through verbal and written descriptions. From an evolutionary point of view it seems plausible that the first of these is primary. For the hunter-gatherer the ability to act on the world – to run swiftly, to pick

---

[97] A.D. Milner and M.A. Goodale, "The Visual Brain in Action", Oxford: Oxford University Press, 1995

[98] While I concentrate exclusively on the visual system because it is the most extensively studied and best understood, I do not wish to imply that the surprising data obtained are not also to be found in other sensory functions like hearing.

things up or to dig them up, indeed all the commonplace motor actions enabled by the visual system, must have been an essential for survival.

Milner and Goodale argue that there is a fundamental distinction between conscious seeing and what Merleau-Ponty calls motor intentional action or what is now often called visuomotor action. That is, we wrongly conflate the two functions described in the preceding paragraph. The two are normally conflated because we believe that when we grasp something, we do so via conscious seeing, which we think governs our grasping. We believe when, say, we delicately direct a sharp pencil to make a mark between two other closely spaced marks, the process consists of a guiding of our hand and successive adjustments controlled and monitored by our conscious observation. The experimental evidence, contrary to common sense, is that this is not so. We achieve such feats by a mental mechanism of which we are unaware. The role of conscious vision is only in this case to choose the task. The fact is that our conscious vision instructs an *alter ego* (namely us in our other, unconscious visual mode) to perform the task. The evidence for this is complex and is set out in the remainder of this section.

Primates have a complicated arrangement of visual areas that have been shown to have different but often overlapping functions. These areas occupy about half of the posterior cerebral cortex. The connections between these

different areas are complex, but there are two broad streams originating from the primary visual area within the occipital lobe (V1)[99].

A small digression is necessary here. The information from the retina to the other parts of the brain has already undergone significant processing. The retina is not a mere transducer; it embodies a complex processing network. The electromagnetic radiation striking the retina has been transformed into physiological signals of various kinds, and components diverge to a number of distinct targets. One of these targets is the primary visual cortex (V1) referred to in the previous section, but there is immense complexity because there are interactions between the other locations in the brain, where some information from the retina was sent, and the outputs of the primary visual cortex.

Returning to the two broad streams – there is a ventral stream from V1 projecting to the inferior temporal (IT) cortex and a dorsal stream projecting to the posterior parietal (PP) cortex. (Of course, the IT and PP regions do receive inputs from a number of other visual structures.)

Ungerleider and Mishkin[100] proposed roles for these two streams. The ventral stream, they said, is critical in the identification and recognition of objects and

---

[99] Very, very roughly, the frontal lobes house problem solving, personality, voluntary movement and language; the temporal lobes are concerned with hearing and memory. The occipital lobes handle visual perception and reading ability. The parietal lobes handle touch, pressure, temperature, and pain. These rough allocations are obviously neither exhaustive nor exclusive, and there are complex connections between areas.

[100] L.G. Ungerleider and M. Mishkin, "Two cortical visual systems", in "Analysis of visual behaviour", D.J. Ingle, M.A. Goodale and R.J.W. Mansfield (Eds), Cambridge: MIT Press, 1982; however, see also M. Jeannerod, "The Cognitive Neuroscience of Action", Oxford: Blackwell, 1997 for a more detailed account of the complex history of the development of these ideas.

thus for purposes allied to representation and commitment to memory. The dorsal stream is needed for the localization of these objects (the role of the guide dog in my metaphor). These are sometimes referred to as the "what" and the "where" streams. This categorization fitted the then data primarily obtained from experiments with monkeys. Lesions were made either in the posterior parietal cortex (the "where" stream), which produced defects in performance in spatially demanding tasks, or in the inferior temporal cortex (the ventral, "what" stream) which produced deficits in the monkey's ability to make discriminations based on the visual features of objects.

More recently, experimental advances, including the ability to make neuronal studies from awake rather than anaesthetized monkeys, have produced evidence to change this picture. Milner and Goodale state that the additional evidence of the nineteen-nineties is consistent with the ventral and dorsal streams being associated not with different sub-domains of perception but with the difference between perception (ventral) and the guidance of action (dorsal). Perception here is defined to mean not generally any processing of visual input, but "any process which allows one to assign meaning and significance to external objects and events"[101]. In other words, pupillary contraction, for example, is not then included in perception. This accords with the usual use of perception in philosophical writings[102].

---

[101] A.D. Milner and M.A. Goodale, *ibid*, p 1

[102] In fact, this is now a complex issue and the subject of much philosophical discussion. Noe for example says, "Perceivers are familiar, at least implicitly, with the way appearances change as a result of movement. This familiarity reveals itself in the automaticity with which we squint, pat our pockets in search of our glasses, and move closer when we are trying to

This Milner-Goodale proposal is quite radical. The relation between the individual and the object is not a single system, with separate regions for specialist processing. Instead there are two (more or less) independent visual worlds that often coincide, but can offer different versions of reality.

The distinction between the two systems is immensely important here. The visuomotor system does not transform symbolic inputs into symbolic outputs, where the former could for example be stored in memory. The nervous system, the body and the environment are in a constant real time interactive coupled relationship in performing visuomotor acts.[103] It is plausible to think of the visuomotor system as an older, earlier system in evolutionary terms when viewed in this way. It is a truism in evolutionary theory that earlier systems are retained - incorporated or modified - as development takes place. The idea

---

get a better look at something. To be a perceiver is, among other things, to know how to do all this. In this way our capacity to keep track of the dependence of perception and action is firmly integrated with personal-level conceptual and inferential skills." (A. Noë "Perception, action, and nonconceptual content" a comment on S.L. Hurley, "Consciousness in Action", Cambridge: Harvard University Press, 1998 (http://host.uniroma3.it/progetti/kant/field/hurleysymp_noe.htm). The issues raised by different views of perception are a significant part of the burden of this chapter, as might be evident from the use of the phrase "know how" in the passage just quoted. Of course, what I have just said is part of the wider argument for knowing-how given in this thesis. Heidegger makes the point in the more general context of physical skills as follows, "not a bare perceptual cognition, but rather that kind of concern which manipulates things and puts them to use: and this has its own kind of 'knowledge' " (M. Heidegger, "Being and Time", trans J. Macquarie and E. Robinson, New York: Harper and Rowe, 1962, p 95)

[103] This is not a crude division of the two streams into a computational model for the ventral stream and a connectionist model for the dorsal stream. Both models probably contain some partial truths about brain functioning. However, the visuomotor system, on the Milner-Goodale thesis, is better described via dynamical modelling (which can, but need not, be connectionist) than computational modelling. There is a good general account of dynamical modelling, using the Watt Governor as an example, in T. Van Gelder, "What Might Cognition Be, If Not Computation", in "Mind and Cognition", W.G. Lycan (Ed), Oxford: Blackwell, 1999, p 170 – 191. Van Gelder contrasts the Watt Governor with an obvious alternative; a computer based governor that measured steam pressure and flywheel velocity and adjusted the former on the basis of the latter. The computer-based system, unlike the Watt Governor, is representational.

would be that the higher primates evolved their complex conscious processing from their simpler precursor species.

There are two sorts of evidence for the Milner-Goodale model. One is a study of the properties of the neurons in the two streams and the other the performance of individuals with lesions in the ventral or dorsal streams.

The relevant evidence on the visual properties of neurons comes primarily from the study of monkeys. Neurons in the ventral (perceptual) stream areas, according to Milner and Goodale, are tuned to the features of objects and are little affected by the monkey's motor behaviour. Neurons in the dorsal (action guiding) stream behave differently and are activated by visual stimuli as a function of the different kind of motor response the monkeys make. The experiments, their results and their interpretation are extremely complex and controversial. There is also debate as to the applicability to humans of the results on monkeys. Perhaps the most that a lay reader can say is that the results can plausibly be interpreted in accord with the Milner and Goodale account.

The experiments on the macroscopic behaviour of animal and human subjects are more accessible to non-specialists. The most dramatic results are those reported by Milner and Goodale on a patient DF who suffered major damage to the ventral stream as a result of carbon monoxide poisoning. She is unable to process simple visual information but demonstrates the ability in actions to make skilled use of visual information of which she has no awareness. Before describing this work, it is useful to look at some

experiments on non-human animals and the impairment of their macroscopic abilities as a result of different lesions.

The basis of the Milner-Goodale view is that the requirements for visuomotor control for actions like grasping or running swiftly through irregular terrain will be egocentrically specified. To grasp an object requires a rapid mechanism that relates its position to that of the grasper. A system for visual perception, on the other hand, needs to provide data on the perceived object that is substantially independent of a particular perspective. The object needs to be recorded in memory in a fashion that allows recognition, whatever its later spatial presentation.

Early studies of bilateral temporal lobe lesions showed that such ventral stream lesioned monkeys had great difficulty in visual recognition but retained motor skills. They could still judge distances when jumping and did not bump into obstacles[104]. This has been confirmed in more recent work where IT lesioned monkeys were skilled at catching gnats despite being deficient in pattern recognition. Confirmation of the Milner-Goodale hypothesis has been found in relation to the dorsal stream by experiments in which microinjection of the drug muscinol into a particular part of the posterior parietal cortex

---

[104] Experiments of H. Kluver and P.C. Bucy, "An analysis of certain effects of bilateral temporal lobectomy in the rhesus monkey, with special reference to psychic blindness", J. Psychol. 5, 33 – 54 (1938), cited in Milner and Goodale, *loc cit*

causes an impairment in hand shaping when the monkey reaches to grasp objects[105].

I turn now to the striking results on the patient DF with ventral stream damage. DF is unable to recognize faces or objects. For example, she is unable to recognize or discriminate between basic geometrical shapes. Only when rectangles have a length to breadth ratio of more than two to one can she distinguish them from squares. She has great difficulty in recognizing simple line drawings and in copying them, although she can draw the same objects from memory quite well. These symptoms of what is called "visual form agnosia" are common in such subjects. By contrast, her sight in a sense is quite satisfactory. Like other visual form agnostics, she is good, for example, at catching a ball.

Miller and Goodale performed a series of careful experiments using an apparatus with a disc in which a slot was cut. The slot could be randomly set at 0, 45, 90 or 135 degrees from the vertical. When DF was asked the orientation of the slot, her answers were hopelessly inaccurate. Indeed, they were more or less randomly distributed in relation to the actual orientation of the slot. This was true whether she was simply asked how the slot was oriented or whether she was asked to orient an identical dummy slot in the same position as the test slot.

---

[105] V. Gallese, L. Fadiga, L. Fogassi, G. Lupino and A. Murata, "A parietal-frontal circuit for hand grasping movements in the monkey: evidence from reversible inactivation experiments", in P. Their and H.O. Karnath (Eds), "Parietal lobe contributions to orientation in 3-D-space" Berlin: Springer-Verlag, 1997

The striking observation came when DF was asked to post something through the slot. This she could do with great accuracy. Moreover, video recordings show that her hand begins to rotate appropriately right from the commencement of the posting action. She knows how to post something through the slot. It is also clear that the subdoxastic states, which enable her to do this posting, are not inferentially integrated with any beliefs about the orientation of this same slot. She has no such beliefs. In a forced choice situation she will make attempts to express beliefs. These, however, bear no relation to the slot orientation.

In the same way, DF displayed ordinary skill in the motor action of the grasping function for a variety of objects, although when asked to indicate her perceptual judgements of the width, say, of the same objects, her abilities were once more very poor.

This brief account does justice neither to the detailed experiments nor to the complexity of their interpretation. I have attempted only to present evidence that I believe is sufficient to show that recent experimental work undermines our commonsense view that we control our visuomotor actions via our conscious visual observations.

Thus, as I said above, it turns out that contemporary research has shown that some of our deliberate acts are not the result of processes under our conscious control. To emphasize the point made earlier, I do not control my fine tuned motor actions by conscious visual guiding, as common sense would tell us. In fact, as I described, my conscious vision instructs my *alter ego* (namely me in my other, unconscious visual mode) to perform the task.

Of course, the experimental evidence does not show that what I have said is true for all visuomotor actions. The mind/brain may operate in task specific ways. Moreover, the experimental evidence does not force the unlikely conclusion that the conscious and unconscious processes are always totally non-interacting (indeed there must be some high level interaction for normal functioning[106]). In any case universal claims would have to await the remote availability of a complete theoretical understanding. However, universality is not required for my argument.

5.5 Knowing-how and the Embodied Person

I think it has now been firmly established that our visuomotor performances are in many cases controlled by an unconscious mechanism. I will assume that the arguments and evidence presented in the earlier sections of this chapter establish this fact. What then are the consequences for my claims about knowing-how? The claim I have made is that knowing-how is ability, but a definable and usefully definable subset of ability. Moreover, as is entailed by casting knowing-how as ability, it is distinct from knowing-that.

I argued in section 5.2 that one's unconscious knowledge of Universal Grammar is not knowledge in the ordinary sense, as Chomsky seems to assert. I argued that it was knowing-how (the way I have defined knowing-how) and that it was at the same time distinct from many abilities, like breathing in a controlled fashion.

---

[106] A.D. Milner and M.A. Goodale, *ibid*, p 203

If then it is agreed that returning a cannon-ball serve from Sampras is done via an unconscious process mediated by the conscious desire to return the serve, what additionally has been established?

Chomsky has the following directly relevant and extremely interesting things to say in his well known "Rules and Representations"[107]:

> "True, there it would be little point to a concept of "cognizing" that did not distinguish "cognizing the rules of grammar" from the bicycle rider's knowing that he should push the pedals or lean into a curve, given what we assume to be the facts of the matter. But it seems easy enough to make the relevant distinction. In my case of riding a bicycle, there is no reason to suppose that the rules in question are represented in a cognitive structure of anything like the sort I have described. Rather, we take bicycle riding to be a skill, whereas knowledge of language and its concomitants, for example, knowledge that reciprocal expressions have the properties I mentioned, is not a skill at all. The skill in question might, perhaps, be based on a certain reflex systems, in which case it would be incorrect to attribute cognitive structure in the sense of this discussion to the person who exercises the skill. In contrast to the case of language, nothing seems to be explained by attributing to the bicycle rider a cognitive structure incorporating the rules with which his practice accords. But suppose we are wrong, and in fact the bicycle rider does have a representation of certain physical principles in his mind and uses them to plan or compute his next act. In this case we should attribute to him a cognitive structure, and in fact, it would be quite appropriate to say that he cognizes these principles as he does the rules of his language. The question, I take it, is basically one of fact.
>
> To help clarify the issue, consider two missile systems, each of which is designed to send a rocket to the moon. One of them operates along lines once proposed by B. F. Skinner; it has several pigeons looking at a screen that depicts what lies directly ahead, trained to peck when the rocket veers off course, their pecking restoring the image of the moon to a focused position on the screen. Consider in contrast a system that incorporates an explicit theory of the motions of the heavenly bodies and information about its position and velocity and that carries out measurements and computations using its internalized theory to adjust

---

[107] N. Chomsky, "Rules and Representations", New York: Columbia University Press, 1980, p 102

> its course as it proceeds. This rocket might hit the very same spot as the servomechanism with the pigeons, but it would do so in a very different way. Mere investigation of behaviour might tell us little, perhaps nothing. A deeper look might be required to distinguish the two systems. In the second case, but not the first, enquiries might lead us to attribute to the missile something like a "mental state." That is, it might lead us to formulate an abstract characterization of perhaps unknown mechanisms, postulating a system which involves the cognizing of certain principles and computation involving these principles and certain representations. In the first case, such an account would be factually wrong. I think that the two cases fall on opposite sides of an important divide, and that the second - the cognizing missile - shares properties with human knowledge. It also lacks crucial properties; for example, it is a task-oriented device, whereas knowledge of language, for example, is not."

There are here four important points:

(i) the separating off of the category of reflex systems;

(ii) the identification of the category, via the pigeon metaphor, of a category not involving representations;

(iii) the identification of the category with the explicit knowledge of the laws of motion, involving representations and computations, and thus involving something like a computational theory of mind (for the relevant part of the modular human brain); and

(iv) the distinction between the "cognizing missile" (in (iii) above) and knowledge of language.

The argument in section 5.4 establishes that, as a matter of experimental fact, human visuomotor actions like grasping are in fact very like Chomsky's rocket. In many visuomotor actions, we are task-oriented automata. It is precisely in this that knowing-how to ride a bicycle consists. This is not to denigrate knowing-how but to distinguish it, on the one hand, from the reflex

act of blinking and, on the other, from knowing-that, say, the square root of two is an irrational number.

However, what the Milner and Goodale evidence makes clear is that in, for example, grasping an object, we do not use a representational system like Chomsky's cognizing missile, but we are well described by Chomsky's metaphor of Skinner pigeons. The latter would be a complex arrangement of what I have earlier called basic tacit mechanisms (see section 1.3.2). The truth is we do not know enough about the mind to characterize its operations in detail. However, at the level where we are talking about things like knowledge, we have a clear distinction between knowing-that and ability, and within the latter, a way of characterizing knowing-how.

Knowing-how at the more cognitive end includes the unconscious subdoxastic states governing our language capacity. At the less cognitive end knowing-how includes many of our fine-tuned motor actions.

Thus *contra* Chomsky, I claim:

(i) Our linguistic ability is knowing-how and incorporates unconscious subdoxastic mental states.

(ii) The important divide is not between these states (which are implicated in linguistic competence) and the unconscious subdoxastic states governing our visuomotor actions, as these are also knowing-how.

(iii) The important divide is between knowing-how and knowing-that.

This is not to deny that the metaphor Chomsky uses to distinguish between the two types of rocket guidance system is a useful one. It is simply to claim that both guidance systems are knowing-how and linguistic capacity falls under the more cognitive (more explicit) one.

To put all this in another way, the behaviour of the subject DF tells us that her understanding of the orientation of the slot is simply a bodily understanding, and, as I have argued, this is a general proposition about much of our fine-tuned motor action. The usual content attitude distinction, which is the hallmark of knowing-that, cannot be applied to these motor actions[108]. The visual experience, as Merleau-Ponty pointed out, is an experience had by the whole animal.

[108] S.D. Kelly, "Merleau-Ponty on the Body", Ratio, 15, 376 – 391 (2002)

CHAPTER 6. KNOWING-HOW AND QUALIA

> *The claim that all knowledge derives from knowledge of experience may have been overblown, but the general point that there is something special about our knowledge of experience has never been overturned.*

David Chalmers[109]

6.1 Qualia as Ability and as Knowing-how

There are philosophers who believe one cannot meaningfully talk about the fact that there is "something that it is like" to have certain sense experiences. That is, they do not believe in the existence of qualia[110]. If there are no qualia, the argument I am about to make - that for individuals to experience certain types of qualia constitutes an example of knowing-how - obviously collapses. I assume that qualia do exist on the basis of numerous powerful arguments[111] and, indeed, in accord with mainstream philosophical opinion.

There are various positions on the problem of qualia, on "what it is like" to be in certain mental states, typically perceptual experiences, like seeing a colour, and bodily sensations, like itches and pains. I claim no more for qualia than

---

[109] D.J. Chalmers, "The Conscious Mind", Oxford: Oxford University Press, 1996, p 197.

[110] Some of the most convincing objections can be found in D.C. Dennett, "The Path Not Taken", in "The Nature of Consciousness", Ed N. Block, O. Flanagan and G. Guzeldere, Cambridge: MIT Press, 1999, Chapter 21 and D.C. Dennett, "Consciousness Explained", Boston: Little, Brown, 1991.

[111] See, for example, D.J. Chalmers, "The Conscious Mind", Oxford: Oxford University Press, 1996, N. Block, "Inverted Earth" in "The Nature of Consciousness", Ed N. Bock, O. Flanagan, and G. Guzeldere, Cambridge MA: MIT Press, 1997, Chapter 43 and D. Papineau, "Thinking about Consciousness", Oxford: Oxford University Press, 2002

this minimal account of them as "what it is like" to have, say, a pain or a colour sensation. That is, I specifically disavow making any claim that qualia are non-physical or ineffable or given incorrigibly, since these claims are irrelevant to my claims about knowing-how. Nor do I believe that any commitment to sense-datum theories is involved in accepting the reality of qualia.

There is a famous argument, commonly called the "Knowledge Argument", made by Frank Jackson against physicalism. The Jackson argument is presented in terms of Mary, a complete expert on a hypothetical completed colour science, who has only ever seen black and white. It is argued that on seeing a red object, Mary learns something new and hence physicalism is false, since Mary now knows something she did not know before[112]. Therefore, it is argued, "there are truths about other people (and herself) which escape the physicalist story"[113]. The point, of course, is that *ex hypothesis* Mary had absolutely all of the knowing-that on colour.

One response to this is to say that what Mary has learned is not some new information but some knowing-how. Now, since she has seen red, she knows what it is like to see red, by the qualia, by what it is like to experience redness.

---

[112] Dennett responds by saying that our intuition that Mary knows something that she did not know before is false. We are misled into this false intuition because we do not fully recognize and take account of the exceptional nature of Mary. Mary, by definition, has superhuman expertise on colour and she would, claims Dennett, know precisely what it is like to see red while never having done so (D.C. Dennett, "Consciousness Explained", Boston: Little, Brown, 1991).

[113] F. Jackson, "What Mary didn't know", Journal of Philosophy, 83, 291 – 295 (1986), p 293

However, it has been said, in defense of the Mary thought experiment being valid in rebutting physicalism, that there is also some knowing-that involved in this new knowledge.

This last is not my concern, nor is the defense of physicalism[114]. My concern is the claim – not negated by the possible co-presence of knowing-that – that we have, in at least some cases of qualia being experienced, exemplars of knowing-how. I will come to these examples shortly. My immediate interest in Jackson's Mary case is that this is in fact ability and not knowing-how. I will contrast this case with more elaborate qualia experiences that do fall under my definition of knowing-how.

Most proponents of the "knowing-how" defense of physicalism are content to characterize what Mary acquires as ability, and, it seems to me, that by the criteria I have advocated thus far, to be able to sense a colour or feel a pain or an itch is simply ability. It is not within the subset of ability I have defined as knowing-how. (Some of the writers use knowing-how and ability as synonyms, whereas I have emphasized that knowing-how is ability, just as suicide is killing. The former is a subset of the latter but they are certainly not synonymous terms.)

---

[114] There is a huge literature around the Knowledge Argument. Numerous writers have defended physicalism against the Knowledge Argument, and Jackson himself has repudiated the argument as conclusive against physicalism. For a brief account see, for example, D. Braddon-Mitchell and F. Jackson, "Philosophy of Mind and Cognition", Oxford: Blackwell, 1996, Chapter 8.

From the definition it is clear that what it is like to feel pain, say, does not qualify as knowing-how; it seems to fail on counts one and five, and, as a consequence of these it fails also on three.

1. A can entertain φing as a purpose.
2. A is acquainted with a set of practical procedures for successful φing.
3. A exhibits or has exhibited reasonable success at φing.
4. A complete propositional account of φing is not available to A.
5. A must have learned to φ via a process of instruction which may include self-instruction (and may be an unconscious process).

Moreover, there is little reason in defending physicalism to claim knowing-how for Mary as distinct from ability

I am, of course, not here entering into the central debate about the Knowledge Argument. I am, however, directly interested in the question as to whether knowledge of qualia, in the sense of knowing what it is like to have a certain sense experience, could sometimes be an instance of knowing-how.

In contrast with Jackson's Mary, consider the case of Annabelle, who by a deliberate effort of intense study and regular concert attendance comes steadily to better appreciate Mozart's symphonies. She acquires a vast amount of propositional knowledge about Mozart's life, about 18th century Vienna, about the structure of the works, about different interpretations; she now can get musical pleasure from reading the scores and imagining the pieces being played and so on. It would be generally agreed that she may

obtain greater enjoyment from a performance, she may be more moved and all of the phenomenological part of what arises from her listening may be enhanced for her Let us suppose for Annabelle that this is the case. (It does not matter for the present discussion that another individual, who undertakes similar study, may not become responsive in this fashion.)

Now some of Annabelle's changed response is knowing-that. For example, she may get intellectual pleasure from reflecting on the role of the strings versus the wind instruments. The essential point for my purposes is that there is now a difference in "what it is like" for Annabelle.

Φing in this case is not cycling or skiing but being moved and joyous on listening, say, to the Jupiter symphony, having the pleasurable sensations music lovers are familiar with on hearing a favourite piece of music under favourable circumstances. Of course, in all three cases there is knowing-that as well as knowing-how in the activities. Annabelle's qualia are as real as the familiar qualia associated with a toothache. The latter come unbidden but in the case of Annabelle's music this form of φing seems to satisfy the definition on all five counts. Φing seems to be cognitive and it is not exclusively propositional knowledge. A quibble might be that she has not learned to φ from scratch, she was able to φ before her study, but she has greatly changed what it is like to listen to Mozart and this seems to be all that is required. After all a talented child can probably stay upright on a bicycle for a few seconds when she first tries.

It seems clear that Annabelle's joy in the music has been enhanced by the deliberate acquisition of propositional knowledge. Her learning to Φ in this

case is different from learning to improve her violin playing or skiing by practice. It is not entirely different, as in both of these an improved knowledge of theory could improve knowing-how in a way that is somewhat different from practicing. The line is, as usual, somewhat fuzzy. Watching videos of tennis games against a forthcoming opponent, with a good coach present, one might learn to stand a little deeper in receiving her first serve - a piece of propositional knowledge. However, in actual play this would be integrated with knowing-how and a skilled player would not make a rule of this sound piece of advice.

The case of qualia does seem to illustrate the point made on a number of earlier occasions about the fuzziness of classification. There are experiences like the taste of Vegemite or Marmite where recognizing these is a simple ability for Australian and British philosophers respectively. On the other hand, there are, even in the same sense modality of taste, cases of ability that involve knowing-how, like the nuanced qualia of the professional wine taster.

I am claiming that some experiences of qualia provide clear instances of knowing-how. We do not need to know any phenomenal propositions about the qualia themselves to know-how to enjoy certain qualia. We can learn to enjoy, and plan, intend, and desire to enjoy, and commonly succeed in enjoying certain qualia.

This is not to claim that there are no phenomenal propositions. We can describe something of our subjective experiences, so that someone who has similar experiences can recognize that they are similar. We might say to someone that the experience of a particular shade of red is rather like the

experience of a shade of orange we shared on some previous occasion. Our companion would then at least have been given a little propositional knowledge to enable her to imagine our experience but is not necessary for her or our experience.[115] There are plenty of propositions about how to play tennis in innumerable books, but they are neither necessary nor sufficient for knowing-how.

6.2 Knowing-how of Qualia v Conventional Knowing-how

It might now seem that I have advocated two classes of knowing-how. There is the knowing-how of the mathematically minded individual, who has a sense, for example, of how to work out the underlying pattern of an unobvious mathematical infinite series. Separately, I have in this chapter advocated the knowing-how of the music lover who cultivates, by a rather similar process, the knowing-how of maximally enjoying qualia. If the learned music lover decided to compose a musical piece, it seems I would credit her with knowing-how (in this role as composer) that is identical with the knowing-how of the mathematician. What is the difference between the music lover qua composer and qua listener? Clearly, for both there is an inescapable component of knowing-that.

At any instant in one's life there is something "that it is like". There is something that it is like to be revising this thesis now. The only reason why the learned qualia of Annabelle fall under the knowing-how subset of ability is

---

[115] This point is made in a different context by Noordhof in a recent article on the Knowledge Argument (P. Noordhof, "Something Like Ability", Australasian Journal of Philosophy, 81, 21 – 40 (2003))

that she has cultivated them, just as she might have learned to ride a bicycle. So why are they knowledge in the sense that speaking one's idiolect is knowledge? Qualia are the least understood of the outcomes of our mind/brain. Perhaps the most that can be said is that it is plausible that they result from the subdoxastic states discussed in Chapter 4 and, as in the language case, the systematized manipulation of information via these states may result in beliefs.

6.3 Qualia of Knowing-how v Qualia of Ability

Is there something definable that is distinctive about the qualia that fall under knowing-how? As I have characterized them, they are such that one must learn to experience them, so pain and itches and perfumes are out and a sophisticated emotional response to Mozart symphonies and a sophisticated olfactory sensing of perfumes or a sophisticated gustatory sensing of wines are in. This is complicated by differences in natural endowment. There is an artificially created substance called phenylthiocarbamide (PTC) that has a very strong distinctive bitter taste for about three quarters of the population and none at all for the remainder. The ability to taste PTC is inherited. Likewise, some peoples' sense of smell is vastly more sensitive than others (hyperosmia); indeed these individuals can smell substances that have no smell for normal people. Some individuals, without having to learn, will have ability, say, for distinguishing perfumes that others of us do not. We can perhaps acquire some degree of this by learning; for us it is then knowing-how. I think a similar argument can be mounted for the experiencing of many other qualia.

Is this last then a defect in the stated definition of knowing-how? I think not; it is inherent in the fuzzy boundary between abilities that are knowing-how and abilities that are not. This simply parallels the fuzzy boundary of the issue of the functions like pupillary contraction and the maintenance of blood pressure that are not ability and those that are, like walking and controlled breathing.

This discussion of qualia, in addition to providing another exemplar of knowing-how, has highlighted the complexity of any analysis of knowledge from a behavioural perspective. This complexity is not surprising given the neurological complexity, for example, of the human visual system alone, as discussed in the previous chapter.

CHAPTER 7. KNOWING-HOW AND VIRTUE

> *Virtues, however, we acquire by first exercising them. The same is true with skills, since what we need to learn before doing, we learn by doing, for example we become builders by building, and lyre players by playing the lyre. So too we become just by doing just actions, temperate by temperate actions, and courageous by courageous actions.*
>
> Aristotle[116]

7.1 Introduction

As mentioned at the end of Chapter 3, there is a view of morality that virtue consists of how it is good to be rather than what it is right to do. This is a view of ethics best known from Aristotle, and in its various contemporary forms is often called neo-Aristotelian ethics or virtue-ethics.

The interest in this chapter is whether acting virtuously can be claimed as a form of knowing-how.

The argument that this might be the case goes as follows.

(i) Actions seem to be closely related to knowing-how; unlike propositional knowing, one seems to know how to do an action, even when it is not a habit.

(ii) There are two models on which one could proceed with some kind of action explanation – or practical reasoning. The first involves a deductive law-like model that explains an action through

---

[116] Aristotle, "Nicomachean Ethics", 1103b

subsumption under a general rule. In the second, explaining an action involves showing how it was the right thing to do from the perspective of the agent, requiring talk of salience and a holistic view of the context.

(iii) The former model arguably seems to makes action explanation a form of, or at least related to, knowing-that. The latter model seems to makes action explanation a form of, or at least related to, knowing-how.

(iv) One argument for knowing-how in the arena of action is the claim that any human action, whether in the area of ethical behaviour or not, cannot be governed by psychophysical laws. This is claimed by Davidson[117] to follow from the claim that belief is based on rationality, which cannot be codified. If this is warranted, the case for knowing-how follows.

(v) If this argument fails to go through in general for human action, ethical behaviour is nevertheless a good candidate for knowing-how, in that it can still be argued in the ethical case, from the stance of virtue ethics, that we have a situation that falls under the second of the two scenarios outlined in (ii) above. That is, moral judgement necessarily involves considerations of agency and cannot be subsumed under general rules.

---

[117] D. Davidson, "Essays on Actions and Events", Oxford: Oxford University Press, 1980

The Davidson argument, briefly sketched in (iv) above, is based on his anomalous monism, a theory of mind quite different from the one underpinning this thesis. As I indicated in the first chapter, I am assuming a representational theory that thought proceeds, at least in part, by computational operations on syntactic elements. The Davidson model allows of a different approach to explaining and licensing the concept of knowing-how. I discuss this briefly in Appendix B.

In this chapter, I discuss virtue in the same spirit as I have considered language and qualia, and do not consider the Davidson picture of required rationality and uncodifiability. Thus, I am arguing for the case outlined in (v) above. If we accept virtue ethics, then virtue is not based on knowing-that, it is based directly on knowing what to do. That is, it is knowing-how.

I consider, in section 7.2, both some aspects of the ethics debate that provide support for taking virtue ethics seriously, and also one of the principal counter arguments. Although, in a sense, it is self-evident that the acceptance of virtue ethics entails the acceptance of knowing-how, I elaborate on this in section 7.3. However, I do believe the big question is whether virtue ethics is tenable.

7.2 Virtue as a Character Trait

7.2.1 Learning by Doing

As just mentioned, there is a view of morality that virtue consists of how it is good to be, rather than consisting of what it is right to do (deontology) or third, what will have, in some sense, the best consequences. If the first of these is

true then it might seem that virtue is knowing-how, contrasted say on the one hand with a Kantian deontology based on the rational agent, or, on the other, with the careful consequential calculation of the utilitarian.

Kant's well-known, cold-hearted benefactor perhaps best illustrates the strength of virtue ethics. Imagine a woman who willingly visits her sick friends in hospital and indeed is considerate towards all her friends in every circumstance. She is, however, lacking in warmth, and does these good acts from a sense of duty, albeit as reliably and willingly as anyone could. There is an intuitive feeling that she is less virtuous than she could be. We want our friends to visit us out of affection. Her behaviour is good but we feel not as good as it could be. Virtue ethics has a ready account of this, whereas a deontological or teleological account is more difficult. Friendship is noble, as Aristotle would say.

It might be said that virtue ethics is a theory of character. Deontology and utilitarianism (in its various forms, the primary exemplar of instrumentalism or consequentialism) are theories of action based on principle. It seems plausible, nevertheless, that virtue ethics and deontology are not mutually exclusive. As I have discussed, some knowing-that is a part of much knowing-how. Indeed, virtue ethicists argue that, for example, kindness, generosity, honesty and courage are virtues, in answer to the charge that their theory gives no moral guidance unless one knows oneself to be virtuous[118]. These qualities are excellences and are virtues in the same sense that it is a virtue

[118] R. Hursthouse, "On Virtue Ethics", Oxford: Oxford University Press, 1999, p 36.

for a knife to be sharp. Moreover, the categories of character-based and action-based are somewhat blurred; there are virtue ethicists who claim that their form of the doctrine is not character-based[119], and some deontologists would see their doctrines as falling into both camps, e.g., the doctrine of original sin and the need to resist temptation can be combined with injunctions like a prohibition on adultery. Moreover, in making judgements we distinguish between the agent and the action via the action explanation. Someone may be blameless in performing an action that is not virtuous, for example, in neglecting a minor duty in the wake of a personal tragedy.

The question is whether there is in ethics an area of knowing-how that can be distinguished from knowing-that. For the present purpose, much of what is commonly discussed as conventional moral doctrine is not relevant. Most philosophers who discuss ethical questions are themselves in comfortable circumstances, and consequently are not tempted to murder or physically assault or steal. Moreover, most of them would rescue a drowning animal or human, if not thereby exposing themselves to great danger. Further, they would be well versed in the conventional moral code of their society. They would know that one should not be cruel or unduly selfish. It is, of course, possible to propose intellectual moral dilemmas – with only enough vaccine to save one person, should it be given to your unremarkable daughter or your equally unremarkable mother or to a young stranger who is a uniquely skilled physician? Clearly, people can and do have extensive debates of a reasoned

---

[119] Examples are given in G. Harman, "Explaining Value", Oxford: Oxford University Press, 2000, p 176.

kind about such questions. However, aside from such debates, there are, among the philosophers who discuss ethical questions, some people who are good and some who are not. There is little doubt that the difference between these is not related to their propositional knowledge of a moral code. Two firm believers in a common doctrine can be entirely different in their moral behaviour.

A widespread contemporary view of this phenomenon is usually couched in terms of the ability to resist temptation. This seems to be particularly important in the Christian tradition, with its doctrine of original sin. The more virtuous individual resists the temptations to which both are subject. The Aristotelian view is that the more virtuous person would, by having learned to practice virtue and having learned by practicing virtue, choose the path of virtue. In other words, the worthier individual would not have the same feeling of temptation to do wrong. The Aristotelian position is that the virtuous person does the virtuous act because it is right and not because she thinks it is right. (She might indeed think it was fun or think it would make her happy or she might think her late grandmother would have been pleased.) Does this person know-how to be good to a greater degree than the other? Is there a parallel between this and, say, riding a bicycle well or poorly? Note that both of the views just discussed incorporate a notion of character traits. There is in the former case a resolute character built on a determination to do what is deemed right, in the latter (Aristotelian) case a disposition to be good built on

repeatedly having done good[120] or, as Hursthouse[121] puts it, based on "a settled state of good character". The clearest statement for my purposes is perhaps that given by McDowell[122],

> "Occasion by occasion, one knows what to do, if one does, not by applying universal principles but by being a certain kind of person: one who sees situations in a certain distinctive way"

That one is this "certain kind of person" does not necessarily depend on Aristotle's view of how one comes to be this kind of person. Mc Dowell's statement of the Aristotelian concept of virtue rightly emphasizes the point that the virtuous person knows the "thing to do". She may well not consider her acts as virtuous or, if she does, may not characterize them as belonging to some specific virtue like kindness or honesty.

Commonly, empirical differences between virtue and its absence are not found in differences in rational deliberation. Good acts are commonplace and in general are an immediate response to a mundane situation. We stop to enquire from a weeping stranger whether we can help, we offer to help carry a pram down some steps, we avoid continuing a conversation that is inadvertently causing embarrassment, we point out to the shopkeeper that we have been given too much change, we hurry and are anxious to avoid keeping a friend waiting, we admit our mistakes and so on.

---

[120] Burnyeat puts this particularly clearly, "Aristotle is not simply giving us a bland reminder that virtue takes practice. Rather, practice has cognitive powers, in that it is the way we learn what is noble and just" (M.F. Burnyeat, "Aristotle on Learning to Be Good", in Essays on Aristotle's Ethics", Ed A.O. Rorty, California: University of California Press, 1980, p 73.)

[121] R. Hursthouse, *ibid*, p 123.

[122] J. McDowell, "Virtue and Reason" reprinted in "Mind, Value, and Reality", Cambridge: Harvard University Press, 1998, p 73.

What I call the Aristotelian position[123] on this observation is that we have learned by our upbringing and previous experiences to behave in certain ways, and these behaviours are not mere habits but are in fact how we navigate the world. Each new situation is dealt with by the learned behaviours from previous experience. We experience the world interactively, our interactions are guided by our perceptions, and our interactions form and transform our perceptions. This is not a simple input-output model but an inescapable ongoing feedback mechanism, so that we develop virtue (or this form of knowing-how) as we grow up. As Varela[124] puts it, reality

> "is perceiver-dependent, not because the perceiver "constructs" it as she or he pleases, but because what *counts* as a relevant world is inseparable from the structure of the perceiver."

Aristotle, of course, notoriously believed that certain perceivers were not capable of attaining virtue by their very nature – notably slaves and women.

This learning by conscious doing applies to our ordinary activities, from mountain climbing to dining in a restaurant to "reading other minds"[125]. It thus

---

[123] I am not concerned here with the issue of whether contemporary virtue ethics, in one or other of its now many strands, can justly be seen to be the position taken by Aristotle himself. The question is only whether a position that claims that virtue is determined by the behaviour, in a given set of circumstances, of a virtuous person, entails that knowing-how is involved in behaving virtuously.

[124] F.J. Varela, "Ethical Know-How", California: Stanford University Press, 1999, p 13.

[125] As John Dewey said as long ago as 1922, "We may indeed 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." (J. Dewey, "Human Nature and Conduct: An Introduction to Social Psychology", New York: Random House, 1957, p 177) The present argument, however, says there is more to this than habit formation. Dewey is clear about this last in the area of virtue, elsewhere he says, "Morals means growth of conduct in meaning; or at least it means that kind of expansion in meaning which is consequent upon observation of the conditions and outcome of conduct" (*ibid*, p 280).

applies to developing abilities as well as to that end of the ability spectrum I have classified as knowing-how. On this view, propositional knowledge, involving conscious learning and/or reflection, or perhaps more generally, learning that does not involve any sensorimotor activity, is in a special category and the norm is knowing-how. In "learning" knowing-how we perceptually experience a situation within which we are perceptually guided, and we generalize from this experience in future similar situations. Importantly, this process goes beyond what we could categorically identify and hence is different from the learning involved in knowing-that. The perennial teacher's question, "Did you understand that?" is not appropriate in the ordinary way, although attentiveness to experience is essential for the getting of wisdom. The process of becoming virtuous consists of developing a disposition to act immediately in certain ways, rather than reflecting on what it is right to do. And here there is indeed a parallel with the cyclist or skier or tennis player.

The distinction drawn in the previous paragraph, between propositional knowing and other knowing, is necessarily fuzzy. The term sensorimotor activity I have used is imprecise. We know that an object two metres away is out of reach. We have learned this by a most complex series of mental processes, most of which are inaccessible to introspection, including a scanning process whereby we have focused attention on different parts of our visual field. So we have learned how to judge distances and we now know that some objects are near and others far, by merely opening our eyes.

The knowing-how of virtue is knowing-how to act, although it must be conceded this action might on occasion be inaction. All that is required for virtue in certain circumstances might be not to react at all, for example, when a friend suddenly remarks that she believes her husband is having an affair with a mutual friend. In these circumstances avoiding the issue might be the best way to minimize suffering, if one knew the belief was true.[126].

7.2.2 Ethical Disagreement

The view that ethics is knowing-how has been propounded from a somewhat different angle by Cross in a careful article entitled "Ethical Disagreement"[127], where he considers precisely the topic indicated by his title. In the sense I am using it, Cross would be classified as a virtue ethicist, since he argues specifically that virtue is knowing-how. He does not do this from an Aristotelian stance; for example, he says nothing about how virtue is acquired. He moves to a position that I would call virtue ethics because he rejects the alternatives.

To make the discussion concrete I consider an example (which is not quite the procedure followed by Cross). Suppose Cardinal Francis Joseph Spellman says abortion is always wrong and Professor Peter Singer says abortion is sometimes right. One well known set of so called non-cognitivist theories says they are not in disagreement; all that is being said is that

---

[126] This is perhaps stated too crudely; the inaction would have to be for the right reasons, although the virtuous person would, of course, not need reasons or consider them. She would simply act appropriately.

[127] R.C. Cross, Philosophy, 25, 301 – 315 (1950)

Spellman never approves of abortion and Singer sometimes does (emotivism). A parallel might be that Singer likes nectarines and the Cardinal does not, and they agree to differ – there is no dispute between them. (It is necessary for my purposes to confine myself to cases where there is agreement about the exact nature of the action, and the disagreement is about the response to the action, given the agreed facts.)

It is clear that the case of abortion and that of liking nectarines are not the same. There is a real disagreement about abortion; and it is a difference in attitude. The Cardinal will use the emotive words "wrong" or "sinful" in describing abortion, in an attempt to persuade us to his way of thinking (on one view because he believes his moral statements are meaningful in prescribing action). He will also perhaps cite the divine authority of Jesus. Likewise Singer (a utilitarian) will speak of the distinction between a person and a foetus, and hence of the moral distinction in terms of consequences, between killing a neo-nate or a foetus, on the one hand, and a person on the other. Singer might also say something about the rights of women to make choices, for example, in the case of a pregnancy resulting from rape, or in an instance where the life of the mother is overwhelmingly threatened by allowing the birth of the child[128]. This attempt to persuade is important and makes for the difference between the preferences of taste versus the views on what is seen as a matter of ethics. Further, this difference in attitude on ethical questions will, in most cases of serious ethical disagreement, be perceived by

[128] I cannot resist adding, irrelevantly, that I find unacceptable both the late Cardinal's attitude to women who have been raped and Professor Singer's attitude to healthy or unhealthy newborn infants.

both parties to have consequences and/or involve acts that in turn are seen to be important.

The crux of the matter is that, in fact, we are not content to say simply that Spellman and Singer differ in their attitude on abortion. On the contrary, we believe that one of them is wrong and the other right. However, it seems impossible to make the true false distinction of propositional knowledge here, as we could if Singer were asserting that π is a rational number or that sheep are on average bigger than horses. In other words, how can one make the true false or wrong right distinction when we are talking about emotions or attitudes and not propositions[129]? Rather than paraphrase or attempt to summarize further it seems best now to quote Cross. He says[130],

> "I do not think what is primarily involved in morals is knowing that something is the case, any more than having certain feelings. I think that what is involved rather is knowing how to act, how to cope with situations, and that our ethical utterances are neither true or false ... nor merely emotive.... Further, the procedure to be adopted does not depend on how we happen to feel, but there are correct procedures and incorrect procedures for dealing with the situations in which we are placed."

Perhaps this is most easily seen to be a reasonable stance when we consider one of the mundane examples I gave earlier. I said it would be good to stop and ask a weeping stranger if we could help in any way. However, this is too crude, and in fact whether a good person would do so depends very much on

---

[129] I am here describing the argument put forward by Cross. It seems to me that there may be a way of making the case that that Singer and Spellman have mental states that are in some way representational and related to the way things are in the world. There might then be a way of asserting that one or other is correct in holding their position.

[130] R.C. Cross, *ibid*, p 314

the circumstances. Is the stranger trying to conceal the fact that she is weeping, is she alone, if not, is she in seemingly sympathetic or hostile company or is she, for example, sitting quietly by a graveside? Moreover, our response would be immediate, and commonly would not involve an analysis of the sort I have just given. We would know-how to behave or we would not know-how.

One might argue against Cross by asking whether a heated dispute between two people who believe that the Beatles are better composers than Mozart and *vice versa* is a knowing-how issue. Or is it a knowing-how issue in a dispute about whether Arsenal is more elegant than Leeds or Paris is more beautiful than Prague? Here it may seem that if we are to take the Aristotelian position, it extends well beyond ethics. Certainly, there is often the same feeling (unlike, at least usually, in matters of the taste of food) that one party is right, but it seems implausible that the superior beauty of Paris or Prague is a matter of fact; it is surely in the eye of the beholder. However, these comparisons with ethical issues are specious. Cross was making the point that ethics is necessarily related to action, even if not immediately. The contrast just drawn between food preferences on the one hand, and preferences between football teams or cities or composers on the other, is deceptive. It is certainly true that the latter arouse stronger passions than the nectarine issue and hence seem more like the abortion issue. However, the differences of opinion, while passionate, do not lead to relevant action in relation to the opinions themselves. That is, the football fans might come to blows, but this is due to their intense attachment to the football teams, not, as

it is in the abortion case, to their views on the disputed difference (between the style of the football teams).

7.2.3 Social Psychology

Gilbert Harman[131] believes that the Aristotelian view, at least in the way I have characterized it, is empty. There are no character traits and hence the notion of the good person as a basis for deciding what is virtuous is untenable.

Harman believes that we attribute character traits to people and that this attribution, in parallel with much other folk wisdom, whether in physics or psychology, is mistaken. There are no relatively long-term stable dispositions to act in certain ways, i.e. to be kind or generous or happy or temperate or courageous.

He offers two major arguments for this seemingly counter intuitive view. First, he says that the mistaken view works quite well because we make it do so in a variety of ways. Criminals are commonly put in situations where criminal behaviour is seen by them to be advantageous. Likewise, clergymen are usually in situations where it is obvious that they are motivated to behave as they are expected to do. There are social pressures to be consistent, to conform, and, moreover, we, as observers of other people's behaviour, want to construct a consistent picture and therefore do so. This last is a theme, interestingly since she seems to be a virtue ethicist, of the major contemporary novelist and philosopher, the late Iris Murdoch, whose novels

[131] G. Harman, "Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error", in "Explaining Value", G. Harman, Oxford: Oxford University Press, 2000, Ch 10.

consistently depict the arbitrariness of human behaviour and the superficiality of our construct of the rational normality of our fellows.

His second and major argument comes from the results of experiments in social psychology. He cites the famous experiments by Milgram[132] in which subjects were asked to administer electrical shocks of increasing severity to putative "learners" in a supposed learning experiment. The widely known result, that a majority of ordinary decent people displayed the most shocking callousness and cruelty, is seen by many social psychologists as inexplicable on the basis of some assumed character traits of the ordinary decency of the subjects. The result is best accounted for by assessing the particular situation – its ambiguity, the gradual process of increased cruelty, the aura of authority of those conducting the experiment and so on. It is notable that most of us believe we would behave differently from the subjects, while not denying that the experimental subjects were "people like us". Harman cites a further set of experiments in which a subject could be a Good Samaritan while *en route* to an engagement that was assigned various degrees of urgency by the experimenters. The only correlation with those who stopped to help the sufferer, as against those who did not, was the stipulated urgency of the fairly mundane scheduled appointment of the subjects. Religious outlook did not matter, nor did it matter if the subject's assignation was to talk on the Good Samaritan parable!

---

[132] S. Milgram, "Behavioural Study of Obedience", Journal of Abnormal and Social Psychology, 67, 371 – 378 (1963)

Harman's view is perhaps supported by what occurred in Germany between 1933 and 1945. Millions of ordinary people acquiesced in the systematic extermination of Gypsies and Jews (Milgram relates his experiments to this fact). Likewise in Rwanda the massacre of over 500 000 Tutsi, on the grounds that they were Tutsi, seems hard to explain other than on the basis that situation, rather than fixed character, determines behaviour.

All of this is not to deny that people may have fixed dispositions, but Harman claims that the experimental evidence is that these dispositions are quite narrowly focused. They do not amount to what we normally see as character traits. For example, he cites evidence that talkativeness at lunch is a fairly stable behaviour pattern – hardly a character trait. Further, Harman does not deny that individuals might behave systematically in certain ways because of personality disorders or mental illness; again this is hardly what Aristotle meant when comparing learning virtue to learning to play the lyre.

> "If there is no such thing as character, then there is no such thing as character building".[133]

There are a number of possible responses to Harman's thesis. One is to deny that we can generalize from the small number of experiments that have been done thus far, in what Harman readily acknowledges is a difficult area. Another is to suggest that the categories of character traits being used are inappropriate. If there are some universal underlying character traits for all individuals, as there are supposed by some to be linguistic universals

---

[133] G. Harman, *ibid*, p177

underlying all the individual idiolects, perhaps we have yet to identify them. If they exist, they may well not be those chosen by Aristotle (or Harman).

The most relevant response for my present purpose is an attempt, from the standpoint of virtue ethics, to establish some substance in the Aristotelian idea of learning by doing in the domain of ethics. It is useful to proceed, as Hursthouse does in another context, from a vice rather than a virtue. Consider a child in South Africa in the 1980s of an Afrikaner family committed to the ideology of apartheid. The constant admonition and advice to fear, hate, suspect and despise Africans, and the constant giving of spurious reasons along the lines of the alleged dishonesty, depravity, idleness, incompetence and stupidity of Africans is, in the climate of a loving parental relationship, an emotional "education" of a powerful kind. There is good evidence of the success of this process in the commitment to racism of the children of racists. The success of the program is happily not universal, and also it is undeniably true that adults not thus wrongly "educated" are sometimes also racist. The point of interest here is that if the racist child becomes an enlightened, kindly adult, can she throw off the effects on her character (on the assumption *contra* Harman that such exists) of the training? It seems plausible that such an adult will have many immediate emotional responses which she will reject, and she will behave properly towards her African fellows, but she will do this by knowing-that. It will be a process of reflection and not the normal immediate responses of goodwill that she shows to the whites with whom she interacts. Hence, it can be argued that Aristotle is right; she cannot display true virtue, because she has been wrongly habituated. She is not

blameworthy but she is not truly virtuous. (It might be argued that Aristotle would allow the re-habituation of the ill brought up person. This seems doubtful, given his view that various categories of people are not able to become virtuous, and the strong sense he conveys of the deeply ingrained nature of the virtuous person. Of course, this does not commit the contemporary virtue ethicist to rejecting the idea of re-habituation. It seems to me that this is a matter of one's (unreliable) intuitions, and my own feeling is that it is unlikely. The immediate response of many of us, who are less than virtuous, to, for example, a hideously facially deformed person[134], and our subsequent proper behaviour of normal respect would be hard to unlearn. To replace this with the desired spontaneous warmth and instant behaviour of normal respect would not be easy.)

The argument about habituation of "bad" character has some force here, but it is not clear that a similar argument for positive traits can be mounted successfully against the evidence from social psychology presented by Harman.

This is an extremely difficult area for empirical research and maybe, anyway, to be a racist or not to be a racist is not a character trait. It may lie somewhere between the traditional, like being courageous, and the specific, like being talkative at lunch[135]. Where, for example, would being punctual lie? I think

---

[134] Perhaps this very concept would not occur to a virtuous person.

[135] The Aristotelian character traits are themselves in any case something of a mixed bag; gratitude seems slightly odd. Moreover constancy is not usual. Temperance is rarely a uniform property of an individual, a phlegmatic husband may be a fanatical football fan and an

Harman would agree that some people are always punctual and others almost universally not. In the ordinary current use of the term it is a character trait. On the other hand, it is not something that the virtue ethicist would regard as having the moral force that would make it a requirement for a virtuous person. Aristotle, one feels, would not have regarded it as a virtue; it is hardly noble, and besides a slave or a woman might well be punctual, and indeed might be forcibly obliged to be punctual.

Whether, then, virtue ethics is a defensible stance remains controversial. On the one hand, we have Harman's view that a "settled character" is a myth. On the other, we have the simple empirical fact[136] that there are individuals who, over a prolonged period and in many different circumstances, have behaved in a way that seems admirable, and others who have consistently behaved in a way it would commonly be agreed is indefensible. It seems to me that the most that one can say in defense of Harman's position is that the best of us may in certain circumstances behave badly. The question is then whether this is something that the virtue ethicist can concede. Does the virtue ethicist need, for her argument, the existence of someone with Kant's unattainable "holy will"?

---

irascible motorist may be a patient botanist. Likewise courage is often variable; a heroic suffragette may taunt the baton-wielding police and yet be afraid of spiders.

[136] Many of us can cite examples from personal acquaintance, but also there would be many who would accept such a contrast between, say, Archbishop Desmond Tutu and the numerous serial child rapists and murderers who prey on the Archbishop's compatriots.

If we accept virtue ethics as a doctrine, we are making the claim that the human agent must take moral actions that are based on grounds that are neither deductive nor inductive.

7.3 Virtue Ethics and the Basis of Moral Judgement

For the virtue ethicist, as I remarked in section 7.1, a utilitarian calculation of moral rightness is impossible. It is an essential part of moral concerns that there should be competing and inconsistent contextually determined factors, which cannot be given a weighting.

In many cases of a non-moral kind we can construct a major premise, a minor premise and the conclusion of a conventional syllogism as follows.

1. Any natural number that is not divisible by another natural number (excluding unity) is prime.

2. The natural number 9699691 is not divisible by a natural number other than itself or unity.

3. Therefore 9699691 is prime.

Compare this with

(1a) Telling lies is to behave wrongly.

(2a) I deliberately told my elderly mother falsely, when she asked, that I liked the hideous new hat she proudly showed me. That is, I told a lie to her.

(3a) I told a lie therefore I behaved wrongly.

In this latter case, where my lie is motivated by kindness and affection, unlike the first syllogism, we do not feel that the argument is valid because the major premise holds only for the most part. It is this contention that is an essential

element in Aristotle's ethics. Moreover, even in the following case we do not see an acceptable parallel with the first case.

> (1b) Consideration and kindness towards friends is right.
>
> (2b) My friend needs my support tonight as she is feeling depressed and providing support would be considerate and kind.
>
> (3b) I will not go to the exciting party I planned for tonight but will stay with my friend because that would be considerate and kind.

This does explain my motivation here, but not why the practice is justified, why it is a virtuous action. It just subsumes particulars under general cases and, as the previous example illustrates, we need further argument that warrants our action. We need to identify certain properties as morally salient - to rank concerns and show which are most central. The best principles about how one should behave hold only for the most part, and indeed if the party in this last example were my wedding party, I would have to neglect my friend that evening.

These differences between 1-3 on the one hand, and 1a-3a, 1b-3b on the other, stem from two sources: the inevitable complexity of particular moral cases, and a certain holism. Even in the most basic moral situation one needs to judge what is relevant and not relevant to the application of the rule. But each situation is unique as the properties come in a unique combination. This holistic interaction of properties in the particular case affects the relevance of the properties cited in the general premise. General principles only pick out properties atomistically, and thus, whilst the general principle may be true, more needs to be said to apply it to the holistic context of complex actual moral situations.

It might be argued that I have chosen a false contrast between 1-3 and the other examples, as I have not dealt in the former with practical reason. However, while this example was chosen for emphasis, the argument readily carries over to the more probabilistic syllogisms common in practical reason. In the non-moral case, I weigh the probabilities of various factors in placing my bet say on Tottenham Hotspur winning the 2004 Premier League. The likelihood of injury to key players, the possible purchases of players, the manager's skills and so on are all factors to which I give an appropriate weight. Once more, this contrasts with the moral case where some factors are given zero weight despite their possible role. As Wiggins[137] points out, in the moral case there is a plurality of human ends and goods. Practical reason, in my case of placing a wager, is a rational judgement of probability. This contrasts with a moral action based on the salience of the many factors impinging on the virtuous agent, who can make no calculation of the ideal moral value, since what in one set of circumstances may be salient, in another may not, because of the complex reciprocal interaction of the surrounding circumstances.

McDowell[138] makes the same point by pointing out that, unlike the utilitarian, the virtue-ethicist does not act as a result of balancing reasons for and against. She undertakes the virtuous action by exercising her sensitivity, and it may be, as mentioned above, that a particular seeming consideration is

---

[137] D. Wiggins, "Deliberation and Practical Reason", in "Essays on Aristotle's Ethics", Ed A.O. Rorty, California: University of California Press, 1980, p 233.

[138] J. McDowell, "Virtue and Reason" reprinted in "Mind, Value, and Reality", Cambridge: Harvard University Press, 1998.

completely silenced. For example, if she must fight in defense of her family, the very real danger ceases to have any salience whatsoever. This is not to say that the danger is not apprehended, otherwise she would merely be behaving rashly. As McDowell sees it, we cannot write down what might constitute the major premise of a putative practical syllogism for moral action. To do so would require that we recapitulate within it the whole of the moral training that Aristotle sees as the basis for virtue.

If we accept virtue ethics, then virtue, like playing the lyre, is knowing-how. There are, however, significant arguments that give rise to legitimate doubts about the antecedent of this proposition.

In summary, then, it has been claimed in this chapter that if one accepts virtue ethics, moral judgements exemplify knowing-how on any of the various positions one can adopt in the area of philosophy of mind (consistent with virtue ethics). This does not, of course, exclude a component of knowing-that. If one does not accept virtue ethics, it still may be true that moral judgements exemplify knowing-how, but it is not clear-cut. In general, utilitarian stances would claim that moral judgements were knowing-that.

Korsgaard[139], a Kantian ethicist, puts the issues very clearly. First, from the Aristotelian perspective,

> "In Greek thought, becoming excellent is as natural as growing up. We need to learn virtue; but it is as we learn language, because we are human and that is our nature."

---

[139] C.M. Korsgaard, "The Sources of Normativity", Cambridge: Cambridge University Press, 1996, p 3

If the ancient Greek view of morality were right, then virtue would, as I have described, be knowing-how, at least in the case of a well-bred and well brought up citizen. (Incidentally, in Chapter 5 I argue that language is knowing-how because learning language is our nature.) However, as Korsgaard then remarks, many now see virtue very differently - because of our autonomy, a common modern view of the world sees ethics in terms of obligation imposed by the agent's own will. For Kant, an unreflective "knowing-how" kind of good act falls short of true virtue. True virtue must be, indeed is, reflective.

CHAPTER 8. CONFRONTING THE OBJECTIONS TO RYLE'S KNOWING-HOW

8.1 Introduction

I concentrate here on the paper by Stanley and Williamson[140] that I have already referred to in a number of instances. It is the most serious recent challenge to Gilbert Ryle's views. Rather than suggesting that knowing-how is mere ability, these writers claim, in an interesting and substantial account of Ryle's position, that knowing-how can be subsumed under knowing-that.

I want first to deal with a minor but nevertheless important source of confusion. This confusion, as I have remarked previously, is to take the common usage, as in "my small daughter knows how to multiply up to the seven times table", to indicate an instance of Rylean knowing-how. In fact, this skill is based on knowing-that. My daughter knows 1 x 2 = 2, 2 x 2 = 4, 3 x 2 =6 ..... 12 x 2 = 24 .....12 x 7 = 84. This is an explicit finite set of propositions that she has memorized. Stanley and Williamson confusingly write[141] about elephants knowing-how to ford a river, when their context makes it clear that the elephants know that to ford a river they go to a place where it is comparatively shallow. To quote, "They go to the only ford and walk across." Stanley and Williamson then make my point - that this is ordinary propositional knowledge - very clear. They remark that if the river

---

[140] J. Stanley and T. Williamson, "Knowing How", Journal of Philosophy, 98, 411 – 444 (2001)

[141] J. Stanley and T. Williamson, *ibid*, p 439

were to be dredged, unbeknown to the elephants, the elephants would then have a false belief[142].

I have also already referred in section 2.4 to the claim that one can know-how, in Ryle's sense, to do something but be unable to do it. As I said there, where it was immediately relevant to the defining of knowing-how, this is not Ryle's knowing-how. Stanley and Williamson write[143] of a ski instructor who knows-how to perform a complex manoeuvre, but cannot do it. In fact, all she can do is to describe in propositional terms what the manoeuvre consists in; she does not know-how.

8.2 Stanley and Williamson's Linguistic Argument Against Knowing-how

Stanley and Williamson do not share the view of some opponents of Ryle; namely, that bicycle riding is an ability rather than knowledge. Their view is that bicycle riding, skiing and the like are examples of knowing-that.[144]

I earlier quoted Fodor's remark[145],

---

[142] The pros and cons of this argument do not require that we go into the vexed question of whether animals have knowledge, as the example could as well have been about the local villagers.

[143] J. Stanley and T. Williamson, *ibid*, p 416.

[144] It is worth noting that once the argument that knowing-how is merely ability is abandoned, as it is by Stanley and Williamson, an argument of Ryle's, distinct from the regress argument, gains some force. Ryle says, "Efficient practice precedes the theory of it; methodologies presuppose the application of the methods, of the critical investigation of which they are the products." (G. Ryle, "The Concept of Mind", *ibid*, p 30) In other words, we can imagine that nobody had ever skied. Someone then attached some smooth flat boards to her feet and by trial and error learned to go rapidly down the snow covered mountain in the absence of any propositional knowledge. Of course, Stanley and Williamson might respond that she has acquired propositional knowledge in stages by successive experiments, and that is possible. But for Stanley and Williamson to be vindicated, they must claim that it is necessary to obtain propositional knowledge by experiment, not merely that it can and does happen, but that it always happens in such circumstances. This is not a credible proposition, since in general one does not have ordinary propositional knowledge of, for example, how to ride a bicycle.

> "If I know how to do x, it does not follow that I know how to explain how x is done. There is at least that much distinction between knowing-how and knowing that".

While I am claiming that knowing-how is a subset of knowledge and also a subset of ability, Stanley and Williamson claim that knowing-how is a subset of knowing-that. In denying the existence of the difference described by Fodor, they appear therefore be saying either: (i) knowing-how involves being able to explain; or (ii) that this is not a distinction which prevents knowing-how, which one may not be able to explain, being a subset of knowing-that. In other words, (ii) above would be the claim, *inter alia*, that one can have propositional knowledge that one cannot articulate.

As I have just discussed, their remark about the ski instructor who knows-how to perform a complex stunt but cannot do it (this is claimed in circumstances where there is no physical impediment to performance) might suggest (i) above. They take it to be possible to know-how without being able to perform, but simply being able to explain. This is not what Ryle is writing about. This interpretation of their stance seems to be exemplified in the claim that "Hannah knows how she could ride a bicycle" is equivalent to "Hannah knows how to ride a bicycle"[146]. Compare "I know how I could beat Sampras at tennis" to "I know how to beat Sampras at tennis"; the first is easy - improve

---

[145] J.A. Fodor, "The Appeal to Tacit Knowledge in Psychological Explanation" Journal of Philosophy, 65 (1968) p 634

[146] J. Stanley and T. Williamson, *ibid*, p 425

the speed and accuracy of my serve twenty fold and so on - the second is extraordinarily difficult.

Elsewhere, however, their use of the examples of pianist and the cyclist contradicts this interpretation, and instead suggests that they subscribe to (ii) above; namely, that one can have propositional knowledge that one cannot articulate. Indeed, they say that knowledge of the proposition "Hannah knows how to ride a bicycle" is ascribed to Hannah under "what we call a *practical* mode of presentation"[147] (authors' italics). This is justified as being propositional knowledge by a grammatical parallelism between the knowing-how case and the identity of the following two statements[148]:

"John believes that that man has burning pants" and,

"John believes that he himself has burning pants."

under a particular set of circumstances where "that man" refers to John. The identity is vindicated by appeal to the relevant mode of presentation in each case.

I have not done justice here to their complex argument based on their application of Kartutunen's account[149] of the semantics of questions. However, it seems to me that they have extrapolated unduly from a linguistic theory that is nowhere claimed to be complete, precise and without

---

[147] J. Stanley and T. Williamson, *ibid*, p 429

[148] J. Stanley and T. Williamson, *ibid*, p 428

[149] L. Kartunnen, "Syntax and Semantics of Questions", Linguistics and Philosophy, 1, 3 – 44 (1977)

exception[150]. It is as though this theory resembles in precision the contemporary theory of classical mechanics in the domain of velocities that are small compared with the velocity of light. In reality, linguistics is more nearly in the state of the theory of light at the time of Huygens and Newton in the 17th century. As in most issues of language, there are some very curious uses of what is phrased like an embedded question; consider Fred leaving the cinema in disgust, saying, "I'm leaving, who would stay for a picture like that."

While it is true that there is no syntactic difference between "Hannah knows how Bill rides a bicycle" and "Hannah knows how to ride a bicycle", a proponent of a Rylean knowing-how would be reluctant to suggest that because the latter is an example, so is the former. Stanley and Williamson appear to argue[151] that because the former is clearly knowing-that, so is the latter.

To fully rebut Stanley and Williamson, it would be ideal to give a detailed alternative account of their treatment of embedded questions, and I have not succeeded in this difficult task. Alternatively, a more specific questioning of the precise structure of their extensive exposition would be an advance on the rather general worries I have expressed.

---

[150] It would thus seem desirable to test the argument for a number of languages, each from a family different from that of English, say, Japanese, Welsh and Swahili. The legitimacy of this demand depends on the usage of the capacity sense of know in many languages. Edward Craig claims to have studied this and affirms that it is indeed the case (E. Craig, "Knowledge and the State of Nature", Oxford: Oxford University Press, 1990, p 152).

[151] J. Stanley and T. Williamson, *ibid*, p 419

Koethe[152] has given a clear and succinct account of the Stanley and Williamson argument, and shown it to be questionable. His brief paper is difficult to summarize and I will give a fairly full account, as the issue is quite central to my purpose.

Stanley and Williamson consider a case where Hannah does not know how to ride a bicycle, but then consider Hannah in relation to Bill (as cited above), where she knows from Bill that that is the way to ride a bicycle, ascribed to her under "a practical mode of presentation"

As Koethe says[153], this suggests

> "what Stanley and Williamson's account of what it is for someone to know how to F really comes to, with its talk of entertaining a proposition under a practical mode of presentation, is that someone knows how to F just in case he knows, of some way w, that w is a way for him to F, and he knows how to instantiate w himself (and thereby F)".

This seems quite correct. Because if we return to the point I made earlier and assume, as Stanley and Williamson appear to assume, that someone can know of some way w for her to F, that w is a way for her to F, but not know how to instantiate w herself, then she would know how to F without being able to F, despite the fact that there is no disability preventing her F-ing. As Koethe points out, Stanley and Williamson[154] reject as "absurd" the suggestion that

---

[152] J. Koethe, Stanley and Williamson on Knowing-how", Journal of Philosophy, 99, 325 – 328 (2002)

[153] J. Koethe, *ibid*, p 326

[154] J. Stanley and T. Williamson, *ibid*, p 443

Mary (Jackson's[155] imaginary colour scientist) knows how to imagine an experience of red and yet is unable to imagine an experience of red.

Koethe[156] then goes on from this to say

> "It appears, then, that entertaining the proposition that w is a way for one to F under a practical mode of presentation involves knowing how to instantiate w one's self. But if so, then Stanley and Williamson's account of knowing-how appeals to the very notion it seeks to explicate."

This seems right. Then Koethe continues[157], from Stanley and Williamson,

> "One must know how to F in order to F. If knowing-how to F is to know, of some way w, that w is a way for one to F and to know how to instantiate w one's self, then knowing-how to instantiate w one's self requires knowing, of some way w', that w' is a way to instantiate w one's self and knowing how to instantiate w' one's self; or it does not."

Koethe points out that if it does, we are led to exactly the fatal regress that Ryle was addressing, whereas if it does not we have an instance of knowing–how which is not an instance of knowing-that, *contra* Stanley and Williamson.

For a thoroughgoing externalist such as Williamson, there is an entirely different argument for a distinct knowing–how that might have some appeal. One can make a case that knowing-how supervenes on the (embodied) brain always, whereas knowing-that does not always do so. This would make a distinction between knowing-how and knowing-that, while allowing knowing-how to be a subset of knowledge. The idea has some attractions. In particular,

---

[155] F. Jackson, "Epiphenomenal Qualia", Philosophical Quarterly, 32, 127 – 136 (1982); "What Mary didn't know", Journal of Philosophy, 83, 291 – 295 (1986)

[156] J. Koethe, ibid, p 327

[157] J. Koethe, *ibid*, p 328

it allows that my knowing-how (if I did) to return the Sampras serve incorporates some of the not conscious, motor intentional, perceptual inputs and outputs that we know exist. It also accommodates the benefit of practice in improving knowing-how, and it very naturally accounts for the failure to manifest knowing-how in the case of particular brain injuries. Moreover, the externalist arguments that are given for the claim that knowing-that does not supervene on the brain do not apply to knowing-how because there is no analogue of false belief.

8.3 Qualia and Stanley and Williamson

How can we evaluate the Stanley and Williamson argument that knowing-how is a subset of knowing-that, in the context of qualia? Even if it is not accepted that Koethe has a knockdown argument against their claim, it seems impossible to see how the knowing-how of qualia can be ascribed under "a practical mode of presentation". The very essence of qualia is that this cannot be done[158].

Consider first, for simplicity, some qualia which fall under ability and not knowing-how. We can take a simple experiment where a subject is looking through spectacles, green over one eye and red over the other. If she now looks at an image consisting of a house in green and a face in red, the image

---

[158] It should be said, in fairness to Stanley and Williamson, that they say quite explicitly that "practical modes of presentation are not essential to our analysis" (J. Stanley and T. Williamson, *ibid*, p 434, footnote 38). They say that a philosopher who rejects the semantic relevance of modes of presentation may easily accept their analysis. In fact, it seems to me, that Stanley and Williamson's best defence against the argument I have presented in relation to qualia is to say that for Annabelle to know "what it is like" to experience various moments during a Mozart symphony is neither knowing-how nor knowing-that. In other words, they would reject my definition of knowing-how.

alternates in time for the subject between a house and a face, and she has no control over this. There is no change in the external stimulus. One could hardly say, "That is the way w to feel what it is like to see a picture of a beautiful Georgian house on an English summer day", when, as often as not, this is not the case.

As another example, we can consider the taste of phenylthiocarbamide. By selective breeding, it is in principle possible to decrease the number of people who can taste PTC. Finally, we can imagine there is only one individual, Maria, who can do this[159]. She tastes it and tries to describe this unique taste to others. She can say it tastes bitter, perhaps that it tastes rather like aloes, but she alone knows what it is like to taste PTC. Of course this once more is an ability, and not an ability of the knowing-how subset, since no learning process is involved. Nevertheless, as in the previous example, there is no way w for anyone to be informed about the ability.

The only proposition about qualia that one could convey to Jackson's Mary is something along the lines of, "Seeing a red fire engine is like seeing a black and white fire engine, and imagining it impinging on your retina as though it

---

[159] There are grounds for caution about such thought experiments, so it is worth considering briefly the conditions of utility and applicability of thought experiments. (There is a good account of the grounds for caution in K.V. Wilkes, "Real People", Oxford: Oxford University Press, 1988.) This is a large topic but three things stand out. First, thought experiments that are physically possible in the real world, even if impractical on logistical grounds, generally have *prima facie* greater merit than those that are not. Second, bizarre thought experiments relying on intuitions vastly outside our normal experience are extremely suspect – teleporting our brain content and the like. Third, thought experiments are most useful for clarifying the meaning of our concepts. My Maria case looks satisfactory on these counts. We certainly need nothing more than a unique sensory experience, which is only a small extrapolation from a rare one (compare the frictionless plane with the very smooth plane). The concept whose applicability we are examining is that of being ascribable under a "practical mode of presentation".

was reflecting light of a wavelength of about 75 millionths of a centimeter, and this makes it more dramatic and striking and gives you all sorts of associations with your childhood ambition to be a fireman". However, of course, Mary already knows this. Mary already knows everything that can be said, and there is no practical mode of presentation.

Stanley and Williamson do not claim that ability is propositional knowledge and I have been discussing qualia relating only to ability, not knowing-how. However, it seems to me that the ideas that are clear in the above examples translate quite readily into the more sophisticated learned appreciation of qualia that I have characterized as knowing-how. Without my conceding their argument against knowing-how in these cases, Stanley and Williamson can claim justly that for cycling or skiing or playing tennis, a demonstration can be given. However, it is not possible to give a demonstration of the way w to be moved by a Mozart symphony, because the same individual can be moved or not under identical external circumstances and can appear the same in the two instances. Stanley and Williamson's practical mode of presentation faces similar difficulties to Skinner's behaviourism, which was demolished by Chomsky in 1959[160]. One of the reasons is in essence the same; there need be no behavioural manifestation of different qualia.

One of the problems is that Stanley and Williamson have perhaps examined too narrow a view of knowing-how; riding bicycles, skiing and the like.

---

[160] N. Chomsky, "A review of B.F. Skinner's Verbal Behaviourism", Language 35: 26—58 (1959)

Consider knowing-how to keep a secret[161]. This cannot be characterized as a series of propositions (i) always remaining sober (ii) being generally distrustful (iii) being generally taciturn (iv) not becoming crazily infatuated with any man or woman (v) being incorruptible (vi) being stoic under torture, and so on. It is a matter of temperament and some learned wisdom. There is no way w of giving a practical demonstration of the art unless one accepts that a way w can include years of performance under observation.

[161] This is discussed as not being an action but a course of action in G. Ryle, "Courses of Action or the Uncatchableness of Mental Acts", Philosophy, 75, 331 – 344 (2000)

CHAPTER 9. CONCLUSION

I have argued that Ryle was right in identifying knowing-how as distinct from knowing-that. Further, I have argued that knowing-how is a subset both of ability and of knowledge. I do not attempt to summarize these lengthy arguments here but comment very briefly on the consequences.

There is a great deal in Ryle's well-known book, "The Concept of Mind", that is not touched on in this thesis. However, there is very little in this thesis that is not at least implicit in Gilbert Ryle's masterly account of the inadequacy of the view that all knowledge is propositional knowledge. Nevertheless, the intervening half-century has provided a basis for reconsideration. As Dennett[162] has pointed out, Ryle's questions are at the personal level of explanation and not at the sub-personal level. They are about what people do and not about what people's brains allow them to do.

There are several additional reasons for the present lengthy exegesis of Ryle's concept of knowing-how.

First, there is clear evidence in the philosophical literature that the concept is by no means universally accepted.

Second, while it is true that a close reading of Ryle shows that he had a comprehensive notion of knowing-how, he does not spell that notion out in any detail. This has allowed a great deal of confusion, demonstrated for

---

[162] D.C. Dennett, Introduction to G. Ryle, "The Concept of Mind", Harmondsworth: Penguin Books, 2000

example in the use of the term "knowing-how" when purportedly referring to Ryle's concept, for what are patently examples of propositional knowledge.

Third, it seems desirable for theories of knowledge either to encompass something beyond propositional knowledge or at least to acknowledge that there is more to knowledge than propositional knowledge. If this is the case, a more precisely spelled out account of knowing-how is a necessary preliminary for a comprehensive understanding of knowledge.

Knowing-how is not justified true belief or any of its near relatives, but it is knowledge. Individuals with severe visual agnosia know-how to do things and, contrary to my firm view, it is perhaps believable that this can be indiscriminately classified with ability. However, a learned acquisition of mathematical nous or a learned responsiveness to complex music is something distinct from feats like the ability to stand on one leg.

It seems plausible that knowledge is a cluster concept like Wittgenstein's family of games. The various sorts of knowledge do not share any one property or set of properties, but we can recognize knowledge. We can see the family relation between unconscious knowledge of our Oedipal desires and of grammaticality, of conscious knowledge of the atomic number of gold and of *modus ponens*. Given this, we need a theory for each type of knowledge that can properly be specified and distinguished.

Since Ryle's work, a number of philosophers like Merleau-Ponty[163] and Polanyi[164] have argued that knowing-how is the basic form of knowledge. This view receives substantial support from the neuro-physiological work of the last thirty years. It seems unlikely that a satisfactory theory of propositional knowledge will be possible without a better theoretical understanding of knowing-how. Whether such an improved theory will arise from closer philosophical theoretical analysis or from further experimental data on the brain is an open question at present.

---

[163] M. Merleau-Ponty, "The Phenomenology of Perception", London: Routledge, 1962

[164] M. Polanyi, "Personal Knowledge", New York: Harper and Row, 1964; "Knowing and Being", Ed M. Grene, Chicago: University of Chicago Press, 1969

BIBLIOGRAPHY

T. Alter, "Know-How, Ability, and the Ability Hypothesis", Theoria, 67, 229 – 239 (2001)

Aristotle, "Nicomachean Ethics", Translated, T. Irwin, Indianapolis, IN: Hackett Publishing Company, 1985

S. Baron-Cohen, T. Joliffe, C. Mortimore and M. Robertson, "Another advanced test of theory of mind: evidence from very high functioning adults with autism or Asberger Syndrome", Journal of Child Psychology and Psychiatry, 38, 813 – 822 (1997)

N. Block, "Inverted Earth", in "The Nature of Consciousness", Ed N. Bock, O. Flanagan, and G. Guzeldere, Cambridge MA: MIT Press, 1997, Chapter 43

G. Botterill and P. Carruthers, "The Philosophy of Psychology", Cambridge: Cambridge University Press, 1999

D. Braddon-Mitchell and F. Jackson, "The Philosophy of Mind and Cognition", Oxford: Blackwell, 1996

M.F. Burnyeat, "Aristotle on Learning to Be Good", in Essays on Aristotle's Ethics", Ed A.O. Rorty, California: University of California Press, 1980, p 69 - 92

D. Carr, "Knowledge in Practice", American Philosophical Quarterly, 18, 53 – 57 (1981)

D.J. Chalmers, "The Conscious Mind", Oxford: Oxford University Press, 1996

W. Child, "Causality, Interpretation, and the Mind", Oxford: Oxford University Press, 1994

N. Chomsky, "A review of B.F. Skinner's Verbal Behaviourism", Language 35, 26 - 58 (1959)

N. Chomsky, "Reflections on Language", New York: Pantheon, 1975

N. Chomsky, "Rules and Representations", New York: Columbia University Press, 1980

N. Chomsky, "Knowing Language: Its Nature Origin and Use", New York: Praeger, 1986

N. Chomsky, "Language and Problems of Knowledge", Chomsky Symposium, Madrid, 1986, cited in ""A Review of D'Agostino, "Chomsky's System of Ideas", "Noam Chomsky Critical Assessments", Ed C.P. Otero, London: Routledge and Kegan Paul, 1994

E. Craig, "Knowledge and the State of Nature", Oxford: Oxford University Press, 1990

R.C. Cross, Philosophy, 25, 301 – 315 (1950)

D. Davidson, "Essays on Actions and Events", Oxford: Oxford University Press, 1980

D. Davidson, "Knowing One's Own Mind", reprinted in "Mind and Cognition, Ed G. Lycan, Oxford: Blackwell, 1999

M. Davies, "Tacit knowledge and the structure of thought and language", in "Meaning and Interpretation" Ed C. Travis, Oxford: Blackwell, 1986, Chapter 5

M. Davies, "Tacit Knowledge and Subdoxastic States" in "Reflections on Chomsky", Ed A. George, Oxford: Blackwell, 1989

M. Davies and T. Stone (Editors), "Folk Psychology", Oxford: Blackwell, 1995

A.D. de Groot, "Thought and Choice in Chess", The Hague: Mouton Publishers, 1978

S. Dehaene, "The Cognitive Neuroscience of Consciousness", Amsterdam: Elsevier, 2001

D.C. Dennett, "Styles of Mental Representation", Proc Aristotelian Soc, 83, 213 – 226 (1983)

D.C. Dennett, "Consciousness Explained", Boston: Little, Brown, 1991

D.C. Dennett, "The Path Not Taken", in "The Nature of Consciousness", Ed N. Block, O. Flanagan and G. Guzeldere, Cambridge: MIT Press, 1999, p 417 – 419

D.C. Dennett, Introduction to G. Ryle, "The Concept of Mind", Harmondsworth: Penguin Books, 2000

M. Devitt and K. Sterelny, "Language and Reality", Cambridge: MIT Press, 1999

J. Dewey, "Human Nature and Conduct: An Introduction to Social Psychology", New York: Random House, 1957

W. Enard, M. Przeworski, S.E. Fisher, C.S. Lai, V. Wiebe, T. Kitano, A.P. Monaco and S. Paabo, "Molecular evolution of FOXP2, a gene involved in speech and language", Nature, 418, 869 - 72 (2002)

G. Evans, "Semantic theory and tacit knowledge", in G. Evans "Collected Papers", Oxford: Oxford University Press, 1985

G. Evans, "The Varieties of Reference", Oxford: Oxford University Press, 1982

J.A. Fodor, "The Appeal to Tacit Knowledge in Psychological Explanation", Journal of Philosophy, 65, 627 - 640 (1968)

J.A. Fodor, "The Language of Thought", Hassocks, Harvester Press, 1975

J.A. Fodor, "the mind doesn't work that way", Cambridge: MIT Press, 2000

C.R. Gallistel, A.L. Brown, S. Carey, R. Gelman and F.C. Keil, "Lessons From Animal Learning for the Study of Cognitive Development", in "The Epigenesis of Mind", S. Carey and R. Gelman (Eds), Hillsdale, NJ: Lawrence Erlbaum Associates, 1991

C.R. Gallistel, "The Replacement of General Purpose Learning Modules with Adaptively Specialized Learning Modules" in "The New Cognitive Neurosciences", M.S Gazzaniga (Ed), 2nd Edition, Cambridge: MIT Press, 2000

V. Gallese, L. Fadiga, L. Fogassi, G. Lupino and A. Murata, "A parietal-frontal circuit for hand grasping movements in the monkey: evidence from reversible inactivation experiments", in P. Their and H.O. Karnath (Eds), "Parietal lobe contributions to orientation in 3-D-space", Berlin: Springer-Verlag, 1997

E. Gettier, "Is Justified True Belief Knowledge?" Analysis, 25, 121 – 123 (1963)

A.I. Goldman, "Philosophical Applications of Cognitive Science", Boulder: Westview Press, 1993

G.H. Harman, "Psychological aspects of the theory of syntax", The Journal of Philosophy, 64, 75 - 87 (1967)

G. H. Harman, "Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error", in "Explaining Value", Ed G.H. Harman, Oxford: Oxford University Press, 2000

K. Hawley, "Success and Knowledge-How", American Philosophical Quarterly, 40, 19 – 31 (2003)

M. Heidegger, "Being and Time", trans J. Macquarie and E. Robinson, New York: Harper and Rowe, 1962

D. Hume, "A Treatise of Human Nature", 2nd edn, L.A. Selby-Bigge and P.H. Nidditch (eds), P.H. Nidditch, (rev), 1978

R. Hursthouse, "On Virtue Ethics", Oxford: Oxford University Press, 1999

F. Jackson, "Epiphenomenal Qualia", Philosophical Quarterly, 32, 127 – 136 (1982

F. Jackson, "What Mary didn't know", Journal of Philosophy, 83, 291 – 295 (1986)

M. Jeannerod, "The Cognitive Neuroscience of Action", Oxford: Blackwell, 1997

L. Kartunnen, "Syntax and Semantics of Questions", Linguistics and Philosophy, 1, 3 – 44 (1977)

S.D. Kelly, "Merleau-Ponty on the Body", Ratio, 15, 376 – 391 (2002)

H. Kluver and P.C. Bucy, "An analysis of certain effects of bilateral temporal lobectomy in the rhesus monkey, with special reference to psychic blindness", J. Psychol. 5, 33 – 54 (1938)

J. Knowles, "Knowledge of grammar as a propositional attitude", Philosophical Psychology, 13, 325 – 353 (2000)

J. Koethe, "Stanley and Williamson on Knowing-how", Journal of Philosophy, 99, 325 – 328 (2002)

H. Kornblith, "Knowledge in Humans and Other Animals" in "Philosophical Perspectives 13 Epistemology", Ed J. E. Tomberlin, Oxford: Blackwell, 1999, p 327 - 346"

H. Kornblith, "Knowledge and Its Place in Nature", Oxford: Oxford University Press, 2002

C.M. Korsgaard, "The Sources of Normativity", Cambridge: Cambridge University Press, 1996

S. Kripke, "Wittgenstein on Rules and Private Language", Oxford: Blackwell, 1982.

G. Lang, "The Rule-Following Considerations and Metaethics: Some False Moves", European Journal of Philosophy, 9, 190 – 209 (2001)

R. Larson and G. Segal, "Knowledge of Meaning", Cambridge: MIT Press, 1995

M. Luntley, "Contemporary Philosophy of Thought", Oxford: Blackwell, 1999

D. Marr, "Vision", San Francisco: Freeman, 1982

C. B. Martin and M. Deutscher, "Remembering", The Philosophical Review, 75, 161 – 196 (1966)

J. McDowell, "Mind, Value, and Reality", Cambridge: Harvard University Press, 1998

M. Merleau-Ponty, "The Phenomenology of Perception", London: Routledge, 1962

S. Milgram, "Behavioural Study of Obedience", Journal of Abnormal and Social Psychology, 67, 371 – 378 (1963)

A.D. Milner and M.A. Goodale, "The Visual Brain in Action", Oxford: Oxford University Press, 1995

S. Mineka and R. Keir, "The effects of flooding on reducing snake fear in rhesus monkeys: 6 month follow up and further flooding", Behavioural Research and Therapy, 21, 527 – 535 (1983)

M. Minsky, Behavioural and Brain Sciences, 3, 439 – 440 (1980)

T. Nagel, "Searle: Why We Are Not Computers", in "Other Minds", Oxford: Oxford University Press, 1995

A. Noë, "Perception, action, and nonconceptual content", a comment on S.L. Hurley, "Consciousness in Action", Cambridge: Harvard University Press, 1998 (http://host.uniroma3.it/progetti/kant/field/hurleysymp_noe.htm)

P. Noordhof, Australasian Journal of Philosophy, 81, 21 – 40 (2003)

R. Nozick, "Philosophical Explanations", Oxford: Oxford University Press, 1981

D. Papineau, "Thinking about Consciousness", Oxford: Oxford University Press, 2002

C. Peacocke, "Explanation in computational psychology: language, perception and level 1.5", Mind and Language, 1, 101 –123 (1986)

H. Pietersma (Ed) "Merleau-Ponty: Critical Essays", Lanham MD: University Press of America, 1990

M. Polanyi, "Personal Knowledge", New York: Harper and Row, 1964

M. Polanyi, "Knowing and Being", Ed M. Grene, Chicago: University of Chicago Press, 1969

G. Rattan, "Tacit knowledge of grammar: a reply to Knowles", Philosophical Psychology, 15, 135 – 154 (2002)

A.S. Reber, "Implicit Knowing and Tacit Knowledge", Oxford: Oxford University Press, 1993

G. Ryle, "The Concept of Mind", London: Hutchinson, 1949

G. Ryle, "Knowing How and Knowing That", Collected Papers, Volume 2, London: Hutchinson, 1971

G. Ryle, "On Thinking", Oxford: Blackwell, 1979

G. Ryle, "Courses of Action or the Uncatchableness of Mental Acts", Philosophy, 75, 331 – 344 (2000)

I. Scheffler, "Conditions of Knowledge", Chicago: University of Chicago Press, 1965

W.S. Sellars, "Empiricism and the Philosophy of Mind", in "Science Perception and Reality", London: Routledge and Kegan Paul, 1963

J. Stanley and T. Williamson, "Knowing How", Journal of Philosophy, 98, 411 – 444 (2001)

S.P. Stich, "Beliefs and Subdoxastic States", Philosophy of Science, 45, 499 – 518 (1978)

S.P. Stich and S. Nichols, "Folk Psychology: Simulation or Tacit Theory", in "Folk Psychology", Ed M. Davies and T. Stone, Oxford: Blackwell, 1995

C. Taylor, "Embodied Agency" in "Merleau-Ponty: Critical Essays", Ed H. Pietersma, Lanham MD: University Press of America, 1990

C. Taylor, "McDowell on Value and Knowledge", The Philosophical Quarterly, 50, 242 – 248 (2000)

G. B. Thomas, "Knowing How: A Study of the Concept of a Human Ability", PhD Thesis, Harvard University 1963

A.M. Turing, "Computing Machinery and Intelligence", Mind, 49, 433 – 460 (1950) p 451

L.G. Ungerleider and M. Mishkin, "Two cortical visual systems", in "Analysis of visual behaviour", D.J. Ingle, M.A. Goodale and R.J.W. Mansfield (Eds), Cambridge: MIT Press, 1982

T. Van Gelder, "What Might Cognition Be, If Not Computation", in "Mind and Cognition", W.G. Lycan (Ed), Oxford: Blackwell, 1999, p 170 - 191

F.J. Varela, "Ethical Know-How", California: Stanford University Press, 1999

M. Welbourne, "Knowledge", Bucks: Acumen Publishing Limited, 2001

L. Weiskrantz, "Outlooks for blindsight: explicit methodologies for implicit processes", Proc Roy Soc, B239, 247 – 278 (1990)

A.N. Whitehead, "Science and the Modern World", New York: New American Library, 1925

D. Wiggins, "Deliberation and Practical Reason", in "Essays on Aristotle's Ethics", Ed A.O. Rorty, California: University of California Press, 1980

K.V. Wilkes, "Real People", Oxford: Oxford University Press, 1988

L. Wittgenstein, "Philosophical Investigations", Oxford: Blackwell, 1953

L. Wittgenstein, "The Blue and Brown Book", Oxford: Blackwell, 1958

L. Wittgenstein, "Zettel", Oxford: Blackwell, 1967

APPENDIX A. LINGUISTIC PRESUPPOSITIONS

Throughout Chapters 4 and 5, I have implicitly taken it for granted that the account of language given by Chomsky is essentially correct. It is a natural question to ask how much of my claim for linguistic ability as knowing-how depends on this assumption.

There is considerable dispute about what is often referred to as Chomsky's innateness hypothesis (an expression Chomsky rejects because he believes it commonly encourages confusion about his naturalistic approach by introducing essentially irrelevant epistemological issues). This is the view that a child is born with latent knowledge of certain facts about the constraints on speaking a human language. There is a vast literature on this topic. There are questions about what is meant by this claim, and in precisely what domain Chomsky makes this claim and so on. It is neither necessary nor appropriate to enter fully into this debate here, although one point is worth making, since criticism of nativism is often superficial. It is not a claim of those who believe that some cognitive structures are innate, that environmental factors do not contribute to the acquisition of the relevant cognitive structures[165]. On the contrary, they are essential. However, leaving these questions aside, the account of language I have been relying on makes a quite restricted claim.

---

[165] Traits can be a very complex mixture of innate and learned. Monkeys fear snakes. However, experiments show that this is not true of monkeys raised in the laboratory away from snakes (this shows they need the learning stimulus of snakes). They do come to fear snakes if they see a video of other monkeys reacting to snakes (they learn not just from the snake but from the other monkeys). But they do not learn to fear flowers by watching a video making it look like other monkeys are giving the fear response to a flower. Thus, they don't learn fear for just anything; there is something innate to being able to learn about snakes (S. Mineka and R. Keir, "The effects of flooding on reducing snake fear in rhesus monkeys: 6 month follow up and further flooding", Behavioural Research and Therapy, 21, 527 – 535 (1983))

There is more or less universal agreement with the thesis that humans are born with a disposition to learn a language – they do so if they are given linguistic data and elephants do not. The further modest thesis I need to defend my position, that I have a clear case of knowing-how, is that humans are innately constrained to learn a language that conforms to a Universal Grammar (this is not to claim that our present understanding of this grammar is correct or complete, but only that it is constituted by an internal structure of computational principles and representations). The evidence for this thesis seems to be at least as convincing as that for most widely accepted scientific theories of which I am aware. It is given in great detail in the various references to the work of Chomsky I have cited (and in the work of numerous other contemporary linguists who in many details do not agree with Chomsky[166]). It is important to note that it is not necessary for my purposes to accept that the initial state of the infant embodies Universal Grammar rules and incorporates a specific language acquisition module[167]. I happen to accept Chomsky's view on these last matters but (i) they are widely doubted[168] and (ii) they are unnecessary for my argument. It is the case that a

---

[166] See for example R. Larson and G. Segal, "Knowledge of Meaning", Cambridge: MIT Press, 1995

[167] The idea of module here is that there is a Language Acquisition Device (LAD), which is well described as governed by local mental processes shared with one another but not with typical instances of global ones. These local mental processes are generally taken to conform to Turing's theory that thinking is computation. This is not to commit to the idea that all thinking can be described in this way or that the brain is entirely modular. Chomsky's claim is restricted to the existence of a biological object, something like a LAD, which can be the subject of fruitful experimental study.

[168] See for example M. Devitt and K. Sterelny, "Language and Reality", Cambridge: MIT Press, 1999 and A.S. Reber, "Implicit Knowing and Tacit Knowledge", Oxford: Oxford University Press, 1993. Reber, for example, gives a lengthy counter argument to Chomsky's well-known poverty of stimulus argument. He says nothing in rebuttal of Chomsky's powerful argument about Creolization (although this too is admittedly the subject of controversy).

modular language acquisition device gives a very natural explanation of the inapplicability of the Generality Constraint to subdoxastic states. However, to repeat, it is not necessary to assume such a module; Evans's argument stands up perfectly well without it.

On the other hand, there is recent evidence that a gene FOXP2 which, among other roles it has, is involved (of course along with other genes) in human language capacity. This gene in humans differs from FOXP2 in non-human primates and mice by a small and it is believed, recent, modification (less than 200 000 years ago and possibly as little as 50 000 years ago). A mutation of FOXP2 has been linked to a severe language disorder (Developmental Verbal Dyspraxia) that causes difficulty with many aspects of speech and language, including articulation, language processing, and grammatical skills (W. Enard, M. Przeworski, S.E. Fisher, C.S. Lai, V. Wiebe, T. Kitano, A.P. Monaco and S. Paabo, "Molecular evolution of FOXP2, a gene involved in speech and language", Nature, 418, 869 - 72 (2002)). This work, along with earlier work on heritable language disorders, provides support for Chomsky's arguments.

APPENDIX B. KNOWING-HOW AND ANOMALOUS MONISM

B1. Uncodifiability

If we accept virtue ethics as a doctrine, we are making the claim that the human agent must take moral actions that are based on grounds that are neither deductive nor inductive. Davidson's anomalous monism says that this is indeed what humans do, since their behaviour is governed by rationality and rationality is uncodifiable. Consequently, Davidson's position leads directly to the idea of knowing-how.

As I have tried to make clear in Chapter 7, we do not need Davidson's anomalous monism to claim that virtue may be knowing-how. The purpose of this appendix is simply to give a brief account of Davidson's argument, since for some authors the belief in knowing-how does stem from Davidson's views. I am thus not making any commitment to anomalous monism. It seems to me to be perfectly possible that a community could have predominantly false beliefs. We could still manage, for example, by making the (now false) assumption that most of their beliefs were true, and tempering this by observation of behaviour.

What precisely is the doctrine of uncodifiability? The central idea is that there is no complete set of rules that can serve in general as a precise guide for human action. Human action is based on the premise of rationality.

> "Just as the satisfaction of the conditions for measuring length or mass may be viewed as constitutive of the range of application of those sciences that employ these measures, so the satisfaction of conditions of consistency and rational coherence may be viewed as constitutive of

> the range of application of such concepts as those of belief, desire, intention, and action."[169]

This is not to claim that all individuals are always rational or hold only true beliefs, but that it is not possible to make sense of the world if it is assumed that individuals are generally irrational. We need to assume, for example, if someone desires X and believes Y will result in X, they will do Y unless there is some hindrance. We seek to optimize the fit between our estimate of what someone will believe and what we judge it would be rational for him or her to believe.

However, while there are some rules of rationality, such as if (P&Q) then P, and it is rational to believe at least as strongly in P as it is to believe in (P&Q), it is the case that there is no accepted general theory of rationality. It is then claimed that there can in principle be no general theory of rationality.

That is, rationality cannot be codified, or, it is uncodifiable.

I return to the merits of this claim below, but for the moment I need to consider the consequence of this claim.

The consequence is that one cannot, in the case of propositional attitude psychology, obtain the sort of precise explanation that is possible in basic physics. There are no strict psychophysical laws. That is, if rationality plays an essential role in determining mental states, and cannot be codified, a

---

[169] D. Davidson, "Essays on Actions and Events", Oxford: Oxford University Press, 1980, p 236

specification of physical states cannot explicitly completely determine all mental states.

This does not mean that there are not laws relating single physical and mental events. Individual mental events are related to physical events in the brain. Severe blows to the head are known to cause unconsciousness and, of course, there are other such laws, but it is the absence of a complete set of laws between mental event types and physical event types that is denied.

The point is that, by contrast, the entities that physics tries to subsume under its universal laws are not subject to the constraint of rationality, and hence a complete specification can, at least in principle, be made.

A counter argument has been advanced that propositional attitude psychology is not different, in this unavailability of a completely specifiable system of laws, from any of the other special sciences. These too must make use of *ceteris paribus* clauses. However, these clauses are, on the Davidson argument, ineliminable in propositional attitude psychology, whereas a finite definable set of *ceteris paribus* clauses can be specified in the special sciences. Or at least if it cannot be specified, it is believed that it could be, because we do not have the obstacle of the uncodifiability of rationality. One can say, for example, at sufficiently low pressure and high temperature in the inert gases, PV = nRT. The conditions of applicability of the law have been specified and they could be specified to any required degree of precision, or, if not, it is the conviction of the chemist that a law could be determined which would always allow the prediction of the temperature of an inert gas, given the pressure and

volume (at least this would be true over some restricted range of these variables).

It can be argued, however, that the post Newtonian laws of physics do not provide a strict system of laws in the sense we have been using this idea. This is not to do with the probabilistic nature of quantum theory (which can be dealt with), but with the view of contemporary physics that says that there are no compelling grounds for proposing such strict lawfulness in the fundamental concepts. William Child[170], in an extensive review, offers three responses to this argument. First, fundamental physics is nevertheless qualitatively different because it approaches this "ideal" in a way that propositional attitude psychology cannot. Second, a search for such laws can serve as a "regulative ideal" in physics, whereas Davidson's argument rules this out for human action. Third, akin to my argument about the behaviour of ideal gases, there are sciences of matter that legitimately aim for strict laws, and such an aim is not tenable in propositional attitude psychology.

How convincing is the basis for all this – the assumption of uncodifiability on the grounds that we have at present no acceptable theory of rationality?

Alan Turing[171] commented on this question, in a paper that remains controversial, long before Davidson's well-known work and he is worth quoting[172].

---

[170] W. Child, "Causality, Interpretation, and the Mind", Oxford: Oxford University Press, 1994, Chapter 2

[171] A.M. Turing, Computing Machinery and Intelligence", Mind, 49, 433 – 460 (1950) p 451

> "From this it is argued that we cannot be machines. I shall try to reproduce the argument, but I fear I shall hardly do it justice. It seems to run something like his. "if each man had a definite set of rules of conduct by which he regulated his life he would be no better than a machine. But there are no such rules, so men cannot be machines." the undistributed middle is glaring[173]. I do not think the argument is ever put quite like this, but I believe this is the argument used nevertheless. There may however be a certain confusion between "rules of conduct" and "laws of behaviour" to cloud the issue. By "rules of conduct" I mean precepts such as "stop it you see red lights," on which one can act, and of which one can be conscious. By "laws of behaviour" I mean laws of nature as applied to a man's body such as "if you pinch him he will squeak." If we substitute "laws of behaviour which regulate his life" for "laws of conduct by which he regulates his life" in the argument quoted the undistributed middle is no longer insuperable. For we believe that it is not only true that being regulated by the laws of behaviour implies being some sort of machine (though not necessarily a discrete-state machine), but that conversely being such a machine implies being regulated by such laws. However, we cannot so easily convince ourselves of the absence of complete laws of behaviour as of complete rules of conduct. The only way we know of for finding such rules is scientific observation, and we certainly know of no circumstances under which we could say, "We have searched long enough. There are no such rules."

Another way of looking at a problem for Davidson's position is to recognize that it is part of his view that we can in fact distinguish between the rational and what is not rational[174]. We are not able to say how we do this in any complete way, but this does not tell us that there is no method in principle. Indeed, it can be argued that there must be a method, otherwise we are faced with the claim that our means of doing so is via some form of miracle. In other words, one might claim that we do in reasoning what we do in language; we

---

[172] The paper proposes a test for intelligence in machines that is now generally agreed not to be necessary (a fact conceded by Turing). There is still argument as to whether it is sufficient.

[173] The "glaring" undistributed middle is apparent if the argument is expressed as follows:
All things that are not mere machines don't follow rules.
People don't follow rules. Etc.

[174] D. Braddon-Mitchell and F. Jackson, "The Philosophy of Mind and Cognition", Oxford: Blackwell, 1996, p 156

have implicit knowledge that enables us to determine the rational. Of course, this does not imply that we will ever be able to explicitly obtain the codifiable version of our rationality.

While acknowledging that the uncodifiability of rationality is not susceptible to proof, Child, in his perceptive review[175], makes a number of points in answer to arguments of the type I have just given. One proposal is to challenge those who believe that rationality is codifiable to give reasons for their belief, and to undermine the motivations for seeking such reasons. If knowledge can result from rationality in the absence of codifiability, one motive would be removed. Similarly, if rational explanation can provide causal explanations in the absence of codifiability, another is removed. He goes on to plausibly justify both contentions. Another, visibly weaker, argument is that since we have looked unsuccessfully for codifiability for centuries, it is reasonable to assume it does not exist (this might be construed as a response to the Turing argument given above). He also cites the arguments of philosophers like Kuhn about the impossibility of rationally comparing theories in science based on different paradigms. This is part of Kuhn's controversial theory about scientific revolutions. These arguments within the philosophy of science are hotly contested, and it would take us too far afield to enter into this debate here.

The question is complex and my own conclusion is that one cannot give a sufficiently compelling set of arguments to force one to accept uncodifiability

---

[175] W. Child, *ibid*

in general. Nor am I claiming that one can clearly reject uncodifiability in general. However, I believe that it is not necessary to assume that the judgements about rationality, relevant to propositional attitude statements in, say, mathematics, need be the same as those in morals or aesthetics. Someone who adopts Turing's position about our fundamental machine-like nature is not committed to the view that a machine could replicate our aesthetic judgements. It is perfectly possible that human aesthetic judgements are incoherent. There is near universal agreement that if you strongly want a beer the answer is to go and get one from the refrigerator, if you believe you have one in the refrigerator[176]. On the other hand, there is no general agreement on whether plastic gnomes from the local department store enhance the appearance of one's front lawn. Even within a particular community, for example, the educated middle class in Cape Town, there is complete disagreement as to whether an abstract impressionist like Jackson Pollock is a great artist, a charlatan, or lies somewhere in between. Moreover, it is possible for one person to entertain the first view of Jackson Pollock and then the second and then the first again, all within the space of a few minutes, and all while looking at one of Pollock's works in the National Gallery in

---

[176] Of course a critic might say, "If the refrigerator will explode when opened, the presence of the beer in the refrigerator is no longer salient, so what is the difference between this case and the case of morality?" The answer of the moral particularist is that the presence of the beer in the refrigerator remains a reason for going there. There is an overriding counterargument. In the moral case, exactly the same feature may be an argument pro or con, depending on other features. It may seem kind to chop off someone's poisoned finger to greatly improve their chance of survival until one realizes they live only for the pleasure of playing the violin exquisitely. In the former case the beer's presence remains germane to getting beer. The amputation remains germane to survival but to what extent it is germane to kindness depends interactively on all the features of the case. The virtuous person would, given the facts, know what to do.

Canberra. Moreover, this is not a mere logical possibility, it is something that not only can occur but also in fact does occur quite frequently.

The question of the essentiality of rationality is made even more complex by the possibility that we may not always truly know what we think, even in the cases where we have a conscious thought. Often we think we want X rather than Y until the moment when we are actually given the choice and we choose Y[177].

Davidson argues, that at every point, rationality compels us to think whatever the evidence best warrants, given everything else that we already think. Rationality genuinely requires, or demands, that we think in the way we do in the ordinary cognitive domain. It doesn't always require the same things of all people and it does not apply in the same way to all domains of thought and it does not apply to all the functions of the mind/brain. Nevertheless, despite Davidson's views, rationality could be codifiable but not constitutive of our beliefs, say, in aesthetics.

---

[177] Davidson acknowledges that this is the case but says that, "the possibility that one may be mistaken about one's own thoughts cannot defeat the overriding presumption that a person knows what he or she believes; in general the belief that one has a thought is enough to justify that belief" (D. Davidson, "Knowing One's Own Mind", reprinted in "Mind and Cognition, (Ed G. Lycan) Oxford: Blackwell, 1999, p 383). This response seems question begging, one does not know when one does not know, and it is not clear how one would quantify the occasions when one does and does not. (Davidson states, and elsewhere defends, the position that these unknown beliefs could not be standard, but this is a different issue if one makes a common sense interpretation of standard.) I can have indefinitely many false beliefs of which I am unaware. Someone might believe that it is obvious that every point on a finite line segment can be represented by a rational number.

B.2 Wittgenstein and the View from Nowhere

I have tried to outline the case for the centrality of an agent's perspective in ethics, based on the nature of ethical judgements; namely, that they are contextual and subject to the conflicting values of the agent. Of course, as I stated, if one accepts Davidson's anomalous monism the agent's perspective necessarily enters.

There is a subtly different account to that of Aristotle, which I believe makes clear how it is that we are forced to accept knowing-how, provided my earlier arguments about uncodifiability in the specific case of morality are accepted. Wittgenstein[178], in talking about the genuineness of expressions of feeling, says the following

> "Correct prognoses will generally issue from the judgements of those with better knowledge of mankind. Can one learn this knowledge? .....What one acquires here is not a technique; one learns correct judgements. There are also rules but they do not form a system, and only experienced people can apply them right. Unlike calculating-rules."

In other words, these judgements are knowing-how, as I have characterized it. It is a precise account of the position of the virtue ethicist. In particular, Wittgenstein is implying that no propositional account can be given of moral rules.

However, contrary to the argument I have presented, it is commonly argued (as it is by McDowell) that Wittgenstein's rule following argument, which I outline in a moment, implies that the knowing-how argument I have just made

---

[178] L. Wittgenstein, "Philosophical Investigations", Oxford: Blackwell, 1958, p 227

can be applied generally to judgements of all sorts. If this were so my present argument for knowing-how (in the most difficult case of the presumption of codifiable laws of rationality, *contra* Davidson and McDowell) would clearly be false. It would be false because beliefs are based on rationality which is codifiable, and moral beliefs would likewise follow from a completely specifiable set of principles, and hence fall under knowing-that. I now argue that this is not so.

If one accepts (a) the Davidson view of belief formation (b) the view that *contra* Davidson (and McDowell), rationality is codifiable and (c) with McDowell that moral judgements and judgements in mathematics are objective in the same way, then it follows that moral judgements can be determined by rules in the same way as mathematical judgements. In this context it is (c) I want to dispute.

Wittgenstein says that we can continue the series 2 4 6 8 10 in infinitely many ways, and indeed this is obvious. Hence he says,

> "This was our paradox: no course of action could be determined by a rule, because every course of action can be made out to accord with the rule. The answer was: if everything can be made out to accord with the rule, then it can also be made out to conflict with it. And so there would be neither accord nor conflict here."[179]

McDowell says the normal account of the person who continues with 12 14 ....1000 1002 .... is that there is an underlying psychological mechanism that churns out these numbers in a way analogous to a reliable machine. He

---

[179] L. Wittgenstein, Philosophical Investigations, Oxford: Blackwell, 1953, 201

argues that this supposition is otiose, and this is indeed Wittgenstein's point. Nor is Wittgenstein making a sceptical point; we do not doubt that we ourselves, or indeed, others could continue the simple series in the accepted fashion.

The point is that it is our shared "forms of life" that govern the usual continuation of the series. Put another way, the custom that determines the continuation of the series exists within a community of speakers. And it is an essential part of McDowell's position that this community of insiders can discern how to proceed, and analogously so they can in matters of ethics. The mathematical example is commonly, McDowell argues, misinterpreted as though there was not an identically valid insider-based approach to morals. In McDowell's view, an outsider would no more grasp our mathematics than she would our morals and both, like redness, are equally objective because there is no "view from nowhere", no "real" objectivity to be contrasted with this objectivity. McDowell attempts to dispose of the argument that a practice like, for example, always being kind, cannot be just as objective as mathematics; they are both of the same ilk, he maintains. Hence they are both evident in the same way to an insider in the community.

This is the conclusion I want to deny, or at least I want to deny that it is a necessary conclusion. We can, for the sake of argument, accept the idea that our practice of mathematics is community defined, and we can accept that redness, say, is a response dependent property, outside the purview of the

natural sciences, and so are moral judgements. But as Lang[180] points out, what is comic or erotic or nice-tasting is also response dependent.

This particular commonality, the fact that there is no "view from nowhere", does not mean that our judgements in mathematics and morals are similarly based and similarly objective. There is a difference between what is kind or funny or beautiful on the one hand and what is red or what constitutes a mathematical proof, on the other. It seems a different proposition to say that ripe lemons would still be yellow if we all went blind, than to say that pictures of Marilyn Monroe would still be erotic if there were no heterosexuals or lesbians any more, but only gay men. It is not necessary to take McDowell's cognitivist view of morals as a consequence of Wittgenstein's rule following. It is also, of course, not necessary to take McDowell's (and my) view of morals as uncodifiable from Wittgenstein's rule following, but that is another and not immediately relevant question[181].

---

[180] G. Lang, "The Rule-Following Considerations and Metaethics: Some False Moves", European Journal of Philosophy, 9, 190 – 209 (2001), p 201.

[181] G. Lang *ibid* p 190