

# **A Practical Guide to Intellectualism**

Yuri Cath

A thesis submitted for the degree of Doctor of Philosophy of  
The Australian National University

PHILOSOPHY PROGRAM  
RESEARCH SCHOOL OF SOCIAL SCIENCES  
COLLEGE OF ARTS AND SOCIAL SCIENCES  
THE AUSTRALIAN NATIONAL UNIVERSITY  
CANBERRA ACT 0200

**AUGUST 2008**

Except where otherwise acknowledged in the text, this thesis represents the original research of the author.

Yuri W. Cath

August 2008.

## Abstract

---

In this thesis I examine the view—known as *intellectualism*—that knowledge-how is a kind of knowledge-that, or propositional knowledge. I examine issues concerning both the status of this view of knowledge-how and the philosophical implications if it is true. The ability hypothesis is an important position in the philosophy of mind that appeals to Gilbert Ryle’s famous idea that there is a fundamental distinction between knowledge-how and knowledge-that. This position appears to be inconsistent with the truth of intellectualism. However, I demonstrate in this thesis that the ability hypothesis can be restated using the intellectualist view of knowledge-how. With regards to the status of intellectualism, I argue that the two main traditional arguments against intellectualism do not succeed. I also provide new and, I claim, successful arguments against intellectualism. These arguments point to a new view of knowledge-how that is distinct from both the standard intellectualist and Rylean views of knowledge-how.

## Acknowledgements

---

This thesis has benefited greatly from the input of others. I especially want to thank my supervisors Daniel Stoljar, David Chalmers and Martin Davies for their intellectual generosity, encouragement and patience. In writing this thesis I have also been aided by feedback on the ideas in it from numerous people including Jonathan Bengson, Berit Brogaard, Andy Egan, Al Hájek, Jonathan Ichikawa, Brendan Jackson, Frank Jackson, Derk Pereboom, Denis Robinson, Jonathan Schaffer, Nico Silins, Paul Snowdon, Nicholas Southwood, Jason Stanley, David Wall, Jordi Valor, an anonymous referee for *Noûs*, and audiences at The Australian National University, The University of North Carolina Chapel Hill and The University of St Andrews. My apologies to anyone I ought to mention here that I have forgotten.

I would also like to thank all of my friends within the RSSS Philosophy Program, and others outside it, for their support and friendship during my candidature at the ANU. Some of these I have mentioned above but others that I especially want to thank include: Ben Blumson, Carl Brusse, Jacek Brzozowski, Brett Calcott, Kirsty Douglas, Chris Haskett, Bryce Huebner, Ben Jeffares, Eric Mandelbaum, John Matthesson, Mark Phelan, Karen Riley, Lila San Roque and Stewart Saunders.

This thesis is dedicated to my parents Gerben and Nerida, without whom I would never have started it (for many reasons other than just the obvious one), and my wife Stephanie, without whom I would never have finished it.

# Contents

---

ABSTRACT .....	I
ACKNOWLEDGEMENTS .....	II
CONTENTS.....	III
INTRODUCTION .....	1
CHAPTER 1 INTELLECTUALISM AND THE INSUFFICIENCY OBJECTION.....	7
1.1 The Proper Target of Intellectualism .....	7
1.2 Simple Intellectualism.....	8
1.3 Simple Intellectualism and Ryleanism .....	10
1.4 The Insufficiency Objection to Simple Intellectualism .....	14
1.5 Sophisticated Intellectualism .....	14
1.6 Brogaard's Defence of Simple Intellectualism .....	25
1.7 Assessing the Insufficiency Objection .....	29
CHAPTER 2 THE ABILITY HYPOTHESIS AND INTELLECTUALISM.....	35
2.1 The New Knowledge-how.....	35
2.2 The Ability Hypothesis.....	39
2.3 Mary's 'New' New Knowledge-how .....	42
2.4 Where are the Abilities?.....	45
2.5 Old Facts New Modes.....	51
2.6 Conclusions .....	58
CHAPTER 3 KNOWING HOW WITHOUT KNOWING THAT .....	62
3.1 Three Counterexamples .....	62
3.2 The No Knowledge-that Claims.....	70
3.3 The Knowledge-how Claims .....	76
3.4 Toward a New Theory .....	85
CHAPTER 4 REGARDING A REGRESS.....	93
4.1 Ryle's Regress.....	94
4.2 The Contemplation Regress .....	109
4.3 The Employment Regress.....	118
4.4 Intellectualist Responses to the Employment Premise.....	128
4.5 Related Regress Arguments .....	136
4.6 Assessing the Regress Objection.....	147
BIBLIOGRAPHY .....	153

## Introduction

---

The subject of this dissertation is intellectualism. Intellectualism is a controversial view about the nature of the knowledge one has when one knows how to do something—what is commonly referred to as *knowledge-how*. Examples of knowledge-how then include knowing how to swim, knowing how to add 2 and 4, knowing how to cook tuna pie, and so on. The standard description of intellectualism is that it is the view that knowledge-how is a kind, or sort, or species, of *knowledge-that*. Knowledge-that, of course, is the knowledge one has when one knows that something is the case, for example, when one knows that 2 + 4 is 6, or one knows that tuna pie is delicious.

Knowledge-that is also called ‘propositional knowledge’ because it is thought that to know that Athens is the capital of Greece, say, is to stand in a certain relation to the proposition that Athens is the capital of Greece. In Chapter 1 I will offer a more precise characterization of intellectualism, but for now it will suffice to say that intellectualism is the view that knowing how to do something is also a matter of standing in this same relation—I will call it simply *the knowledge-that relation*—to some proposition. That is, intellectualism tells us that knowing how to swim, like knowing that Athens is the capital of Greece, is a matter of knowing that p, for some proposition p.

Gilbert Ryle (1946, 1949) famously argued that intellectualism was an untenable position, and he also advanced his own positive view, according to which knowledge-how is a kind of dispositional capacity or ability. Ryle’s two-fold view of knowledge-how has always had its dissenters, but nonetheless it has effectively been the orthodox view of knowledge-how in analytic philosophy; that is, the orthodox position has been that intellectualism is false, and that the correct account of knowledge-how is something like Ryle’s account, whereby to know how to do something is to possess a certain ability or disposition. Furthermore, Ryle’s distinction between knowledge-how and knowledge-that became a familiar

tool in the kit of the contemporary philosopher, with this distinction playing an important role in debates in the philosophy of mind, the philosophy of language, and other areas.

Recently, however, there has been a resurgence of interest in intellectualism. In particular, Jason Stanley and Timothy Williamson's (2001) paper "Knowing How" has played a highly influential role in reviving the intellectualist view of knowledge-how. Stanley and Williamson presented an intellectualist account of knowledge-how and provided linguistic arguments in support of it. Furthermore, they, and others such as Paul Snowdon (2004), have argued that Ryle did not present a successful regress argument against intellectualism, and that his own account of knowledge-how is subject to clear counterexamples. Intellectualist accounts of knowledge-how have also been advanced by John Bengson and Marc Moffett (2007) (see also Bengson et al. forthcoming), and Berit Brogaard (forthcoming a, forthcoming b).

#### *Four Questions*

The recent enthusiasm for intellectualist accounts of knowledge-how raises a number of interesting philosophical issues. In particular, I take the following to be four key research questions concerning intellectualism:

- (1) Is the knowledge-that that intellectualists appeal to sufficient for knowledge-how?
- (2) Is the knowledge-that that intellectualists appeal to necessary for knowledge-how?
- (3) Is there some successful regress argument against intellectualism?
- (4) What are the philosophical implications of intellectualism?

Question 1 is related to what is perhaps the most obvious form of objection one might raise to intellectualism, namely, that merely knowing that some proposition is true does not appear to be a sufficient condition for knowing how to do something. Consider, for example, the action of riding a bicycle. Is there really some proposition  $p$  such that merely knowing that  $p$  is sufficient for knowing how to ride a bicycle? Many critics of intellectualism, including Ryle himself, have argued that there is not, on the grounds that no matter what this proposition  $p$  is one could know that  $p$  and yet still fail to know how to ride a bicycle.

Now it is not true that there is no proposition  $p$  such that knowing that  $p$  is a sufficient condition for knowing how to ride a bicycle. For example, if  $p$  is the proposition that one knows how to ride a bicycle then obviously knowing that  $p$  entails that one knows how to ride a bicycle. Still, knowing how to ride a bicycle is surely not a matter of knowing that one knows how to ride a bicycle. What this kind of insufficiency objection really amounts to then is something like the following claim: for any *prima facie* plausible account of knowledge-how whereby knowing how to  $F$  is a matter of knowing that  $p$  (for some proposition  $p$ ), one will be able to describe a possible scenario where someone knows that  $p$  but fails to know how to  $F$ . In other words, the real issue is whether the kind of knowledge-that that intellectualists appeal to is sufficient for knowledge-how.

Question 2 concerns the issue of whether knowing that some proposition is true could really be a necessary condition for knowing how to do something. Any intellectualist account of knowledge-how will be committed to some claim of the form: one knows how to  $F$  only if one knows that  $p$ , for some proposition  $p$ . But one might suspect that no matter what this proposition  $p$  is, it will be possible for one to know how to  $F$  even when one fails to know that  $p$ . Or, again, one might at least suspect that for any *prima facie* plausible intellectualist account of knowledge-how, whereby knowing how to  $F$  is a matter of knowing that  $p$  (for some proposition  $p$ ), one will be able to describe a possible scenario where someone knows how to  $F$  but fails to know that  $p$ .

Question 3 concerns the most prominent, but also elusive, objection that has been made against intellectualism, namely, Gilbert Ryle's famous objection that intellectualism must be false because the assumption that it is true leads to an infinite and vicious regress. As mentioned earlier, intellectualists have argued that Ryle's regress argument fails, but many critics of intellectualism have replied that intellectualists have either misinterpreted Ryle's argument or that there is some other related regress argument which does succeed. The issue then of whether or not there is a successful regress argument against intellectualism is still rather obscure and in need of further examination.

Unlike the previous questions, Question 4 concerns not the status but rather the consequences of intellectualism. There has been a lot written on the issue of whether or not intellectualism is true, but very little on what follows if it is true. But the truth of intellectualism has potential implications for any area of philosophy where the knowledge-how versus knowledge-that distinction has played an important role. Indeed, Stanley and Williamson (2001: 441–4) have claimed that the truth of their intellectualist account of knowledge-how undermines certain philosophical positions that rely on Ryle's distinction between knowledge-how and knowledge-that. However, while Stanley and Williamson's account of knowledge-how has received a great deal of critical attention, these claims about the implications of their account have not been subjected to the same kind of scrutiny.

### *Overview*

The purpose of this dissertation is to address the four questions outlined above. In particular, my main concern is to develop three independent but closely related essays in Chapters 2–4, which respectively address the question of what the philosophical implications of intellectualism are, the question of whether knowledge-that is necessary for knowledge-how, and the question of whether there is some successful regress argument against intellectualism.

In Chapter 2, ‘The Ability Hypothesis and Intellectualism’, I address the question of what the philosophical implications of intellectualism are. In particular, I examine a supposed implication of intellectualism for what is probably the most well-known application of the knowledge-how versus knowledge-that distinction in philosophy. The ability hypothesis, endorsed by David Lewis (1998) and Laurence Nemirow (1980, 1990) amongst others, is an important form of reply to Frank Jackson’s (1982, 1986) famous knowledge argument against physicalism. And this reply crucially relies on a distinction between knowledge-how and knowledge-that. Not surprisingly then, Stanley and Williamson argue that the truth of their intellectualist account of knowledge-how is inconsistent with the ability hypothesis reply to Jackson’s argument. However, I shall argue that we can restate the core claims made by the ability hypothesis using Stanley and Williamson’s account of knowledge-how.

In Chapter 3, ‘Knowing How Without Knowing That’, I address the question of whether the kind of knowledge-that appealed to by intellectualists is necessary for knowledge-how. I argue that, given certain very standard assumptions about the nature of knowledge-that, intellectualism is subject to three different kinds of counterexample. Each putative counterexample is a scenario where someone knows how to F but they fail to stand in the knowledge-that relation to any proposition p, such that their knowing how to F might plausibly be a matter of their knowing that p. The counterexamples differ with respect to the reason why the subject fails to possess the relevant knowledge-that.

After presenting and defending these arguments, I then go on to suggest that these new arguments against intellectualism point to a new view of knowledge-how. This alternative view of knowledge-how is distinct from both intellectualism and Ryleanism, although it shares important features with both of these views.

In Chapter 4, ‘Regarding a Regress’, I address the elusive question of whether there is a successful regress argument against intellectualism. I begin by returning to Gilbert Ryle’s regress argument against what he called *the intellectualist legend*, as this argument is the source of the idea that there is a successful regress argument

against intellectualism. Ryle's regress argument is demonstrably not such an argument. However, one might think that Ryle's argument points to some related and successful regress argument against intellectualism. I go on to consider two regress arguments against intellectualism that are related to Ryle's argument: what I call *the contemplation regress argument* and *the employment regress argument*. I argue that neither of these arguments succeeds. Furthermore, I show that the regress arguments against intellectualism presented by Stephen Hetherington (2006) and Alva Noë (2005) are also undermined by the same kind of problems faced by the contemplation and employment regress arguments.

Before proceeding to the discussion in Chapters 2–4, however, it will be useful to take a closer look at intellectualism. In Chapter 1, 'Intellectualism and the Insufficiency Objection', my aim is to set the scene for the discussion to follow by considering the intellectualist view of knowledge-how in more detail and to also address the first of our four questions; that is, the question of whether the kind of knowledge-that that intellectualists appeal to is sufficient for knowledge-how. These two tasks are closely intertwined because, as we will see, the major differences between the various existing intellectualist accounts of knowledge-how are related to the different ways that intellectualists have attempted to respond to the insufficiency objection to intellectualism. I shall argue that once we distinguish between two broad kinds of intellectualism it is clear that the insufficiency objection could not form the basis of any successful argument against intellectualism in general.

# Chapter 1 Intellectualism and the Insufficiency Objection

---

In the introduction I set out four questions concerning intellectualism. My aim in this chapter is to address the question of whether the insufficiency objection shows us that intellectualism is false. To address this question, however, we need to first take a closer look at the intellectualist view of knowledge-how. In §1 I begin by making a clarification about the proper target of intellectualist accounts of knowledge-how. In §2 I identify one important kind of intellectualism, what I call *simple intellectualism*. In §3 I discuss the relationship between simple intellectualism and Rylean accounts of knowledge-how. In §4 I describe how simple intellectualism clearly faces an insufficiency objection. The standard response to this objection is to reject simple intellectualism in favour of some version of what I call *sophisticated intellectualism*. In §5 I characterize this response to the insufficiency objection and show how the intellectualist accounts of knowledge-how given by Stanley and Williamson (2001) and Bengson and Moffett (2007) are both versions of sophisticated intellectualism. In §6 I discuss an interesting attempt to defend a version of simple intellectualism by Berit Brogaard (forthcoming a, forthcoming b). In §7 I argue that the insufficiency objection strongly suggests that there is no plausible version of simple intellectualism, but that it does not show us that there is no plausible version of sophisticated intellectualism.

## 1.1 The Proper Target of Intellectualism

The standard description of intellectualism is that knowledge-how is a kind, or sort, or species, of knowledge-that. It is worth noting a sense in which this description, whilst standard, is somewhat misleading. The reason is that the term ‘knowledge-how’ is quite naturally interpreted in a broad way, whereby it refers to the knowledge attributed by any sentence of the form ‘S knows how ...’ and not

merely to the knowledge attributed by sentences of the form ‘S knows how *to* ...’. On this broad interpretation then, both (1) and (2) attribute knowledge-how to Ari:

(1) Ari knows how to swim.

(2) Ari knows how Ian Thorpe swims.

But the debates about whether ‘knowledge-how’ is a kind of knowledge-that have focused almost exclusively on the issue of whether the knowledge attributed by sentences like (1) is a kind of knowledge-that, and not the issue of whether the knowledge attributed by sentences like (2) is a kind of knowledge-that. Accordingly, when proponents of intellectualism state their accounts of ‘knowledge-how’ they provide analyses of the form ‘S knows how *to* F if and only if ...’. And, insofar as they consider the matter, even some ardent opponents of intellectualism are happy to grant that the knowledge attributed by sentences like (2) is a kind of knowledge-that (for example, see Noë 2005: 284). It would be more accurate then to talk of knowledge-how-*to*, rather than knowledge-how, when discussing intellectualism. But given the entrenched use of the term ‘knowledge-how’ in the literature on intellectualism, I will continue to use this term in this dissertation. As I use this term, however, it should be understood in a narrow way, whereby it refers only to the kind of knowledge one has when one knows how to do something.

## 1.2 Simple Intellectualism

The proper target of intellectualism is knowledge-how, in the sense of knowing how to do something. With this clarification in place we can now take a closer look at the intellectualist view of knowledge-how. In the introduction, I characterized intellectualism as the view that knowing how to do something is a matter of standing in the knowledge-that relation to some proposition. But what proposition? Intellectualists hold that knowing how to perform some action F is a

matter of standing in the knowledge-that relation to some proposition, the content of which concerns a *way* to F, where the relevant sense of the word ‘way’ is the sense whereby it denotes something like a method, technique or procedure for performing an action. The standard idea is that S’s knowing how to F is a matter of there being some way *w* such that S stands in the knowledge-that relation to the proposition that *w* is a way to F, or the proposition that *w* is a way for S to F, or some variant thereof.

What we can call *simple intellectualism* is the view that knowing how to perform some action is *solely* a matter of standing in the knowledge-that relation to some such proposition. The simple intellectualist then, is someone who is committed to the truth of some instance of the following equivalence claim:

The Simple Equivalence Thesis

Necessarily, S knows how to F if and only if there is some relevant proposition p concerning a way to F such that S stands in the knowledge-that relation to p.

Now, someone who endorses some instance of this equivalence claim could also go further and endorse the corresponding instance of the following claim:

The Simple Identity Thesis

To know how to F is to know that p (for some relevant proposition p concerning a way to F).

But I think the best way of characterizing simple intellectualism is to use an equivalence claim. For one thing, when intellectualists give the official statement of their view they typically offer some analysis of the form ‘S knows how to F if and only if ...’. And, for reasons that I will discuss in Chapter 4, someone might conceivably endorse such an equivalence claim whilst arguing that this need not entail the corresponding identity claim (for an example of this stance see Bengson et al. forthcoming: fn. 3). It is important to make note of the simple identity thesis, however, as intellectualists often do take the further step of endorsing

some identity claim. And in Chapter 4 we will see that regress arguments against intellectualism are normally formulated as arguments against some instance of the simple identity thesis.

### 1.3 Simple Intellectualism and Ryleanism

When Gilbert Ryle argued that intellectualism was an untenable view of knowledge-how, what he had in mind was what I have called ‘simple intellectualism’. Ryle not only rejected simple intellectualism but he also offered an alternative account of knowledge-how. But what is the Rylean view of knowledge-how and what is its relationship to simple intellectualism?

Gilbert Ryle is often interpreted as claiming that to know how to F is to possess the ability to F, which in turn is to possess a certain complex of abilities. For example, Stanley and Williamson (2001: 411) write that: “According to Gilbert Ryle ... knowledge-how is an ability, which is in turn a complex of dispositions”. They go on to show that they interpret Ryle as claiming that to know how to F is to possess the *corresponding* ability to F. In other words, Ryle is interpreted as being committed to the following three identity claims:

(ID.1) To know how to F is to possess the ability to F.

(ID.2) To possess the ability to F is to possess a certain complex of dispositions.

(ID.3) To know how to F is to possess a certain complex of dispositions.

To parallel this presentation of simple intellectualism, however, it will be useful to focus on the equivalence claims entailed by these identity claims:

(EQ.1) S knows how to F if and only if S possesses the ability to F.

(EQ.2) S possesses the ability to F if and only if S possesses a certain complex of dispositions.

(EQ.3) S knows how to F if and only if S possesses a certain complex of dispositions.

Intellectualists frequently argue against Ryle's account of knowledge-how by providing counterexamples to EQ.1. For example, Stanley and Williamson present the following cases as counterexamples to Ryle's account of knowledge-how:

[a] ski instructor may know how to perform a certain complex stunt, without being able to perform it herself. Similarly, a master pianist who loses both of her arms in a tragic car accident still knows how to play the piano (Stanley and Williamson 2001: 416).

Such cases<sup>1</sup> appear to be counterexamples to EQ.1 because they are cases where, intuitively, someone knows how to F even though they do not possess the ability to F. In which case, contra EQ.1, possessing the ability to F is not a necessary condition for knowing how to F. Snowdon (2004) and Bengson and Moffett (2007) also provide examples that are meant to be cases where intuitively someone has the ability to F but does not know how to F. Snowdon, for example, presents the following scenario as an example of such a case:

A man is in a room, which, because he has not explored it in the least, he does, as yet, not know how to get out of. In fact, there is an obvious exit which he can easily open. He is perfectly able to get out, he can get out, but does not know how to (as yet) (Snowdon 2004: 11).

And Bengson and Moffett present the following scenario:

---

<sup>1</sup> Stanley and Williamson credit the ski instructor to Jeff King. Others have appealed to similar cases to support the claim that one can know how to F without possessing the ability to F, including: Noam Chomsky (1998), Carl Ginet (1975: 8) and Paul Snowdon (2004).

Suppose that Irina is seriously mistaken about how to perform a salchow. She believes incorrectly that the way to perform a salchow is to take off from the front outside of her skate, jump in the air, spin, and land on the front inside edge of her skate. (The correct sequence is to take off from the *back inside edge* and land on the *back outside edge* of the opposite foot after one or more rotations in the air.) However, Irina has a severe neurological abnormality that makes her act in ways that differ dramatically from how she actually thinks she is acting. Whenever she actually attempts to do a salchow (in accordance with her misconceptions) this abnormality causes her to reliably perform the correct sequence of moves. So, although she is seriously mistaken about how to perform a salchow, whenever she actually attempts to do a salchow (in accordance with her misconceptions) the abnormality causes Irina to perform the correct sequence of moves, and so she ends up successfully performing a salchow. Despite the fact that what she is doing and what she thinks she is doing come apart, she fails to notice the mismatch. In this case, it is clear that Irina is (reliably) able to do a salchow. However, due to her mistaken belief about how to perform the move, she cannot be said to know how to do a salchow (Bengson and Moffett 2007: 46).

Both of these examples are meant to demonstrate that, contra EQ.1, having the ability to F is not a sufficient condition for knowing how to F. Intellectualists have argued then that Ryle's claim that S knows how to F if, and only if, S possesses the ability to F is doubly wrong, for one can know how to F but fail to possess the corresponding ability to F, and one can possess the ability to F but fail to know how to F.

However, while Ryle was clearly committed to EQ.3—that is, the claim that S knows how to F if and only if S possesses a certain complex of dispositions—it is actually not entirely clear that Ryle equated knowing how to F with simply possessing the *corresponding* ability to F. In which case, it is not clear that the counterexamples given to EQ.1 are actually counterexamples to Ryle's account of knowledge-how. We will have reason to return to this issue in Chapter 3. For now, I will simply borrow a term from Bengson and Moffett (2007) and call any view that identifies or equates knowing how to F with possessing the ability to F *neo-Ryleanism*, as a way of acknowledging that while this view is closely associated with Ryle it may turn out that Ryle himself was not committed to it.

Now, while it is not entirely clear that neo-Ryleanism formed an essential part of Ryle's view of knowledge-how, it is the most prominent alternative to intellectualism. This may explain why intellectualists often appear to regard their counterexamples to neo-Ryleanism as offering strong support to their own view of knowledge-how. For example, as Bengson et al. (forthcoming: 14) write of the counterexamples they offer: "Understood as counterexamples to neo-Ryleanism, these vignettes *serve to lend substantial plausibility to radical intellectualism*" (emphasis added).<sup>2</sup>

However, the conclusion that neo-Ryleanism is false only really provides very indirect support for intellectualism. To see why, compare neo-Ryleanism with simple intellectualism. Neo-Ryleanism tells us that S knows how to F if and only if one possesses the ability to F, whereas simple intellectualism tells us that one knows how to F if and only if, for some relevant proposition p concerning a way to F, one knows that p.

Now, these two equivalence claims are at best contraries, for while one might reasonably argue that they cannot both be true it clearly could be the case that they are both false. And the same point clearly applies to the equivalence claim entailed by EQ.3, that is, the claim that one knows how to F if and only if one possesses a certain complex of dispositions.<sup>3</sup>

The point that intellectualism and Ryleanism are, at best, contraries is obvious but important to mention because in the literature intellectualism and Ryleanism are often the only alternative accounts of knowledge-how that are discussed. This indicates, I suspect, that many participants in the debates about the nature of

---

<sup>2</sup> Bengson and Moffett (forthcoming: 1) use the term 'radical intellectualism' for the view that "S knows how to  $\psi$  if, and only if, S possesses a certain sort of propositional knowledge concerning  $\psi$ ".

<sup>3</sup> Actually, in the case of simple intellectualism and EQ.3 it is not even clear that they are contraries. Supposing one accepted some dispositional account of knowledge-that, one would then have to allow that simple intellectualism and EQ.3 could even turn out to be equivalent.

knowledge-how assume that intellectualism and Ryleanism are the only serious alternative views of knowledge-how worth considering. I shall argue in Chapter 3, however, that this assumption is mistaken.

#### 1.4 The Insufficiency Objection to Simple Intellectualism

One of the most prominent objections to simple intellectualism is that merely standing in the knowledge-that relation to some proposition concerning a way to F does not appear to be a sufficient condition for knowing how to F. To illustrate this objection, consider a version of simple intellectualism whereby S knows how to F if and only if there is some way  $w$  such that S knows that  $w$  is a way to F.

Imagine now that you are watching the Tour de France on TV with Hannah, who has never learnt to ride a bicycle. Pointing to one of the cyclists you gently rib her by remarking: “That’s a way for you to ride a bicycle”. Consequently, Hannah comes to know that that way is a way to ride a bicycle. So, there is a way  $w$  such that Hannah knows that  $w$  is a way to ride a bicycle. But clearly Hannah still does not know how to ride a bicycle.

The general point is that it appears quite easy for someone to gain the kind of knowledge-that that simple intellectualists equate knowledge-how with, whilst failing to know how to F. But then simple intellectualism looks to be clearly false, as simple intellectualism tells us that standing in the knowledge-that relation to some such proposition is a sufficient condition for knowing how to F.

#### 1.5 Sophisticated Intellectualism

How might an intellectualist respond to the insufficiency objection? They could try to offer some defence of simple intellectualism, and in the next section I will consider an attempt to provide such a defence by Berit Brogaard. But the standard response to the insufficiency objection is to actually grant that merely standing in the knowledge-that relation to some proposition concerning a way to F is not a sufficient condition for knowing how to F. What we can call *sophisticated*

*intellectualism* then is the view that knowing how to do something is only *partly* a matter of standing in the knowledge-that relation to some relevant proposition concerning a way to perform that action. The sophisticated intellectualist agrees with the simple intellectualist that standing in the knowledge-that relation to some relevant proposition concerning a way to F is a necessary condition for knowing how to F. However, unlike the simple intellectualist, they hold that satisfying such a condition does not suffice for knowing how to F. Rather, they hold that knowing how to F is a matter of standing in the knowledge-that relation to the right kind of proposition and also satisfying some further condition. The sophisticated intellectualist then is someone who is committed to the truth of some instance of the following equivalence claim:

The Sophisticated Equivalence Thesis

S knows how to F if and only if, for some relevant proposition p concerning a way to F:

- (i) S stands in the knowledge-that relation to p, and
- (ii) S satisfies X (for some further condition X).

And, as with simple intellectualism, the sophisticated intellectualist may not only endorse an instance of the sophisticated equivalence claim, but they may also go a step further and endorse the corresponding instance of the following claim.

The Sophisticated Identity Thesis

To know how to F is to:

- (i) To know that p (for some relevant proposition p concerning a way to F)
- (ii) To satisfy X (for some further condition X).

Although, as with the simple equivalence thesis, a sophisticated intellectualist may choose only to endorse an instance of the sophisticated equivalence thesis, and not the corresponding instance of the sophisticated identity thesis.

Now to see how this kind of response to the insufficiency objection works in practice I will consider two different versions of sophisticated intellectualism, namely, the accounts of knowledge-how offered by Stanley and Williamson (2001) and by Bengson and Moffett (2007) respectively.

*Stanley and Williamson's version of sophisticated intellectualism*

Stanley and Williamson's (2001: 441) account of knowledge-how as "simply a species of propositional knowledge" is by far the most prominent intellectualist account of knowledge-how. Their account of knowledge-how will be a major focus of the discussion in Chapters 2 and 3. For these reasons, it will be useful to look at their account of knowledge-how in some detail.

Stanley and Williamson state their account as an analysis of the truth conditions of knowledge-how ascriptions. On the basis of linguistic arguments that I will consider in Chapter 2, Stanley and Williamson offer an initial account of the truth conditions of 'S knows how to F' ascriptions, whereby they conform to the following schema:

'S knows how to F' is true relative to a context  $c$  if and only if there is some contextually relevant way  $w$  for S to F such that S stands in the knowledge-that relation to the proposition that  $w$  is a way for S to F.

However, as Stanley and Williamson recognise, this account is subject to the insufficiency objection. For example, in the Tour De France scenario Hannah will not only know that that way (i.e. the way the cyclists are riding) is a way to ride a bicycle, but she will also know that that way is a way for her to ride a bicycle. For recall that you tell Hannah: "That's a way for *you* to ride a bicycle." In other words, it is a context in which intuitively (1) is false even though (2) is true:

(1) Hannah knows how to ride a bicycle.

(2) Hannah knows that that way is a way for her to ride a bicycle.

Stanley and Williamson explicitly acknowledge the possibility of such a scenario. Their response to this insufficiency objection utilizes the idea “that one and the same way can be entertained under different modes of presentation” (*ibid.* 428). To explicate this idea Stanley and Williamson appeal to sentences like the following:

(3) John knows that he himself has burning pants.

(4) John knows that that man has burning pants.

There can be contexts in which it appears that (3) is false and (4) is true (at the same time) even though (3) and (4) arguably attribute to John knowledge of the very same singular proposition (when ‘that man’ refers to John). For example, we can imagine that John does not realize that that man he sees with burning pants is actually himself reflected in a mirror.

One solution to the problem presented by (3) and (4) is to say that these sentences attribute knowledge of the same proposition but under different modes of presentation: a first-personal and a demonstrative mode of presentation, respectively. This is meant to dissolve any tension between (3) being false and (4) true, since knowing a proposition under one mode of presentation does not entail one’s knowing it under any other mode of presentation.

According to Stanley and Williamson, analogous points apply to the problem presented by (1) and (2). On their preferred development of this strategy “verbs such as ‘believes’ and ‘knows’ express three-place relations between persons, Russellian propositions, and ways of thinking of Russellian propositions” (*ibid.* 427). Stanley and Williamson apply this framework so that (1) attributes knowledge to Hannah of at least one proposition that contains a “way of riding a bicycle” (*ibid.* 427). Furthermore, they claim that in order for an assertion of (2) to be true or appropriate there has to be some way *w* such that Hannah stands in the knowledge-that relation to the Russellian proposition that *w* is a way for Hannah to ride a bicycle under a particular type of mode of presentation—what they call a

*practical mode of presentation*. The idea here is that Hannah entertains this proposition under a practical mode of presentation by entertaining the way  $w$ , which is a constituent of this proposition,<sup>4</sup> under a practical mode of presentation.

The qualification ‘true or appropriate’ reflects the fact that Stanley and Williamson wish to remain neutral on the issue of whether this requirement—that one entertain the relevant proposition under a practical mode of presentation—should be viewed as relevant to the semantics of knowledge-how ascriptions, or as only relevant to the pragmatics of using such ascriptions.

Stanley and Williamson’s explanation then of why it appears that (1) is false even though (2) is true, would be that in the Tour de France scenario there is no way  $w$  such that Hannah stands in the knowledge-that relation to the proposition that  $w$  is a way for Hannah to ride a bicycle *under a practical mode of presentation*. In which case, on their view, it follows that either (1) is false or (1) is true. But an assertion of (1) is inappropriate, as it pragmatically conveys a falsehood, namely, that there is some way  $w$  such that Hannah stands in the knowledge-that relation to the proposition that  $w$  is a way for Hannah to ride a bicycle under a practical mode of presentation.

As Stanley and Williamson point out then, their final account of knowledge-how ascriptions can be described either of two ways depending on whether or not practical modes of presentation are semantically relevant:

So, here is our complete account of knowing-how. Suppose modes of presentation are semantically relevant. Then [1] is true relative a context  $c$  if and only if there is some contextually relevant way  $w$  such that Hannah stands in the knowledge-that relation to the Russellian proposition that  $w$  is a way for Hannah to ride a bicycle, and Hannah entertains this proposition under a practical mode of presentation. If modes of presentation are not semantically relevant then the truth of [1] does not require that

---

<sup>4</sup> Stanley and Williamson take ways to be properties of token events, for example, the way  $w$  of riding a bicycle such that Hannah knows that that way is a way to ride a bicycle, is understood to be a property of a token event of riding a bicycle.

Hannah entertain the proposition in question under a practical mode presentation, though a use of [1] pragmatically conveys that she does (Stanley and Williamson 2001: 430).

Now, this passage suggests that if practical modes of presentation are semantically relevant then it is Stanley and Williamson's view that the truth conditions of knowledge-how ascriptions conform to the following schema:

'S knows how to F' is true relative to a context  $c$  if and only if there is some contextually relevant way  $w$  for S to F such that:

- (a) S stands in the knowledge-that relation to the proposition that  $w$  is a way for S to F, and
- (b) S entertains the proposition that  $w$  is a way for S to F under a practical mode of presentation.

But this analysis cannot be quite the right representation of Stanley and Williamson's view, the previous quote notwithstanding. To see why, consider John again, who knows that that man has burning pants but who does not know that he himself has burning pants. John could conceivably entertain the thought that he himself has burning pants without thereby coming to know that he himself has burning pants. For example, John might consider, but not accept, the supposition that that man is himself. This would produce a scenario where all of the following conditions hold: (i) John stands in the knowledge-that relation to the Russellian proposition containing John and the property of burning pants under a demonstrative mode of presentation; and (ii) John also entertains this proposition under a first-personal mode of presentation; but (iii) John does not stand in the knowledge-that relation to this proposition under a first-personal mode of presentation.

Analogously, one would expect that there should be some possible scenario where all of the following conditions hold: (iv) there is some contextually relevant way  $w$  such that Hannah stands in the knowledge-that relation to the proposition that  $w$  is a way for Hannah to ride a bicycle under a non-practical mode of

presentation; and, (v) Hannah also entertains this proposition under a practical mode of presentation; but, (vi) there is no contextually relevant way  $w$  such that Hannah stands in the knowledge-that relation to the proposition that  $w$  is a way for Hannah to ride a bicycle under a practical mode of presentation.

According to the analysis of the truth conditions of knowledge-how ascriptions offered above, in such a scenario Hannah knows how to F, for given (iv) and (v) it follows that Hannah satisfies conditions (a) and (b) of this analysis. But I assume that Stanley and Williamson would not want to say that Hannah knows how to F in such a scenario, and they should allow that such a scenario is possible, given the analogy they draw between first-personal and practical modes of presentation. The moral here is that even though they occasionally appear to endorse it, we should not interpret Stanley and Williamson as making the conjunctive claim that Hannah knows how to ride a bicycle just in case there is some way  $w$  such that Hannah knows that  $w$  is a way for her to ride a bicycle *and* she entertains  $w$  under a practical mode of presentation. Rather, their view is that Hannah knows how to ride a bicycle just in case there is some way  $w$  such that Hannah knows that  $w$  is a way for her to ride a bicycle under a practical mode of presentation. It is true that if one knows that  $p$  under a practical mode of presentation, this entails that one knows that  $p$  *and* that one entertains  $p$  under a practical mode of presentation, but the reverse is not the case.

For such reasons, I take it that a more accurate representation of Stanley and Williamson's final account of knowledge-how (again assuming that practical modes of presentation are semantically relevant) is that the truth conditions of knowledge-how ascriptions conform to the following schema:

'S knows how to F' is true relative to a context  $c$  if, and only if, there is some contextually relevant way  $w$  for S to F such that:

(c) S stands in the knowledge-that relation to the proposition that  $w$  is a way for S to F, and

(d) In standing in this relation S entertains the proposition that  $w$  is a way for S to F under a practical mode of presentation.

Unlike the previous analysis, this analysis does require that for S to know how to F it must be the case that S knows, under a practical mode of presentation, that *w* is a way for S to F,<sup>5</sup> as opposed to merely requiring that S know that *w* is a way for S to F *and* that S entertain the proposition under a practical mode of presentation.

What is involved in entertaining a way *w* under a practical mode of presentation? Stanley and Williamson admit that giving a non-trivial characterisation of a practical mode of presentation is a substantive philosophical task, one they themselves do not claim to have offered.<sup>6</sup> However, they do suggest that entertaining a way under a practical mode of presentation will entail the possession of certain dispositions:

Thinking of a person as oneself entails being disposed to behave in certain ways, or form certain beliefs, given relevant input from that person ... Analogously, thinking of a way under a practical mode of presentation undoubtedly entails the possession of certain complex dispositions. It is for this reason that there are intricate connections between knowing-how and dispositional states (Stanley and Williamson 2001: 429–30).

Presumably, the idea here is that if there is some way *w* such that Hannah knows, under a practical mode of presentation, that *w* is a way for her to ride a bicycle,

---

<sup>5</sup> One could state the same view by replacing conditions (c) and (d) with just one condition of the form ‘S stands in the knowledge-that relation to the proposition that *w* is a way for S to F under a practical mode of presentation’ or ‘S knows, under a practical mode of presentation, that *w* is a way for S to F’. But the advantage of stating the analysis using (c) and (d) is that it emphasizes the important point that on this view the fact that S stands in the knowledge-that relation to the proposition that *w* is a way for S to F, is only a necessary and not a sufficient condition for S’s knowing how to F.

<sup>6</sup> Stanley and Williamson (2001: 429) do not think that the difficulty of this task is a problem for their view because they hold that the same is true of first-personal modes of presentation, and yet in “both cases ... one can provide an existence proof for such modes of presentation”. They argue that if, as is often assumed in the philosophy of language, there is a sound argument from (3) and (4) to the existence of first-personal modes of presentation, there should also be a sound argument from (1) and (2) to the existence of practical modes of presentation.

this entails that she possesses certain complex dispositions *related to the action of riding a bicycle*. However, Stanley and Williamson would not want to say that it entails that Hannah possesses the corresponding ability to ride a bicycle, for as we have seen, they deny that possessing the ability to F is a necessary condition of knowing how to F. But clearly the role of practical modes of presentation is to explain the natural thought that Hannah failing to know how to F is connected in some way to the fact that she fails to possess certain abilities or dispositions related to the action of riding a bicycle. On Stanley and Williamson's view this thought is explained by the fact that possessing some such dispositions is a necessary condition of entertaining a way of riding a bicycle under a practical mode of presentation.

*Bengson and Moffett's versions of sophisticated intellectualism*

Bengson and Moffett (2007) (see also Bengson et al., forthcoming) offer the following analysis of knowledge-how:

S knows how to F if and only if for some way *w* of F-ing such that:

- (e) S stands in the knowledge-that relation to the proposition that *w* is a way of F-ing, and
- (f) S minimally understands *w*.

According to Bengson and Moffett then, Shane knows how to bowl a googly just in case there is some way *w* such that he stands in the knowledge-that relation to the proposition that *w* is a way of bowling a googly, and he also minimally understands *w*.

There are some fairly minor differences between Bengson and Moffett's equivalence thesis and Stanley and Williamson's equivalence thesis. For example, Bengson and Moffett's condition (e) differs slightly from Stanley and Williamson's corresponding condition (c); whereas (c) refers to propositions of the form '*w* is a way for S to F', (e) refers to propositions of the form '*w* is a way of F-ing'. The most significant difference between the two accounts, however, is

that where Stanley and Williamson's condition (d) requires one to entertain the relevant proposition concerning a way to F under a *practical mode of presentation*, Bengson and Moffett's condition (f) requires one to *minimally understand* the relevant way.

This minimal understanding condition is what Bengson and Moffett appeal to when responding to the insufficiency objection to simple intellectualism.<sup>7</sup> For example, in response to the Tour De France case they would grant that Hannah satisfies condition (d) of their analysis—that is, that there is a way *w* of F-ing such that Hannah knows that *w* is a way of F-ing. However, Bengson and Moffett would argue that Hannah fails to minimally understand *w* and that, therefore, the case is not a counterexample to their analysis.

What does it mean to say that Hannah does not minimally understand that way of riding a bicycle? Bengson and Moffett explicate this notion of minimally understanding a way by appealing to a complex set of views on the nature of concepts and concept possession. I do not wish to examine these views in any detail, as to do so would require a lengthy discussion that would be tangential to my aims here. But Bengson and Moffett's basic strategy in replying to the

---

<sup>7</sup> This is not Bengson and Moffett's only motivation for including this condition in their analysis of knowledge-how. Their account of knowledge-how is motivated in large part by what they take to be a certain puzzle regarding knowledge-how attributions, namely, that some, but not all, attributions of the form 'S knows how to F' entail the corresponding attribution of the ability 'S has the ability to F' depending on what value we give to 'F'. Bengson and Moffett's solution to this puzzle appeals to their minimal understanding condition; very roughly, the idea is that for some actions—what they call 'select activities'—one can only minimally understand a way to perform that action if one possesses the ability to perform that action. Their examples of select activities include: adding *m* from *n*; subtracting *m* from *n*; inferring *q* from (*p* and *q*); inferring *q* from *p* and (if *p* then *q*). I am not convinced that this puzzle is genuine, as it seems to me that Bengson and Moffett do not establish that there is not still a gap between knowledge-how and ability, even in the case of their select activities (but see Bengson and Moffett 2007: 36–7 for further discussion).

insufficiency objection can be explained without going into the full details of these views.

With regards to the Tour De France case, I take it that Bengson and Moffett's reply to such an example relies, in the first instance, on the intuitive idea that while Hannah knows that that way is a way of riding a bicycle, her understanding of that way is, in some sense, inadequate or incomplete when compared to the understanding of, say, the cyclist who she sees riding that way. In Bengson and Moffett's terminology, the cyclist minimally understands this way and Hannah does not. This terminology is somewhat misleading, however, as to say that Hannah does not minimally understand that way could be interpreted as saying that she does not grasp or understand it *at all*, whereas their idea is only that her understanding of that way is somehow less than optimal when compared to the understanding of someone like the cyclist—that is, someone who actually knows how to ride a bicycle.

Bengson and Moffett would then analyse this difference—between Hannah's understanding and the cyclist's understanding—in terms of a difference between either their respective concepts of the way in question, and/or their mastery of their respective concepts of that way. For example, Bengson and Moffett (2007: 52) suggest that in a scenario like the Tour De France case Hannah's conception of that way would be less accurate or complete than the cyclist's conception of that way, as only the cyclist's conception could be used to guide someone in riding a bicycle. For Hannah's conception is based solely on demonstrative concepts of that way gained via her visual perception of someone riding that way, and they suggest that such concepts alone could not form a correct and complete conception of the kinaesthetic properties involved in that way of riding a bicycle.

However, as I stated above, the details of how Bengson and Moffett analyse their minimal understanding condition are not important for our purposes. What is important is the structure of their full account of knowledge-how. Bengson and Moffett hold that knowing how to F is not only a matter of there being some way *w* such that one stands in the knowledge-that relation to the proposition that *w* is

a way of F-ing. Rather, it is a matter of satisfying this condition and also possessing an adequate understanding of that way. In which case, their account of knowledge-how is clearly a version of sophisticated intellectualism.

It is also worth noting that Bengson and Moffett can appeal to their minimal understanding condition to explain the natural thought that Hannah's failing to know how to ride a bicycle is connected somehow to her failure to possess certain abilities or dispositions related to the action of riding a bicycle. For Bengson and Moffett clearly think that the best way of coming to minimally understand a way of riding a bicycle is to actually practise riding a bicycle oneself.<sup>8</sup>

## 1.6 Brogaard's Defence of Simple Intellectualism

Adopting some version of sophisticated intellectualism is the standard response to the insufficiency objection. But Berit Brogaard (forthcoming a, forthcoming b) has offered an interesting defence of a version of simple intellectualism to this objection. Brogaard offers the following analysis of knowledge-how, which is clearly a version of simple intellectualism:

S knows how to F if and only if there is some way *w* such that S stands in the knowledge-that relation to the proposition that *w* is how to F.

---

<sup>8</sup> Bengson and Moffett claim that for most actions, including actions like riding a bicycle, it is possible to know how to perform that action without possessing the ability to perform it (the exceptions being 'select activities' see fn. 7). It is important to note that there is no lurking inconsistency here. For suppose one claimed that the *only* way one could come to minimally understand a way of riding a bicycle was to practise riding a bicycle until one possessed the ability to ride a bicycle. Now, I think Bengson and Moffett would actually reject this claim, but they could in principle accept it, as it is perfectly consistent with the idea that subsequently one could lose the ability to ride a bicycle whilst retaining one's minimal understanding of a way to ride a bicycle.

According to Brogaard then, Shane knows how to bowl a googly just in case there is some way  $w$  such that he stands in the knowledge-that relation to the proposition that  $w$  is how to bowl a googly.

Brogaard's analysis is a version of simple intellectualism for it tells us that standing in the knowledge-that relation to the right kind of proposition is both a necessary and *sufficient* condition for knowing how to F. Of course, Brogaard is well aware of the insufficiency objection to simple intellectualism, as she writes:

Suppose I have never practiced playing the piano but that I have taken numerous theory lessons. There is then an  $x$  such that I know  $x$  is how to play the piano. Still, it would seem that someone could correctly claim that I don't know how to play the piano. Likewise, if Mary—a mono-lingual speaker of English—sees Danny curse out his cousin in Italian, she might correctly say [whilst pointing] '*that* is how to curse someone in Italian'. Yet someone could correctly say 'Mary doesn't know how to curse out someone in Italian'. After all, Mary doesn't even speak Italian (Brogaard forthcoming b: 47).

However, Brogaard does not respond to such cases by adopting some version of sophisticated intellectualism. Rather, Brogaard claims that the lesson of such cases is that 'S knows how to F' attributions are ambiguous between two readings, and that this ambiguity provides us with a response to the insufficiency objection:

There is, however, a straightforward reply to these sorts of objections ... knowledge-how ascriptions that embed infinitive clauses are ambiguous between a reading that requires that the subject possess an ability (first-person) and a reading that does not require that the subject possess an ability (third person). For example, 'John knows how to play the piano' may be read as saying that John knows how JOHN *may* play the piano, or as saying that John knows how ONE *may* play the piano. So, on the analysis offered here, 'John knows how to play the piano' can be read as saying that there is a  $w$  such that John knows that  $w$  is how John may play the piano or as saying that there is a  $w$  such that John knows that  $w$  is how one may play the piano. If John has never practiced playing the piano, it is false that there is a  $w$  such that John knows that  $w$  is how John may play the piano' but it may well be true that there is a  $w$  such that John knows that  $w$  is how one may play the piano. So, 'John knows how to play the piano' is false when given the first reading but it may well be true when given the second reading. Likewise, if Mary doesn't speak Italian, then it will be false that there is a  $w$  such that Mary knows

that  $w$  is how Mary may curse out someone in Italian but it may be true that there is a  $w$  such that Mary knows that  $w$  is how one may curse out someone in Italian (Brogaard forthcoming b: 47).

Let us consider then how Brogaard would analyse the Tour De France case. To begin with, Brogaard would claim that (1) can be disambiguated in two different ways:

(1) Hannah knows how to ride a bicycle.

In particular, Brogaard would claim that (1) can be disambiguated as either (1a) or (1b):

(1a) There is a way  $w$  such that Hannah knows that  $w$  is how *Hannah* may ride a bicycle.

(1b) There is a way  $w$  such that Hannah knows that  $w$  is how *one* may ride a bicycle.

Brogaard's response to the Tour de France scenario, in short, would be to claim that in such a scenario (1a) is false and so, given that (1a) is a legitimate interpretation of (1), there is a good sense in which (1) is false, which explains our intuition that Hannah does not know how to ride a bicycle.

Note that the claim that (1) is ambiguous between (1a) and (1b), does not actually play any role in Brogaard's response to the insufficiency objection. The important claim with respect to the Tour De France scenario is simply that (1a) is a legitimate interpretation of (1), and that (1a) is false in this scenario. The further claim that (1b) is also a legitimate interpretation of (1) is superfluous to Brogaard's response to this kind of case. And this is probably a good thing because (1b) does not look to be a good interpretation of (1). It is true that the syntactic structure of (1) allows for it to be disambiguated as either (1a) or (1b), for there is a covert pronoun in the structure of (1) which can either receive its

interpretation from the subject of the main clause (i.e. Hannah) or can be interpreted as ‘one’.<sup>9</sup> But (1b) looks to be implausible as a semantic interpretation of (1), for if (1b) is a legitimate interpretation of (1) then it would follow that there is a legitimate sense in which anyone who merely knows that there is some way that people ride bicycles thereby knows how to ride a bicycle.

Brogaard’s response to the insufficiency problem requires that (1a) be a legitimate interpretation of (1), and that (1a) be false in a scenario like the Tour De France case. But why think that (1a) is false in the Tour De France scenario? For note that the interpretation of (1) as (1a) is very close, if not equivalent, to Stanley and Williamson’s initial account of the truth conditions of (1), which tells us that Hannah knows how to ride a bicycle just in case there is some way such that Hannah knows that *w* is a way for her to ride a bicycle. To see this compare (1a) with (1c):

(1a) There is a way *w* such that Hannah knows that *w* is how *Hannah* may ride a bicycle.

(1c) There is a way *w* such that Hannah knows that *w* is a way for *Hannah* to ride a bicycle.

Stanley and Williamson’s initial account of knowledge-how tells us that an assertion of (1) is equivalent to an assertion of (1c), and in turn (1a) and (1c) look to be equivalent in meaning. That is, replacing ‘*w* is how Hannah may ride a bicycle’ in (1a) with ‘*w* is a way for Hannah to ride a bicycle’ does not appear to change the meaning of (1a) at all. For Stanley and Williamson (2001: 424–5) interpret ‘to ride a bicycle’ in (1c) as expressing something like ‘can ride a bicycle’, and Brogaard also clearly uses ‘may’ in (1a) to express something like the ability sense of ‘can’ (rather than the deontic sense of ‘may’ where it expresses something like permissibility).

---

<sup>9</sup> See Chapter §2.1 for further discussion.

But while (1a) and (1c) look to be equivalent Stanley and Williamson disagree with Brogaard about what the truth value of (1a/1c) would be in a scenario like the Tour de France case. Stanley and Williamson (*ibid.* 428) explicitly state that such a scenario would be a case where (1c) would be *true*. As we have seen, this is why Stanley and Williamson think that their initial account of knowledge-how ascriptions is inadequate, and why they suggest that a more adequate account of knowledge-how ascriptions must appeal to practical modes of presentation—in either the semantics or the pragmatics of such ascriptions. Brogaard, however, thinks that the Tour De France case would be a scenario where (1a) would be false.

Bearing in mind that Brogaard uses ‘may’ to express something like the ability sense of ‘can’, I take it that her rationale for thinking that (1a) is false in a scenario like the Tour De France case goes like this: if there were some way *w* such that Hannah knows that *w* is a way how Hannah may ride a bicycle, this would entail that *w* is a way how Hannah herself may ride a bicycle. But, Hannah clearly lacks the ability to ride a bicycle, so there is no way *w* such that *w* is a way how Hannah may ride a bicycle and, therefore, there is no way *w* such that Hannah knows that *w* is a way how she may ride a bicycle.

## 1.7 Assessing the Insufficiency Objection

We are now in a proper position to assess the question of whether the insufficiency objection shows us that intellectualism is false, for we have seen that to answer this question we actually have to address two distinct questions. The first insufficiency question is: does the insufficiency objection show us that simple intellectualism is false? And the second insufficiency question is: does the insufficiency objection show us that sophisticated intellectualism is false?

In answer to the first question, I think the insufficiency objection strongly suggests that simple intellectualism is false. As we have seen, Brogaard gives an interesting defence of simple intellectualism that relies on the fact that the standard cases used to motivate this objection are cases where a subject not only

fails to know how to F but they also fail to possess the ability to F. For example, Hannah not only does not know how to ride a bicycle, but there is also a clear sense in which she does not possess the ability to ride a bicycle. Brogaard's basic strategy for handling such cases then is to equate knowing how to F with knowing that p, for a proposition p such that knowing that p entails that one may/can F. In particular, Brogaard claims that S knows how to F if and only if there is some way *w* such that Hannah knows that *w* is a way how S may/can F.

This strategy may appeal when responding to examples like the Tour de France case, because this is a scenario where intuitively Hannah not only fails to know how to ride a bicycle, but she also fails to possess the ability to ride a bicycle. In which case, it is relatively easy to motivate the idea that there is no way *w* such that Hannah knows that *w* is how she can ride a bicycle. But the problem is that the insufficiency objection can also be motivated using cases where it seems clear that a subject S both fails to know how to F and that there is a way *w* such that S knows that *w* is a way how S can F.<sup>10</sup>

For example, suppose that a chess grandmaster is coaching two of her students, Sacha and Boris, as they near the end of a game of chess. The grandmaster says to Sacha: "Sacha there is no way for Boris to win from here but there is a way that *you* can win from here". Sacha thereby comes to know that there is a way that Sacha can win the game, but Sacha still does not know how to win the game.<sup>11</sup>

---

<sup>10</sup> In fact, even in the Tour De France case I am inclined to agree with Stanley and Williamson that Hannah would know that that way (i.e. the way the cyclist on the television is riding his bicycle) is a way how Hannah may/can ride a bicycle. At least, I think it is natural to imagine the case in such a way that Hannah would possess this knowledge-that. If we suppose that Hannah is such that she could learn how to ride a bicycle in that way—she has the relevant limbs, motor control, coordination etc.—and that Hannah knows that she is such a person, then it seems to me that she would not only know that that way is how one may ride a bicycle, but she would also know that that way is how she may ride a bicycle. In which case, the Tour De France case is still a counterexample to Brogaard's analysis of knowledge-how.

<sup>11</sup> Thanks to Daniel Stoljar for suggesting this example.

Such a scenario is surely possible, for Sacha may know that the way the grandmaster is referring to is a way she can win the game and yet, unlike the grand master, she may still fail to entertain or grasp this way in the manner required for her to also know how to win the game. Indeed, it is tempting to say that Sacha would still fail to know how to win the game because she only entertains this way under some non-practical mode of presentation.

Similarly, recall Snowdon's (2003) case of the man who has the ability to get out of a room, but who does not yet know how to get out of the room because he has not yet explored it. Suppose that a friend calls this man on his mobile phone and says: "Don't worry there is a way you can get out of the room". The man now knows that there is a way that *he* can get out of the room but, I submit, intuitively he still does not know how to get out of the room given that he has not yet explored the room.

Such cases appear to show that the insufficiency objection can be made even against Brogaard's version of simple intellectualism. Perhaps Brogaard could offer some response to these cases. But, in answer to our first insufficiency question, I think it is safe to say that the insufficiency objection at least strongly suggests that any plausible version of simple intellectualism will turn out to be false.<sup>12</sup>

---

<sup>12</sup> Perhaps Brogaard could respond to such cases by appealing to the context sensitivity of the modals 'may' or 'can'. For example, in a context where we are discussing what Hannah can do given the proper instruction it would be correct to say, "Hannah can ride a bicycle". But in other contexts such an assertion would be incorrect, for example, in a context where we are discussing what Hannah can do prior to receiving any instructions or lessons. Similarly, we will happily judge that someone can F *in view of such-and-such*, whilst denying that they can F *in view of so-and-so* (for discussion see Kratzer 1977). One might argue then that the fact that we can elicit the intuition that Hannah knows that that way is a way how she may/can ride a bicycle, or that Sacha knows that there is a way she can win the game, merely reflects the context sensitivity of 'may' and 'can'. But appealing to the context sensitivity of modals like 'may' and 'can' will only help to defend Brogaard's account of knowledge-how if our intuitions as to whether S knows how to F 'line up' with our intuitions as to whether there is some way *w* such that S knows that *w* is a way for S to F. And examples like the chess case suggest that this is not the case. (Thanks to Berit

However, in answer to our second insufficiency question, it is clear that the insufficiency objection does not show us that sophisticated intellectualism is false. As we have seen, many intellectualists simply grant that merely standing in the knowledge-that relation to some relevant proposition concerning a way to F does not constitute a sufficient condition for knowing how to F. Sophisticated intellectualists claim that knowing how to F is not only a matter of standing in the knowledge-that relation to the right proposition but it is also a matter of satisfying some further condition.

Of course, one might argue that particular versions of sophisticated intellectualism do not offer an adequate response to the insufficiency objection. Some philosophers have argued that this is the case with respect to Stanley and Williamson's account of knowledge-how, on the grounds that they tell us little about what practical modes of presentation are (for example, see Noë 2005, Rosefeldt 2004). But even if these objections to Stanley and Williamson's specific version of sophisticated intellectualism are reasonable, they do not show us that there is some principled problem with sophisticated intellectualism in general.

I suspect that many critics of intellectualism consider the insufficiency objection to be an objection to intellectualism in general, because they think that the only way for the intellectualist to respond to examples like the Tour de France case is to appeal to some Rylean condition on knowing how to F. For example, Tobias Rosefeldt (2004: 375) seems to raise this kind of concern when, after considering Stanley and Williamson's appeal to the notion of a practical mode of presentation as a way of responding to the insufficiency objection, he concludes that their "talk about practical modes of presentation is simply disguised talk about abilities."

There are two things to say about this kind of concern with intellectualist responses to the insufficiency objection. First, it is not at all obvious that the right

---

Brogaard for suggesting that one might appeal to the context sensitivity of the modals 'may' and 'can' as a way of responding to the concerns raised above for her defence of simple intellectualism, any errors in presenting this response in the best light are my own.)

diagnosis of why Hannah fails to know how to ride a bicycle is simply that she lacks certain abilities or capacities. As Stanley and Williamson's ski instructor and pianist cases suggest, having the ability to F is not a necessary condition of knowing how to F; and examples like Bengson and Moffett's salchow case suggest that it is not even a sufficient condition for knowing how to F. It may well turn out then that the best explanation of why Hannah fails to know how to ride a bicycle is not simply that she fails to possess certain abilities, but that she does not entertain certain propositions under the right mode of presentation, or that she does not possess an adequate conception of the way to ride a bicycle.

Second, even if the intellectualist does have to directly appeal to certain abilities or dispositional states in order to respond to the insufficiency objection this is not inconsistent with some version of sophisticated intellectualism being true. For example, suppose that in response to the insufficiency objection an intellectualist offered the following analysis of knowledge-how:

S knows how to F if, and only if:

(g) S stands in the knowledge-that relation to p (for some relevant proposition p concerning a way to F), and

(h) S possesses the ability to F.

This analysis of knowledge-how includes the Rylean condition that one possess the ability to F, but it is still a version of sophisticated intellectualism because it also includes the intellectualist condition (g). A neo-Rylean might argue that all that this shows is that (g) is redundant on the grounds that (h) states not only a necessary but also a sufficient condition for knowing how to F. However, as we have seen, there are examples that strongly suggest that (h) does not state a sufficient condition for knowing how to F. Of course, we have also seen that there are examples that strongly suggest that (h) is not a necessary condition for knowing how to F, so I am not recommending the above analysis to intellectualists. The point is simply that there is no obvious route from the insufficiency objection to simple intellectualism to the truth of some form of Ryleanism, as opposed to some form of sophisticated intellectualism. The

insufficiency objection strongly suggests that there is no plausible version of simple intellectualism, but it does not show us that there can be no plausible version of sophisticated intellectualism. Hence, we must look elsewhere if we are to find a conclusive argument against intellectualism in general. In Chapter 3 I will present what I take to be the most powerful arguments for rejecting intellectualism, arguments that, if sound, undermine not only simple but also sophisticated intellectualism. Now, however, I want to consider the philosophical implications of sophisticated intellectualism. In particular, I will consider the implications of Stanley and Williamson's version of sophisticated intellectualism for an important position in the philosophy of mind.

## Chapter 2 The Ability Hypothesis and Intellectualism

---

What follows for the ability hypothesis reply to the knowledge argument (Jackson 1982, 1986) if intellectualists are right that knowledge-how is just a kind of knowledge-that? The obvious answer is that the ability hypothesis is false. For the ability hypothesis says that when Mary—Frank Jackson’s super-scientist—sees red for the first time she gains only knowledge-how and not knowledge-that. In this chapter I argue that this obvious answer is wrong: a version of the ability hypothesis might be true even if knowledge-how is a kind of knowledge-that. In §2.1 I briefly return to Stanley and Williamson’s (2001: 1) account of knowledge-how as “simply a species of propositional knowledge”. In §2.2 I set out the ability hypothesis and explain Stanley and Williamson’s claim that it is a consequence of their account that the ability hypothesis fails. In §2.3 I demonstrate that this claim is not quite right. Stanley and Williamson’s account of knowledge-how is inconsistent with the combination of everything said by standard statements of the ability hypothesis. However, we can restate the core claims of the ability hypothesis—that Mary only gains new knowledge-how and not knowledge-that—within their account of knowledge-how as a species of knowledge-that. In the remainder of the chapter (§§2.4–2.6) I examine the implications of this result for both critics and proponents of the ability hypothesis.

### 2.1 The New Knowledge-how

As discussed in Chapter 1, Gilbert Ryle famously argued that there was a fundamental distinction in kind between knowledge-how and knowledge-that. Proponents of the ability hypothesis standardly appeal to this Rylean view of knowledge-how when stating their reply to the knowledge argument. As we have seen, however, Stanley and Williamson reject the orthodox notion that there is a fundamental distinction between knowledge-how and knowledge-that. To

understand why, consider once again a knowledge-how ascribing sentence, such as:

(1) Mary knows how to ride a bicycle.

Stanley and Williamson's main argument for intellectualism rests on the claim that ascriptions like (1) are constructions where the complement of the verb 'knows' is an *embedded question* formed by the question word 'how' and the infinitive 'to ride a bicycle'. Stanley and Williamson claim then that the syntactic structure (1) is of a kind with (2)–(5):

(2) Mary knows where to find a dollar.

(3) Mary knows whom to call for help in a bush fire.

(4) Mary knows which prize to look for.

(5) Mary knows why to vote for Garrett.

In particular, Stanley and Williamson claim that, abstracting away from certain details, the standardly accepted syntactic structure of (1)–(5) is as follows:

(1') Mary knows [how PRO to ride a bicycle *t*].

(2') Mary knows [where PRO to find a dollar *t*].

(3') Mary knows [whom PRO to call *t* for help in a fire].

(4') Mary knows [which prize PRO to look for *t*].

(5') Mary knows [why PRO to vote for Garrett *t*].

‘PRO’ here is a covert or phonologically null pronoun that is the subject of the infinitive clause, and the occurrences of ‘*t*’ indicate the *traces* of the movement of the phrases ‘how’, ‘where’, ‘whom’, ‘which prize’, and ‘why’ in the structure of (1’–(5’) respectively.

The standard semantics for ascriptions like (2’–(5’) tells us, roughly, that each ascription is true just in case Mary knows an answer to the respective embedded question, where to know an answer to the embedded question is to stand in the knowledge-that relation to some proposition *p* that answers that question. If we apply the same semantics for embedded questions to (1’) then we get the result that (1’) will be true if and only if Mary knows that *p*, for some proposition *p* that is an answer to the embedded question in (1’).

Stanley and Williamson suggest that for ascriptions of the form ‘S knows how to F’ any legitimate answer to the respective embedded question will be a proposition of the form ‘*w* is a way for S to F’.<sup>13</sup> Hence, they hold that if we follow the standard semantics for embedded questions then (1’) will be true if and only if there is some way *w* such that Mary stands in the knowledge-that relation to the proposition that *w* is a way to F. More generally, Stanley and Williamson’s initial suggestion is that the truth conditions of knowledge-how ascribing sentences conform to the following schema:

‘Mary knows how to ride a bicycle’ is true if and only if there is some way *w* for Mary to ride a bicycle such that Mary stands in the knowledge-that relation to the proposition that *w* is a way for Mary to ride a bicycle.

---

<sup>13</sup> Why ways for S to ride a bicycle rather than for someone else? The reason is that Stanley and Williamson interpret ‘PRO’ in, for example, (8) as receiving its subject from the main clause so that ‘PRO’ is interpreted as ‘Mary’. However, like Brogaard, Stanley and Williamson (2001: 424–5) note that ascriptions like (8) also allow for an interpretation whereby ‘PRO’ is interpreted as ‘one’. Furthermore, they note that the infinitive in (8), ‘to ride a bicycle’, not only can be interpreted as having ‘can’-like force (which is the interpretation they focus on), but it can also be interpreted as having ‘ought’-like force.

According to this view then, (1') expresses an existential generalization over propositions; it says that there exists at least one proposition of the form '*w* is a way for Mary to F' such that Mary knows that proposition. Given the accepted picture of both the syntactic structure of ascriptions like (2)–(5) and the semantic interpretation of (2')–(5'), Stanley and Williamson (2001: 431) take their “view of ascriptions of knowledge-how to be the default view.” And the import of this view is that “to say that someone knows how to F is always to ascribe to them knowledge-that” (*ibid.* 426).<sup>14</sup>

Of course, in response to the insufficiency objection, Stanley and Williamson also include in their final account of knowledge-how the further condition that in standing in the knowledge-that relation to some proposition of the '*w* is a way for Mary to F', Mary must also entertain this proposition under *a practical mode of presentation*. More generally, they hold that, if modes of presentation are semantically relevant,<sup>15</sup> the truth conditions of knowledge-how ascribing sentences conform to the following schema:

‘S knows how to F’ is true if, and only if, there is some way *w* for S to F such that:

(a) S stands in the knowledge-that relation to the proposition that *w* is a way for S to F, and

---

<sup>14</sup> See Brown (1970) for a related precursor of this argument.

<sup>15</sup> For ease of exposition, I will treat practical modes of presentation as if they are part of the semantics of ‘S knows how to F’ ascriptions in this discussion. Nothing that I go on to argue, however, would be essentially affected were we to treat practical modes of presentation as part of the pragmatics of using such ascriptions. Also, as we saw in §1.1, on Stanley and Williamson’s view the truth conditions of knowledge-how ascriptions like (1') are relativized to conversational contexts. This is because changes in conversational context can change what counts as a legitimate answer to an embedded question. But as this feature of their view is also inessential for our discussion I have ignored it here.

(b) In standing in this relation S entertains the proposition that  $w$  is a way for S to F under a practical mode of presentation.

And recall that on Stanley and Williamson's preferred development of this view, a practical mode of presentation is a way of thinking of a Russellian proposition, rather than a constituent of a fine-grained or Fregean proposition (throughout this chapter the word 'proposition' should be understood as referring to the Russellian notion of a proposition).

The analysis above constitutes the core of Stanley and Williamson's account of knowledge-how. As we have seen, one might object to this account in various ways, but the truth or falsity of their account is not my concern in this chapter. My concern here is with the consequences of their account for the ability hypothesis. Thus, in the following discussion I assume that Stanley and Williamson's account of knowledge-how is true and then consider what follows (if anything) for the ability hypothesis reply to the knowledge argument.

## 2.2 The Ability Hypothesis

Frank Jackson's knowledge argument (KA) invites us to consider the following thought experiment. Mary is a brilliant scientist who has spent all her life in a black and white room studying all of the (completed) natural sciences; her only access to the outside world is via images on a black and white TV screen. One fateful day, Mary steps outside her room and sees red for the first time. It seems that, upon her release, Mary learns something about the nature of the world, namely, what it is like to see red. However, by hypothesis, Mary knows every physical truth there is to know prior to her release. Therefore, Jackson argues, there is at least one non-physical truth, that is, the truth Mary comes to know upon her release. From this initial conclusion, and taking physicalism to be

committed to the thesis that all truths are physical truths or are a priori entailed by physical truths,<sup>16</sup> Jackson concludes that physicalism is false.

The ability hypothesis offers an elegant reply to KA. Mary does not gain any new knowledge-that upon her release but merely knowledge-how. Proponents of the reply often take the further step of identifying Mary's new knowledge-how with new abilities that she gains when released—hence the name of the hypothesis. We shall come to this further step in due course (see §2.4). Potential candidates for Mary's new knowledge-how include: knowledge how to imagine experiences of red; knowledge how to recognize experiences of red; and knowledge how to remember experiences of red. For our purposes, the exact knowledge-how Mary gains upon release is not crucial.<sup>17</sup> What *is* crucial is that the core of the ability hypothesis centres largely on just two claims, a negative claim about what Mary does not gain after her release, and a positive claim about what she does gain:

(NEG) Upon release Mary does not gain any new knowledge-that.

(POS) Upon release Mary gains new knowledge-how.

NEG establishes that Jackson's thought experiment does not describe a possibility that undermines physicalism. Mary does not learn any new truth about the world upon release because she does not gain any new propositional knowledge. POS serves to explain our deep-seated intuition that Mary does nevertheless learn something, and thereby comes to know something, upon her release. The truth of POS is consistent with the truth of NEG because Mary's new knowledge-how is *not* propositional knowledge. As Laurence Nemirow says, "Mary's knowledge gap is practical not propositional" (Nemirow 1995: 36). For

---

<sup>16</sup> This particular assumption is highly controversial. For his position on this issue see Jackson (2004: 423–6).

<sup>17</sup> From this point forward I will use 'knowledge-how to imagine experiences of red' as a placeholder for one's favoured version of some list like this.

those who take the further step just mentioned, the gap in Mary's knowledge is not propositional because her new knowledge-how is to be identified with new abilities, as opposed to new knowledge-that (a point I will return to in §2.4).

What is the relationship between this response to KA and Stanley and Williamson's account of knowledge-how? As suggested earlier, the obvious answer is that Stanley and Williamson's account of knowledge-how is inconsistent with the ability hypothesis. Stanley and Williamson endorse this obvious answer in their critique of David Lewis's (1998) version of the ability hypothesis:

According to Lewis, the correct account of Jackson's knowledge argument is that Mary does not gain new knowledge-that when she leaves her black and white room, but only knowledge-how. In particular, she gains knowledge how to recognize, remember, and imagine experiences of red. Our discussion shows, however, that Lewis's account is incorrect. Knowing how to imagine red and knowing how to recognize red are both examples of knowledge-that. For example,  $x$ 's knowing how to imagine red amounts to knowing a proposition of the form ' $w$  is a way for  $x$  to imagine red', entertained under a guise involving a practical mode of presentation of a way (Stanley and Williamson 2001: 442).

Are Stanley and Williamson right? Undeniably, there is a conflict between their account of knowledge-how and the letter of *classic statements* of the ability hypothesis, at least of the kind presented by Lewis (1988) and Nemirow (1980, 1990). After all, the classic statements of the ability hypothesis say that Mary does not gain any new knowledge-that because she merely gains new knowledge-how, while Stanley and Williamson's main claim is that knowledge-how is a form of knowledge-that. The interesting question, however, is whether we can give a modified statement of the ability hypothesis that, on the one hand, is consistent with Stanley and Williamson's account of knowledge-how and, on the other hand, retains the spirit of these classic statements.

One obvious modification of the ability hypothesis would be to concede that Stanley and Williamson's account of knowledge-how as a kind of knowledge-that,

makes POS incompatible with NEG, and so to deny POS, as it stands, and replace it with:

(POS') Upon release Mary gains new abilities.

Stanley and Williamson (2001: 442–3) identify this “fallback position” for proponents of the ability hypothesis themselves, and they argue that it fails. I discuss these arguments in §2.4. For now it will suffice to note that such a response to Stanley and Williamson is very unattractive for reasons they do not discuss. If our modified ability hypothesis no longer affirms POS, then it no longer speaks to our intuition that Mary learns something upon release and thereby gains knowledge. Being able to say that Mary *learns* how to do something, and thereby gains *knowledge-how*, seems to me to be a clearly essential feature of the ability hypothesis. Consequently, I will now develop a very different response to Stanley and Williamson. This response claims that, despite all appearances to the contrary, we can accept Stanley and Williamson’s account of knowledge-how and *still* consistently assert versions of both NEG and POS.

### 2.3 Mary’s ‘New’ New Knowledge-how

Here is a reply to KA that is not only consistent with, but relies upon, Stanley and Williamson’s account of knowledge-how—specifically the issues surrounding sentences (1) and (2). Prior to her release, Mary knows that *w* is a way for her to imagine an experience of red, but she does not know how to imagine an experience of red. This is because Mary does not know the proposition that *w* is a way for her to imagine an experience of red under a practical mode of presentation. After her release, Mary comes to know that *w* is a way for her to imagine an experience of red under a practical mode of presentation, and hence gains new knowledge how to imagine an experience of red. In gaining this knowledge-how Mary does not gain knowledge of any new proposition. Mary only comes to know, under a practical mode of presentation, a proposition that she already knew under some other mode of presentation (recall sentences (3) and

(4)). Just as Nemirow (1995) says, the gap in Mary's knowledge was practical not propositional. Physicalism is not refuted by KA, because if Mary does not come to know any new propositions then she does not come to know any new non-physical truths about the world.

Under what non-practical mode of presentation could pre-release Mary have known that *w* is a way for her to imagine an experience of red? Stanley and Williamson themselves suggest one possibility:

If someone entertained a way of riding a bicycle by possessing a complete physiological description of it, that might also give them *de re* knowledge of that way, though not under a practical mode of presentation (2001: fn. 429).

If someone in this scenario could know (*de re*) of some way that it is a way for them to ride a bicycle, even though they do not know how to ride a bicycle, then presumably Mary in her lab (with her knowledge of all of the completed natural sciences) could likewise know that *w* is a way for her to imagine an experience of red, even if she did not know how to imagine an experience of red. In both cases, Mary knows the relevant proposition under something like a theoretical mode of presentation. Alternatively, pre-release Mary might have known that *w* is a way for her to imagine an experience of red under a demonstrative mode of presentation. What matters is that Stanley and Williamson's account explicitly allows for the possibility that pre-release Mary could know that *w* is a way for her to imagine red under *some* mode of presentation other than a practical mode of presentation. The above response to KA simply employs this feature of Stanley and Williamson's account to explain how Mary could gain knowledge-how without gaining knowledge of a new proposition. Let us call this response to KA the ability hypothesis\*, or AH\*.

Is AH\* really a version of the ability hypothesis? Perhaps the answer is not immediately clear. For AH\* reveals an ambiguity in standard statements of the ability hypothesis with respect to NEG:

(NEG) Upon release Mary does not gain any new knowledge-that.

We can interpret NEG as either of the two following claims:

(NEG<sub>1</sub>) Upon release Mary does not gain knowledge of any new proposition.

(NEG<sub>2</sub>) Upon release Mary does not come to be in any new state of propositional knowledge.

With this disambiguation now clearly in view we can legitimately identify different versions of the ability hypothesis. Any version of the ability hypothesis will, of course, have to affirm NEG<sub>1</sub>, since the claim that Mary does not come to know any new proposition is the claim the ability hypothesis uses to block Jackson's conclusion that Mary comes to know a new truth about the world. So, we have just two versions of the ability hypothesis to choose between: a version that endorses POS and only NEG<sub>1</sub>, and a version that endorses POS and both NEG<sub>1</sub> and NEG<sub>2</sub>.

Ability Hypothesis (Version 1)

(NEG<sub>1</sub>) Upon release Mary does not gain knowledge of any new proposition.

(POS) Upon release Mary gains knowledge-how.

Ability Hypothesis (Version 2)

(NEG<sub>1</sub>) Upon release Mary does not gain knowledge of any new proposition.

(NEG<sub>2</sub>) Upon release Mary does not come to be in any new state of propositional knowledge.

(POS) Upon release Mary gains knowledge-how.

It should now be clear that AH\* is just an instance of version 1 of the ability hypothesis, as it asserts POS and NEG<sub>1</sub> and it denies NEG<sub>2</sub>—for if Mary knows an old proposition under a new mode of presentation then she is in a new state of propositional knowledge. Thus, once we disambiguate NEG we see that AH\* is a legitimate version of the ability hypothesis. This is enough to show that Stanley and Williamson are not quite right to suppose that their account of knowledge-how is inconsistent with the ability hypothesis. We can consistently accept their account of knowledge-how whilst replying to KA with a version of the ability hypothesis like AH\*.

To clarify this position, and to reveal some of its implications, I will now look at two features of AH\* that distinguish it from classical statements of the ability hypothesis. Understanding these features improves our understanding of the ability hypothesis as a general form of reply to KA. We will also see that AH\* has virtues other than it simply being consistent with Stanley and Williamson's account of knowledge-how.

## 2.4 Where are the Abilities?

Classic statements of the ability hypothesis identify Mary's new knowledge-how with new abilities. In so doing, they go beyond the core claims NEG and POS. Why do they make this further step? The reason is that there is a *prima facie* challenge of explaining how NEG and POS are compatible. Imagine a person who claimed (for whatever reason) that in a particular situation Fred gains knowledge-*why* but no knowledge-*that*. Obviously, it would be incumbent on this person to explain how Fred could come to know why something is the case (or why to do something, etc.), without thereby also coming to know that something is the case, i.e. without gaining new knowledge-*that*. Likewise, the proponent of the ability hypothesis has to explain how Mary can gain new knowledge-how without gaining new knowledge-*that*.

The version of the ability hypothesis I developed in the previous section, AH\*, answers this challenge in a particular way. AH\* explains how NEG is compatible

with POS by disambiguating NEG to get NEG<sub>1</sub> and NEG<sub>2</sub>, and then pointing out that POS is clearly compatible with NEG<sub>1</sub>. On the other hand, classic statements of the ability hypothesis answer this challenge in a different way. Their answer comes in two parts. First, they identify knowledge-how with abilities. More specifically, they claim that Mary's knowing how to imagine an experience of red is identical to her ability to imagine an experience of red. Second, they claim that these abilities are what I will call 'mere abilities', that is, abilities that are utterly distinct from any propositional knowledge.<sup>18</sup> Together, these two claims provide an explanation of how NEG and POS are compatible. NEG, the claim that Mary does not gain any new knowledge-that, is compatible with POS, the claim that Mary gains new knowledge-how, because Mary's knowledge-how is identified with mere abilities as opposed to propositional knowledge.

So, the version of the ability hypothesis I have given is different in an important respect from classic statements of the ability hypothesis, namely, it offers a different explanation of how NEG is compatible with POS. What is the reason for this difference? The reason is that no version of the ability hypothesis that is consistent with Stanley and Williamson's account of knowledge-how can say that NEG is compatible with POS *because* Mary's new knowledge-how is identical to mere abilities. If knowledge-how is a species of knowledge-that, then Mary's new knowledge-how cannot be identical to abilities that are utterly distinct from any knowledge-that. That is why AH\* answers the challenge of explaining how NEG is compatible with POS in a different way. In short, Stanley and Williamson's account of knowledge-how is inconsistent with the explanation of how POS is compatible with NEG given by classic statements of the ability hypothesis.

---

<sup>18</sup> Lewis does allow that sometimes one's ability to F is simply identical to one's knowledge that *p*, for some proposition *p* (see Lewis 1998: 459 on having the ability to open the combination lock on a bank vault). In the case of knowing what it is like to have an experience, however, Lewis thinks of the relevant abilities as not involving propositional knowledge at all.

We could summarize these points by noting that the following three claims form an inconsistent triad:

(Stanley and Williamson's Main Claim) Knowledge-how is a species of knowledge-that.

(The Identity Claim) Mary's knowledge how to imagine experiences of red is identical to her ability to imagine experiences of red.

(The Distinctness Claim) Mary's ability to imagine experiences of red is utterly distinct from any knowledge-that.

Classic statements of the ability hypothesis explain how NEG is compatible with POS by appealing to both the identity and distinctness claims. This is the grain of truth in Stanley and Williamson's claim that their account of knowledge-how is inconsistent with the ability hypothesis. Their main claim is inconsistent with the standard explanation of how NEG is compatible with POS because it is inconsistent with the conjunction of the identity and distinctness claims. However, AH\* is a version of the ability hypothesis that relies on an alternative explanation of how NEG is compatible with POS. This explanation is consistent with Stanley and Williamson's main claim.

There is still a question remaining, however, concerning Mary's new knowledge-how. In this context, where we are assuming that Stanley and Williamson's main claim is correct, what is the relationship between Mary's new knowledge-how and her new abilities? The relationship will differ depending on whether we choose to retain just the identity claim, or just the distinctness claim, or neither of these claims. I will consider just the first two options.

If we deny the distinctness claim then we can still identify Mary's new knowledge-how with her new abilities. Mary's new ability to imagine experiences of red is identical to Mary's new knowledge that *w* is a way for her to imagine experiences of red, under a practical mode of presentation. On the other hand, if we deny the

identity claim we can maintain the idea that Mary's new abilities are mere abilities. This option may appeal if one is inclined to regard abilities and propositional knowledge as utterly distinct categories. Stanley and Williamson themselves argue, as noted in Chapter 1, that knowledge-how at least does not entail ability, because it is possible to know how to F whilst not possessing the ability to F. Recall that they cite the case of "a master pianist who loses both of her arms in a tragic car accident" (Stanley and Williamson 2001: 416). Stanley and Williamson claim that this person would still know how to play the piano despite having lost the ability to play the piano. If a similar scenario could be conceived for Mary—whereby she knows how to imagine an experience of red but does not have the ability to imagine an experience of red—we would have to reject the identity claim (see Alter 2001 for a development of this idea).

What would the relationship be between Mary's new knowledge-how and her new abilities if not identity? One possibility is that there is a conceptual or constitutive connection, other than identity, between Mary's new knowledge-how and her new abilities. For example, Stanley and Williamson explicitly claim that knowing how to F is not identical to the ability to F (at least in many cases), however, they still acknowledge intimate connections between knowing how to F and having an ability to F. Stanley and Williamson think that if someone has an ability to F, and F is an intentional action, then this entails that they know how to F (see 2001: 415–16, 442–3).<sup>19</sup> This is because for Stanley and Williamson intentional actions

---

<sup>19</sup> Moreover, recall that Stanley and Williamson (2001: 429) do acknowledge that there are some entailments in the other direction, i.e. from knowing how to F to possessing dispositional capacities or abilities: "[t]hinking of a way under a practical mode of presentation undoubtedly entails the possession of certain complex dispositions. It is for this reason that there are intricate connections between knowing-how and dispositional states". Stanley and Williamson must deny that these dispositions that are entailed by S's knowing how to F are identical to S's ability to F, given that they deny that knowing how to F entails having the ability to F (*ibid.* 416). Still, as discussed in §1.1, such dispositional capacities are presumably closely connected to the ability to F.

are “employments of knowledge-how” (*ibid.* 442–3), that is, whenever one successfully F’s one employs one’s knowledge how to F. AH\* could appeal to this same entailment to explain how Mary’s coming to know how to imagine an experience of red is intimately connected to her also coming to possess an ability to imagine an experience of red. The knowledge-how and the ability may not be identical, but the former is a precondition of the latter. Intuitively, knowing how to F and the ability to F (at least intentionally) are intimately related, and any viable account of knowledge-how will accommodate this intuition. One way to accommodate this intuition is to identify knowing how to F with the ability to F, but it is not the only way.

We have just seen that Stanley and Williamson make two important claims about the relationship between knowledge-how and ability that are not entailed by their core account of knowledge-how (as stated in §2.1). These claims are: that knowing how to F does not entail having the ability to F; and that having the ability to F (where F is an intentional action) entails knowing how to F. Consider now the claim that Mary gains new abilities:

(POS') Upon release Mary gains new abilities.

If we reject the distinctness claim we can understand POS' so that it is equivalent to POS, by identifying Mary’s new ability to imagine experiences of red with her new knowledge how to imagine experiences of red. While this suggestion is compatible with Stanley and Williamson’s main claim, it does not fit well with the motivations behind their further claim: that knowing how to F does not entail that one has the ability to F.

On the other hand, if we reject the identity claim then POS and POS' are distinct claims. Nevertheless, we can say that Mary’s new knowledge-how is a precondition of her new abilities—appealing to Stanley and Williamson’s idea that ability entails knowledge-how. Alternatively, we could just assert POS' and deny POS. This is the “fallback position” briefly discussed at the end of §2.2. We are in a good position now—having noted Stanley and Williamson’s claims about

entailment relations between knowledge-how ability—to review their arguments against this response to their account of knowledge-how.

As Stanley and Williamson note, the fallback position could be developed in two ways. The first development says that Mary only gains new abilities because “there is no knowing how to imagine an experience of red. There is just being able to imagine an experience of red” (Stanley and Williamson 2001: 442). Stanley and Williamson reject this development as they think that the ability to imagine an experience of red is clearly an ability to perform an intentional action. But then Mary must know how to imagine an experience of red if she gains the ability to imagine an experience of red—given Stanley and Williamson’s claim that having the ability to F entails knowing how to F.

The second way to develop the fallback position is to say that Mary only gains new abilities because she already had “in her black and white room, knowledge how to imagine an experience of red” (*ibid.* 443). Stanley and Williamson reject this development by appealing to our intuition that Mary does not know how to imagine an experience of red before she leaves her room. Stanley and Williamson ask, “If she knows how to imagine an experience of red, why is she unable to imagine such an experience?” (*ibid.* 443). This is perhaps a somewhat weak objection given that they themselves argue that it is possible that one know how to F without having the ability to F. However, as we saw in §2.2, there is a more fundamental reason to reject either development of the fallback position. If we are to preserve the intuition that Mary learns something, we do need to say that Mary gains new knowledge-how, and so, by Stanley and Williamson’s main claim, that she gains new knowledge-that.

We have seen that the sense in which Mary gains new knowledge-that conflicts with NEG<sub>2</sub> but need not conflict with NEG<sub>1</sub>. Classic statements of the ability hypothesis rely on the conjunction of the identity and the distinctness claims to show that NEG is compatible with POS. In the current context, we must either deny the identity claim, or the distinctness claim, or both. AH\* is compatible with any of these responses to the inconsistent triad. For AH\* does not rely on either

the identity or the distinctness claim to show that  $NEG_1$  is compatible with POS. Rather  $AH^*$  relies on the distinction between knowing a new proposition versus merely coming to be in a new state of propositional knowledge. Let us now examine this idea.

## 2.5 Old Facts New Modes

The ‘old-fact/new-mode’ reply to KA<sup>20</sup> relies on the idea that we can individuate states of propositional knowledge not only by what their propositional objects are, but also by what mode of presentation a proposition is known under.<sup>21</sup> Given that idea, the reply to KA is that Mary comes to be in a new state of propositional knowledge but does not gain knowledge of any new proposition. Mary only gains knowledge of an old proposition—i.e. a proposition she already knew under some mode of presentation—under a new mode of presentation. Proponents of the old-fact/new-mode reply often appeal to examples involving co-referring names to support this reply to KA. For example, (11) can be false even though (12) is true:

(11) Lois knows that Clark Kent is afraid of Kryptonite.

(12) Lois knows that Superman is afraid of Kryptonite.

Old-fact/new-mode theorists claim that if Lois comes to know that Clark Kent is afraid of Kryptonite she will come to be in a new state of propositional

---

<sup>20</sup> Also known as the ‘the old-fact/new-guise reply’ and the ‘two-ways reply’. Proponents include: Horgan (1984), Loar (1990), Lycan (1996), Perry (2001), and Tye (2000).

<sup>21</sup> For reasons of continuity I will state the old-fact/new-mode reply in terms of ‘old-propositions’ rather than ‘old-facts’, but nothing hinges on this choice. The old-fact/new-mode reply can be stated either way; I merely retain the name ‘old-*fact*/new-mode reply’ out of deference to convention. Also, for my purposes, questions about the difference between facts and propositions can be set to one side.

knowledge, but she will not come to know a new proposition. Like Mary, Lois will only come to know an old proposition under a new mode of presentation.<sup>22</sup>

Note that we could also have supported the old-fact/new-mode reply by appealing to an example we saw in Chapter 1: that (13) can be false even though (14) is true.

(13) John knows that he himself has burning pants.

(14) John knows that that man has burning pants.

Suppose that John comes to know that he himself has burning pants. John comes to be in a new state of propositional knowledge but he does not come to know a new proposition. He already knew the relevant proposition under a demonstrative mode of presentation. Like Mary, John only comes to know an old proposition under a new mode of presentation.

Clearly, AH\* is a version of this old-fact/new-mode form of reply to KA. Is this a problem for AH\* conceived as a version of the ability hypothesis reply to KA? One erroneous but revealing objection that could be made to AH\* is the following:

#### Objection A

---

<sup>22</sup> At the very least Lois will not come to know a new ‘coarse-grained’ proposition, like the Russellian notion of a proposition used by Stanley and Williamson. If we were to individuate propositions in a Fregean or ‘fine-grained’ manner (5) and (6) would express different propositions. But it is clearly a coarse-grained notion of a proposition that proponents of the old-fact/new-mode reply have in mind when they say that Mary does not gain knowledge of any new proposition. The old-fact/new-mode reply does not stand or fall, however, on issues about how to individuate propositions. Old-fact/new-mode theorists can grant that Mary will learn new fine-grained propositions or facts (see Lycan 1996: 61), they just deny that Mary would thereby gain new information of a kind that would trouble physicalism.

(A1) AH\* is a version of old-fact/new-mode reply to KA.

(A2) If a reply to KA is a version of the old-fact/new-mode reply to KA then it is not a version of the ability hypothesis reply to KA.

(A3) Therefore, AH\* is not a version of the ability hypothesis reply to KA.

The problem with this objection is that A2 is false. It is true that the old-fact/new-mode reply and the ability hypothesis are typically portrayed as rival responses to KA, for example by Jackson (2005: 318–20). Proponents of one of these forms of reply to KA are often critics of the other.<sup>23</sup> Consequently, one might think that there is a formal tension involved in claiming that AH\* is a version of both of these general forms of reply to KA. But in fact, there is no such tension. AH\* just demonstrates that the line between these two forms of reply to KA can be blurred depending on the details of specific versions of either reply. However, with few exceptions, this point is seldom acknowledged in the literature on KA (Van Gulick 2004: 391–2 is one notable exception; and Pettit 2004: 106–7 also seems sympathetic to this point). Objection A reveals that the gap between the old-fact/new-mode reply and the ability hypothesis reply is not as large as it is typically portrayed as being.

A more substantial objection could be made to AH\* based on the fact that it is a version of the old-fact/new-mode reply. Critics of the old-fact/new-mode reply often claim that it fails for a simple and principled reason. This perception leads David Chalmers (1996: 142) to say of the old-fact/new-mode reply that despite “the fact that it is easily the most popular response to the knowledge argument, it is also easily the weakest of the major replies. It simply does not hold up to scrutiny”. Typically, such critics appeal to what we can call, following Chalmers,

---

<sup>23</sup> For example, Lewis (1998: 454–5) argues against a rather narrowly conceived version of the old-fact/new-mode form of reply. On the other side, many proponents of the old-fact/new-mode form of reply are critics of the ability hypothesis, including Loar (1990) and Lycan (1996).

‘the new fact thesis’ (Chalmers 2004: 289–90).<sup>24</sup> Stated in its strongest form the new fact thesis says that:

(The New Fact Thesis) Necessarily, whenever someone gains new knowledge of an old proposition they must also gain knowledge of some new proposition.<sup>25</sup>

If true, the new fact thesis does appear to cause trouble for the old-fact/new-mode reply. For it is a consequence of the new fact thesis that Mary must come to know some new proposition if she comes to know an old proposition under a new mode of presentation—contra the old-fact/new-mode reply. In effect, the new fact thesis tells us that if  $NEG_2$  is false then  $NEG_1$  must be false as well—contra  $AH^*$ . This suggests the following objection to  $AH^*$ :

#### Objection B

(B1)  $AH^*$  is a version of the old-fact/new-mode reply to KA.

(B2) If the new fact thesis is true then all versions of the old-fact/new-mode reply to KA are unsuccessful.

(B3) The new fact thesis is true.

(B4) Therefore,  $AH^*$  is unsuccessful.

We cannot settle here the highly contested set of issues that would have to be

---

<sup>24</sup> Critics who appeal (more or less explicitly) to this thesis include: Chalmers (1996: 141–2, 2004: 289), Lockwood (1989: 136–7), Thau (2002: 127), and proponents of the ability hypothesis like Jackson (2005: 318) and Lewis (1998: 454–5).

<sup>25</sup> Again, for the sake of continuity I am stating the new *fact* thesis in terms of propositions rather than facts.

addressed to establish the truth or falsity of B3. Notably, it would help to have the correct account of informative identity statements. That the new fact thesis is controversial should not be surprising as it is really just a denial of the idea behind the old-fact/new-mode reply: that one can come to be in a new state of knowledge without coming to know any new proposition. But then why should a proponent of the old-fact/new-mode reply be moved to abandon their position by a thesis that is simply the denial of that position?

Critics of the old-fact/new-mode reply do argue that the new fact thesis should not be seen as controversial. Chalmers (2004: 289) says that the new fact thesis “can be endorsed even by those with very different views about reference and mental content”. Chalmers’ thought is that we need not appeal to any controversial views to establish the new fact thesis. Rather, we need only reflect on the stock examples appealed to by proponents of the old-fact/new-mode reply: that (5) can be false while (6) is true, even though both sentences attribute knowledge of the same proposition to Lois; or that one can know that there is water in one’s glass without knowing that there is H<sub>2</sub>O in it. In all variations of such examples (so the thought goes), one would necessarily learn new and substantive truths about the world if one came to know that Clark Kent is afraid of kryptonite or that there is H<sub>2</sub>O in my glass.<sup>26</sup>

Even if we grant that the new fact thesis is obviously true in all variations of these

---

<sup>26</sup> To illustrate this point consider Bill. Bill knows a lot of everyday propositions about water, but does not know that water is H<sub>2</sub>O. One day, Bill’s chemistry teacher informs him that water is H<sub>2</sub>O, and thereby Bill comes to know an old proposition under a new mode of presentation. But now Bill should also come to know entirely new and substantive truths about his world: that one substance is both water and H<sub>2</sub>O; that the liquid in one’s glass has a particular molecular structure; and so on. Advocates of the new fact thesis think that such examples provide good evidence for this thesis. For my part, I am not convinced that we can reasonably infer that anyone who gains new knowledge of an old proposition must also gain knowledge of some new proposition in *all* such cases—including very atypical cases like Mary’s—from the obvious fact that normally when we come to know a new proposition we also thereby come to know a lot of other propositions.

stock examples it is not obvious that it is generally true. Think of David Lewis' (1979) example of the two gods who occupy the same possible world. Both gods are omniscient with respect to every proposition true at their world, but they still lack knowledge as:

[n]either one knows which of the two he is. They are not exactly alike. One lives atop the tallest mountain and throws down manna; the other lives atop the coldest mountain and throws down thunderbolts. Neither knows whether he lives on the coldest mountain or on the tallest; nor whether he throws manna or thunderbolts (Lewis 1979: 520–1).

According to Lewis (1979: 521): “If the gods came to know which was which, they would know more than they do. But they wouldn't know more propositions”. Now imagine someone who (unlike Lewis) wished to explain the gods' new knowledge in terms of their coming to know old propositions under new modes of presentation. It would clearly be no objection to such an account to point out that it contradicts the new fact thesis.

Chalmers (2004: 289) himself acknowledges that the new fact thesis does not apply to what he calls ‘indexical cases’.<sup>27</sup> In doing so, he seems to acknowledge that John need not necessarily learn any new proposition when he learns that he is that man whose pants are on fire. But John's epistemic gain in such a situation is not trivial and, as we have seen, one explanation of his epistemic gain is that John comes to know an old proposition under a new first-personal mode of presentation. There is then at least one notion of coming to know an old proposition in a new way to which the new fact thesis does not apply.

This will not concern Chalmers, as he has independent reasons for thinking that appeals to indexical knowledge cannot support physicalist responses to KA. However, the inapplicability of the new fact thesis to indexical knowledge is significant as it shows that the new fact thesis cannot ground a principled

---

<sup>27</sup> See Chalmers' statement of the new fact thesis (Chalmers 2004: 289).

objection to all variations of the old-fact/new-mode reply. The new fact thesis may only show us that proponents of the old-fact/new-mode reply have often chosen poor examples to support their reply to KA.<sup>28</sup> This is an important concession but it falls well short of conceding that one cannot reply to KA by claiming that Mary only comes to know an old-fact in a new way.

Our reply to Objection B then is this: if the new fact thesis is taken to apply to all cases of coming to know an old proposition in a new way then it looks false; in which case, we can defend AH\* by denying B3. If, on the other hand, the new fact thesis is restricted so as not to apply to indexical cases, then we can defend AH\* by denying B2, for many versions of the old-fact/new-mode reply to KA do rely on appeals to first-personal or indexical modes of presentation. Now, it could still be the case that the new fact thesis applies to cases of coming to know an old proposition under a new practical mode of presentation. But if there is one exception to the new fact thesis couldn't there be more? At the very least, we can say that if cases of coming to know an old proposition under a new first-personal mode of presentation are exceptions to the new fact thesis then (for all we know) the same could be true for practical modes of presentation.

Furthermore, Stanley and Williamson's practical modes of presentation are modeled on—and are meant to share—important features with first-personal modes of presentation<sup>29</sup> (and also demonstrative modes of presentation). Given the close relationship between first-personal modes of presentation and practical modes of presentation, it seems reasonable to argue that if cases of coming to

---

<sup>28</sup> To be fair, many proponents of the old-fact/new-mode reply, such as Lycan (1996) and Perry (2001), do rely heavily on analogies between Mary's new knowledge and the gaining of new indexical knowledge.

<sup>29</sup> For example, Stanley and Williamson (2001: 433–4) concede that it might be impossible to give any reductive description of knowledge-how in non-indexical involving terms. They argue that this is not a problem for their account on the grounds that it is likely that the same is true for first-personal knowledge, yet first-personal knowledge is still a sub-class of knowledge-that.

know an old proposition under a new first-personal mode of presentation are exceptions to the new fact thesis, the same is likely to be true of practical modes of presentation.<sup>30</sup>

Undeniably, AH\* is vulnerable to objections that do not obviously arise for classic statements of the ability hypothesis. But AH\* is also better off in regards to a host of well-known criticisms of the ability hypothesis. For example, William G. Lycan (1996: 92–5) presents ten prominent arguments against the ability hypothesis. All of these arguments are meant to show “that what Mary gains is indeed propositional knowledge” (*ibid.* 92). That is, they are all arguments for the conclusion that Mary’s new knowledge must be *some* form of propositional knowledge. Without considering the details of Lycan’s ten arguments we can see that none of them affects AH\*, for AH\* endorses their conclusion. AH\* says that Mary’s new knowledge is knowledge-how, which in turn is a form of propositional knowledge. Thus, AH\* shows us that the ability hypothesis cannot be undermined by the claim that Mary’s new knowledge must be some form of propositional knowledge.

## 2.6 Conclusions

Stanley and Williamson (2001: 441) regard the ability hypothesis as an example of a philosophical thesis whose “reliance on the alleged distinction between knowledge-how and knowledge-that is fatal to the thesis advanced”. They reach this conclusion because they assume that their account of knowledge-how is inconsistent with the ability hypothesis. This assumption is not quite right. The heart of the ability hypothesis response to KA is the idea that NEG is compatible with POS. Classic statements of the ability hypothesis maintain that NEG is compatible with POS by defending the identity and distinctness claims—that Mary’s new knowledge-how is ability, and that these abilities are utterly distinct

---

<sup>30</sup> Intuitively, coming to know an old proposition under a *practical* mode of presentation does not sound like a process that necessarily involves coming to know some new proposition.

from propositional knowledge. The conjunction of these two claims is inconsistent with Stanley and Williamson's main claim: that knowledge-how is a species of knowledge-that. An alternative way to maintain the compatibility of NEG with POS is to discern an ambiguity in NEG. This allows one to recognize a version of the ability hypothesis, AH\*, that is consistent with both Stanley and Williamson's main claim as well as their further claims, that ability entails knowledge-how, and that knowledge-how does not entail ability. This is why the obvious answer that we began with is false, for we cannot assume that if knowledge-how is some form of propositional knowledge then it follows that no version of the ability hypothesis is true. AH\* shows us that we should not confuse the ability hypothesis reply to KA with any specific thesis about the nature of knowledge-how.<sup>31</sup>

In closing, I want to highlight two further implications of AH\*. The first of these is an implication for certain proponents of the ability hypothesis. Consider Frank Jackson: the inventor of the knowledge argument, but now a latter-day convert to the ability hypothesis. As Jackson (2004: 439) says, he now agrees "with Laurence Nemirow and David Lewis on what happens to Mary on her release." Jackson also dismisses the idea that the physicalist can respond to KA by saying that Mary only comes to know an old proposition under a new mode of presentation. His argument against the old-fact/new-mode reply rests on a variation of the new fact thesis (see Jackson 2005: 318).<sup>32</sup> But if Stanley and Williamson's account of

---

<sup>31</sup> Of course, there are possible accounts of knowledge-how that are inconsistent with the ability hypothesis. The point is that broadly Rylean accounts of knowledge-how (of the kind that Lewis and Nemirow appeal to) are not the only accounts of knowledge-how that can be used to state the ability hypothesis.

<sup>32</sup> Jackson thinks that every mode of presentation of some  $x$ —where he lets  $x$  stand for any unique thing, event or fact—must be associated with a unique feature or property of  $x$ . As he states: "Although there is one Robinson, for each guise that applies to him, there is a separate feature: it is a feature of Robinson that he is a suspect, that he is a tennis partner, and that he is the buyer of a house, and each of these features are distinct" (Jackson 2005: 318). The upshot of this view is that corresponding to every individual mode of presentation of some  $x$  there will be

knowledge-how is correct then Jackson can no longer affirm both NEG and POS by saying that Mary only gains new “abilities or know-how (as opposed to knowledge-that)” (*ibid.* 320). Stanley and Williamson’s main claim is inconsistent with this conjunction of the identity and distinctness claims. In this context, the only way to maintain that NEG and POS are compatible is to distinguish NEG<sub>1</sub> from NEG<sub>2</sub>, by appealing to something like the distinction between coming to know a new proposition versus merely coming to know an old proposition under a new mode of presentation. Jackson could choose to deny POS and only affirm POS'—the claim that Mary gains new abilities. But as we have seen, this response to Stanley and Williamson comes at a very high price, namely, that one can no longer say that Mary *learns* how to do something.

So, unless he can supply an argument against Stanley and Williamson’s account of knowledge-how, Jackson must choose between his commitment to the new fact thesis and his commitment to the ability hypothesis. The right choice for Jackson, *qua* proponent of the ability hypothesis, is clear. If we should not confuse the ability hypothesis with particular accounts of knowledge-how, then *a fortiori* we should not confuse the ability hypothesis with the views of some of its proponents about the new fact thesis, or any other issues not directly related to the core claims of the ability hypothesis, i.e. NEG and POS. Proponents of the ability hypothesis committed to the new fact thesis, like Jackson (2005) and Lewis (1998), must reconsider this commitment or provide an argument against Stanley and Williamson’s account of knowledge-how.

---

a unique fact about  $x$ . In the terminology we have been using we can say that every individual mode of presentation of  $x$  is paired with a unique proposition. So, Jackson thinks that if someone comes to know an old proposition under a new mode of presentation they must come to know some new proposition; that is, he endorses the new fact thesis.

Finally AH\* also has implications for our understanding of the distinction between knowledge-how and knowledge-that. In arguing that their account of knowledge-how is inconsistent with the ability hypothesis, Stanley and Williamson intend to underscore “the dangers of invoking Ryle’s distinction” (2001: 441) when doing philosophy. Of course, Stanley and Williamson themselves give an account of the distinction between knowledge-how and other forms of knowledge-that, but they regard this distinction as being philosophically insignificant. This claim also supports their denial that they have “just recreated the traditional distinction between knowing-how and knowing-that, but in other terms” (*ibid.* 434).

But we have seen in this chapter that we can deploy Stanley and Williamson’s account of knowledge-how to motivate the ability hypothesis reply to KA—one of the most important applications of the knowledge-how/knowledge-that distinction in philosophy. The moral is that even if some form of intellectualism is true there may still be a philosophically significant distinction between knowledge-how and other kinds of knowledge-that. Of course, there is still the further issue of whether or not intellectualism is in fact true. In Chapter 1 I argued that the insufficiency objection only establishes that simple, and not sophisticated, intellectualism is false. In the next chapter, however, I shall provide arguments that—if sound—demonstrate that no form of intellectualism is true.

## Chapter 3 Knowing How Without Knowing That

---

In this chapter I present three new arguments against intellectualism. In particular, I shall argue that, given certain very standard assumptions about the nature of knowledge-that, intellectualism is subject to three different kinds of counterexample. Each counterexample is a scenario where someone knows how to F but they fail to stand in the knowledge-that relation to any proposition p, such that their knowing how to F might plausibly be a matter (either partly or wholly) of their knowing that p. In each counterexample the subject fails to stand in the knowledge-that relation to some such proposition p because they fail to satisfy some familiar, necessary condition for knowing that p. The counterexamples differ with respect to what this necessary condition is.

In denying that knowledge-how is a matter—either partly or wholly—of standing in the knowledge-that relation to a proposition, I thereby endorse Gilbert Ryle’s conclusion that intellectualism is false, albeit on the basis of different arguments. However, I do not thereby endorse some form of Ryleanism. After presenting (§3.1) and defending my case against intellectualism (§§3.2–3.3), I suggest that the previous discussion points to a new view of knowledge-how that is distinct from both intellectualism and Ryleanism (§3.4), although it shares important features with each of these views.

### 3.1 Three Counterexamples

To see why the examples I will present are putative counterexamples to intellectualism it will help to have some actual intellectualist account of knowledge-how in mind. In what follows I will focus again on Stanley and Williamson’s (2001) intellectualist analysis of knowledge-how:

‘S knows how to F’ is true if, and only if, there is some way *w* for S to F such that:

- (a) S stands in the knowledge-that relation to the proposition that *w* is a way for S to F, and
- (b) In standing in this relation S entertains the proposition that *w* is a way for S to F under a practical mode of presentation.

According to this view of knowledge-how then, Shane knows how to juggle just in case he stands in the knowledge-that relation to some proposition of the form ‘*w* is a way for Shane to juggle’, and in standing in this relation he entertains this proposition under a practical mode of presentation.

With Stanley and Williamson’s account in mind I can introduce our three putative counterexamples. I take each example to be a case where knowledge-how *comes apart* from knowledge-that; that is, a case where someone knows how to F but there is no proposition *p* such that their knowing how to F might be plausibly a matter (either partly or wholly) of their knowing that *p*. The first example is a case where intuitively someone knows how to F, but they do not possess the kind of knowledge-that that such knowledge-how might be plausibly equated with, because their relevant beliefs are only *accidentally true*. Similar cases have been discussed in the literature, by Stanley and Williamson (2001: 435) and Ted Poston (forthcoming). But the possibility that such cases constitute counterexamples to intellectualism has normally been overlooked<sup>33</sup> (a point I will return to at the end of this section).

The second and third cases are each of a kind that has not been discussed before. The second case is a scenario where intuitively someone knows how to F but they do not possess the kind of knowledge-that that this knowledge-how might be plausibly equated with, because their relevant beliefs are *defeated*. The third case is a scenario where intuitively someone knows how to F but they do not possess the

---

<sup>33</sup> Aidan McGlynn, on his blog ‘The Boundaries of Language’, has independently noted the possibility of construing the kind of Gettier style cases that are already found in the literature as counterexamples to intellectualism (see <<http://aidanmcglynn.blogspot.com/2007/08/is-knowledge-how-gettier-susceptible.html>>).

kind of knowledge-that that this knowledge-how might be plausibly equated with, because they *lack* the relevant beliefs.<sup>34</sup> Here then are our three putative counterexamples:

*The Lucky Light Bulb*

Charlie wants to learn how to change a light bulb, but he knows almost nothing about light fixtures or bulbs. So, he consults *The Idiot's Guide to Everyday Jobs*. Inside, Charlie finds an accurate set of instructions describing a light fixture and bulb, and the way to change a bulb. Charlie grasps these instructions perfectly. And there is a way, call it ' $w_1$ ', such that Charlie now believes that  $w_1$  is a way for him to change a light bulb, namely, the way described in the book. However, unbeknownst to Charlie, he is extremely lucky to have read these instructions. For the disgruntled author of *The Idiot's Guide* filled her book with misleading instructions. Under every entry she intentionally misdescribed the objects involved in that job, and described a series of actions that would not constitute a way to do the job at all. However, at the printers, a computer error caused the text under the entry for 'Changing a Light Bulb', in just one copy of the book, to be randomly replaced by new text. By incredible coincidence, this new text provided the clear and accurate set of instructions that Charlie would later consult.

*The Dogmatic Hallucinator*

Lucy occasionally suffers from a peculiar kind of hallucination. On occasion, it seems to her that she remembers an event of learning how to F, when in fact no such event occurred. Furthermore, the way Lucy

---

<sup>34</sup> The three arguments by counterexample given here closely parallel, and are partly inspired by, arguments that Dean Pettit (2002: 519–50) has given for the conclusion that linguistic understanding is not a kind of knowledge-that. I will not examine Pettit's arguments here as to do so would take us into issues beyond the scope of this dissertation. Nevertheless, my intellectual debt to Pettit's arguments should be clear to those familiar with his excellent paper.

'remembers' as being the way to F, is not a way to F at all. On Saturday, a clown teaches Lucy how to juggle. Consequently, she knows how to juggle. And there is a way, call it ' $w_2$ ', such that Lucy now believes that  $w_2$  is a way for her to juggle, namely, the way the clown taught her to juggle. On Sunday, Lucy is about to tell a friend the good news that she knows how to juggle. However, as she begins, the alarm goes off on her false memory detector (or FMD), a remarkable device that is a super-reliable detector of her false memories. This indicates to Lucy that her apparent memory of learning how to juggle is a false memory that is misleading with respect to the way to juggle. Normally, Lucy would revise her beliefs accordingly, and this is what she believes she ought to do now. However, on this occasion she is unable to shake the beliefs she believes she ought to revise. For example, Lucy continues to believe that she knows how to juggle and that  $w_2$  is a way for her to juggle. Of course, Lucy did learn how to juggle yesterday, so her FMD has made an error, albeit one that was highly unlikely.

*The Non-Dogmatic Hallucinator*

Jodie occasionally suffers from a peculiar kind of hallucination. On occasion, it seems to her that she remembers an event of learning how to F, when in fact no such event occurred. Furthermore, the way Jodie 'remembers' as being the way to F, is not a way to F at all. On Saturday, a clown teaches Jodie how to juggle. Consequently, she knows how to juggle. And there is a way, call it ' $w_3$ ', such that Jodie now believes that  $w_3$  is a way for her to juggle, namely, the way the clown taught her to juggle. On Sunday, Jodie is about to tell a friend the good news that she knows how to juggle. However, as she begins, the alarm goes off on her false memory detector (or FMD), a remarkable device that is a super-reliable detector of her false memories. This indicates to Jodie that her apparent memory of learning how to juggle is a false memory that is misleading with respect to the way to juggle. Normally, Jodie would revise her beliefs accordingly, and this is exactly what Jodie does. She no longer believes that she knows how

to juggle or that  $w_3$  is a way for her to juggle. Of course, Jodie did learn how to juggle yesterday, so her FMD has made an error, albeit one that was highly unlikely.

The conclusion that these examples are all counterexamples to Stanley and Williamson's account of knowledge-how rests on two premises. The first premise is that the subjects in these cases each possess the relevant knowledge-how. More precisely, the premise states that the following claims are all correct, where ' $t_1$ ' refers to a moment just after Charlie has grasped the instructions in *The Idiots Guide*, ' $t_2$ ' refers to a moment just after Lucy has resisted revising her beliefs, and ' $t_3$ ' refers to a moment just after Jodie has revised her beliefs:

The Knowledge-how (KH) Claims

- (KH.1) At  $t_1$  Charlie knows how to change a light bulb.
- (KH.2) At  $t_2$  Lucy knows how to juggle.
- (KH.3) At  $t_3$  Jodie knows how to juggle.

The second premise is that the subjects do not possess the kind of knowledge-that which Stanley and Williamson would identify their knowledge-how with. More precisely, the premise states that the following claims are all correct:

The No Knowledge-that (NKT) Claims

- (NKT.1) At  $t_1$  Charlie does not know that  $w_1$  is a way for him to change a light bulb.
- (NKT.2) At  $t_2$  Lucy does not know that  $w_2$  is a way for her to juggle.
- (NKT.3) At  $t_3$  Jodie does not know that  $w_3$  is a way for her to juggle.

The KH claims, I submit, are all intuitively correct. The fact that Charlie is extremely lucky to have read accurate (as opposed to misleading) instructions just seems irrelevant to whether or not he comes to know how to change a light bulb on the basis of reading those instructions. The fact that a number of Lucy's beliefs about juggling are defeated does not seem to be a reason to think that she

has lost her knowledge how to juggle. Indeed, the intuitive thing to say with regard to Lucy's belief at  $t_2$  that she knows how to juggle, is that while this belief is unjustified, it is nonetheless true. Finally, the fact that at  $t_3$  Jodie no longer believes that she knows how to juggle, or that  $w_3$  is a way for her to juggle, does not seem to be a reason to conclude that Jodie has lost her knowledge how to juggle. Indeed, while Jodie's belief at  $t_3$  that she does not know how to juggle is justified, it is also intuitively false.

Moving to the NKT claims, recall that according to Stanley and Williamson if S knows how to F then there is some contextually relevant way  $w$  such that S knows that  $w$  is a way for S to F. But this putative necessary condition for knowing how to F fails to hold in any of our three scenarios. The contextually relevant ways in our three scenarios are clearly just  $w_1$ ,  $w_2$  and  $w_3$ . Now, at  $t_1$  Charlie does believe that  $w_1$  is a way for him to change a light bulb, and this belief is both true and justified. But this belief does not constitute knowledge, for it is only accidentally true, or true only as a matter of mere luck. And it is a familiar lesson from the Gettier literature, that knowledge-that is incompatible with the kind of epistemic luck present in this scenario.<sup>35</sup>

Similarly, at  $t_2$  Lucy does believe that  $w_2$  is a way for her to juggle. But again, this belief does not constitute knowledge, for Lucy knows that her FMD is a super-reliable detector of her false memories, and that these false memories are misleading with respect to the way to perform the relevant action. Lucy believes then that her belief that  $w_2$  is a way for her to juggle is not reliable or epistemically responsible. Furthermore, she is justified in this higher-order belief. In such a situation, Lucy's first-order belief that  $w_2$  is a way for her to juggle, while true,

---

<sup>35</sup> There is no need to offer an explicit analysis here of the kind of epistemic luck that knowledge-that excludes. Notoriously, no such analysis is widely accepted. But there is widespread agreement that knowledge-that excludes that kind of luck—whatever it is exactly—that is at work in Gettier cases. All we require then, is that the case I describe is of a kind with cases found in the Gettier literature; and I think that it clearly is. For an excellent discussion of the relationship between knowledge-that and epistemic luck see Duncan Pritchard (2005).

does not possess the justification or warrant necessary for knowledge.<sup>36</sup> Finally, at  $t_3$  Jodie clearly does not know that  $w_3$  is a way for her to juggle, for she does not even believe that  $w_3$  is a way for her to juggle.

I submit that the KH and NKT claims are all correct. It follows that each of our three examples is a counterexample to Stanley and Williamson's account of knowledge-how.

Furthermore, I submit that these examples will be counterexamples to any plausible account of knowledge-how, whereby knowing how to F is, at least partly, a matter of standing in the knowledge-that relation to some relevant proposition p. On any plausible version of such an account, this proposition p will concern something like a way, method or procedure for F-ing. If so, it would be an easy exercise to redescribe our three examples to emphasize the fact that Charlie's belief that p is only accidentally true, that Lucy's belief that p is defeated, and that Jodie does not believe that p. In other words, for any plausible version of either simple or sophisticated intellectualism—whereby one knows how to F only if one stands in the knowledge-that relation to some proposition p concerning a way to F—we will be able to provide parallel arguments for the corresponding NKT claims of the form: 'At  $t_n$  S does not know that p.' The arguments given here clearly apply then not only to Stanley and Williamson's account of knowledge-how, but also to the intellectualist accounts of knowledge-how offered by Bengson and Moffett (2007) and Brogaard (forthcoming a, forthcoming b) that I discussed in Chapter 1.

---

<sup>36</sup> I assume that the defeater for Lucy's belief that  $w_2$  is a way for her to juggle is her higher-order belief that her belief that  $w_2$  is a way for her to juggle is not reliable. But this assumption is not essential to my argument. It could be that the defeater is Lucy's experience of seeing the readout on her FMD or some relevant proposition. For our purposes, all that matters is that Lucy's belief that  $w_2$  is a way for her to juggle does not constitute knowledge-that in this scenario. Similarly, for ease of exposition, I assume that what gets defeated is Lucy's belief that  $w_2$  is a way for her to juggle. But my argument is perfectly consistent with views according to which it is Lucy's reasons for believing that  $w_2$  is a way for her to juggle that are defeated, rather than the belief itself.

To clarify these three arguments by counterexample, it may help to contrast one of them with a related but weaker form of argument against intellectualism examined by Stanley and Williamson (2001) and Poston (forthcoming), each of whom discusses cases like the lucky light bulb as a means of evaluating this objection. Stanley and Williamson imagine that someone might object to their account of knowledge-how by appealing to a supposed disanalogy between knowledge-how and knowledge-that:

On the analysis we presented in the last section, knowing-how is analyzed in terms of knowing-that. In particular, knowing how to F is a matter of knowing that  $p$ , for a certain proposition  $p$  (as well as entertaining it under the right mode of presentation). So, knowing-how is straightforwardly analysed in terms of knowing-that. But one might worry that significant disanalogies still remain between knowing-how and other kinds of knowing-that. One potential source of disanalogy involves Gettier cases. We can imagine cases of justified true belief that fail to be knowledge-that, because they fail to satisfy some extra condition. It may appear difficult to conceive of Gettier-cases for knowledge-how. But if knowledge-how is really a kind of knowledge-that, there should be such cases (Stanley and Williamson 2001: 435).

Stanley and Williamson dismiss this disanalogy objection by disputing the claim that there are no Gettier cases for knowledge-how.<sup>37</sup> In response, Poston defends the disanalogy objection by defending the claim that there are no Gettier cases for knowledge-how.

The disanalogy objection and my argument that appeals to the lucky light bulb case are importantly different. Suppose we could demonstrate that Poston is right, and there are no Gettier cases for knowledge-how. That is, no cases where one fails to know how to F for the same kind of reason one fails to know that  $p$  in a standard Gettier case. This alone would not establish that intellectualism is false. As it could be the case that knowledge-how is a kind of knowledge-that that is

---

<sup>37</sup> As we will see in §3.2, Stanley and Williamson also raise another reason for rejecting this disanalogy objection, namely, they reject the assumption that all kinds of knowledge-that are susceptible to Gettier cases.

merely disanalogous, in this respect, to other kinds of knowledge-that. That is, for all that we have shown, it could be the case that in any Gettier-like scenario where someone knows how to F, they will also possess the kind of knowledge-that that intellectualists would equate their knowledge-how with.

On the other hand, our argument claims that there is at least one Gettier scenario where someone knows how to F and they also fail to possess the kind of knowledge-that that this knowledge-how might be plausibly equated with. If this is correct, it does follow that knowledge-how is not a species of knowledge-that. Furthermore, the existence of such a Gettier scenario is consistent with the existence of other Gettier scenarios where knowledge-how and knowledge-that go together. This argument does not require then that knowledge-how is never susceptible to the kind of epistemic luck found in Gettier cases. Nor, for that matter, does it require that knowledge-that is always susceptible to such luck.

The crucial issue then, with respect to Gettier scenarios, is not whether or not there is some disanalogy between knowledge-how and knowledge-that with respect to such scenarios. Rather, the crucial issue is whether or not knowledge-how and knowledge-that *come apart* in any such scenarios.

The more general moral is that to respond to any of our putative counterexamples it will not suffice for the intellectualist to merely argue that there are other similar cases where knowledge-how and knowledge-that go together. Rather, the intellectualist must dispute the evaluation offered of these particular examples. There are obviously two ways they could do this. For each case the intellectualist could deny the relevant KH claim, or they could deny the relevant NKT claim. In §3.2 and §3.3 I will discuss both forms of response separately. In doing so, I hope to show that on close examination neither form of response is plausible.

### **3.2 The No Knowledge-that Claims**

The first form of response to our putative counterexamples that I will consider is one that disputes the relevant NKT claim. If we start with the lucky light bulb

case, the question is whether the intellectualist can reasonably deny NKT.1. Recall that the reason for thinking that Charlie's belief that  $w_1$  is a way for him to change a light bulb does not constitute knowledge—that is that this belief is only accidentally true.<sup>38</sup> If the intellectualist is to claim that at  $t_1$  Charlie does know that  $w_1$  is a way for him to change a light bulb, they will have to deny the standard view that knowledge—that is subject to an anti-luck condition. Namely, that if one knows that  $p$  then it is not a matter of mere luck or accident that one's belief that  $p$  is true. Denying NKT.1 appears to be an unattractive response to the lucky light bulb case because it commits the intellectualist to a major revision of our conception of knowledge—that.

The intellectualist might still respond that all that is needed is a 'localized' rejection of the idea that knowledge—that is subject to an anti-luck condition. Stanley and Williamson themselves could be interpreted as suggesting this kind of response in their discussion of the disanalogy objection:

---

<sup>38</sup> One might point out that at  $t_1$  Charlie is better positioned with respect to knowing that  $w_1$  is a way for him to change a light bulb than he was before  $t_1$ . For example, if he now attempts to change a light bulb he will come to know that  $w_1$  is a way for him to change a light bulb more easily than he would have if he did not already believe that this was the case. This is true but beside the point. For it does not alter the fact that at  $t_1$  Charlie does not know that  $w_1$  is a way for him to change a light bulb. Consider an analogy. At morning tea you ask Mary if she knows which bus goes to Kingston. She tells you that the 501 goes to Kingston and you believe her, despite the fact that you know Mary is a compulsive liar. And indeed, Mary did intend to give you false information but, by accident, she gave you correct information. Given your new true belief that the 501 goes to Kingston, you are now in a better position with respect to knowing that the 501 goes to Kingston. For now you are more likely to act in ways that will lead to you gaining further evidence in support of this belief. For example, if you want to get a bus to Kingston you will choose to catch the 501. But your new proximity to knowledge does not change the fact that at morning tea—when your only evidence is testimony from a source you know to be highly unreliable—your belief does not constitute knowledge.

We doubt that every kind of knowledge-that is susceptible to Gettier cases. So it would not worry us if it were not possible to come up with a Gettier case for knowledge-how (Stanley and Williamson 2001: 435).

On one interpretation of this passage, Stanley and Williamson are claiming that they would be unconcerned if they had to deny that knowledge-how is subject to an anti-luck condition, because they think that there are other kinds of knowledge-that which are also not subject to such a condition.<sup>39</sup> And the claim that knowledge-how is not subject to an anti-luck condition is consistent with the claim that other kinds of knowledge-that are subject to such a condition. Stanley and Williamson might then point out that in claiming that Charlie knows that  $w_1$  is a way for him to change a light bulb they need only commit themselves to the claim that one particular kind of knowledge-that is not subject to an anti-luck condition.

However, Stanley and Williamson cannot simply assert that knowledge-how is a distinctive kind of knowledge-that that is not susceptible to Gettier cases. Rather, what they would need to establish is that S's standing in the knowledge-that relation to a proposition of the form ' $w$  is a way for S to F', is a distinctive kind of knowledge-that that is not susceptible to Gettier cases. But why should we think

---

<sup>39</sup> The claim that not all kinds of knowledge-that are susceptible to Gettier cases is somewhat difficult to interpret, as there are at least two quite distinct ways it could turn out to be true. As interpreted above, the idea is that there is at least one kind of knowledge-that such that one can possess this kind of knowledge-that even when one's relevant justified true beliefs are only accidentally true. If this were the case, then this kind of knowledge-that would not be susceptible to Gettier cases because it is not subject to an anti-luck condition. But perhaps Stanley and Williamson's idea here is that there are some kinds of knowledge-that such that one simply cannot describe any scenario where one has the relevant justified true beliefs but they are only accidentally true. If this were the case then this kind of knowledge-that would not be susceptible to Gettier cases, but it would still be subject to an anti-luck condition, for it would trivially satisfy such a condition. I have focused on the former idea above for the simple reason that we obviously can describe scenarios where someone has a justified true belief of the form ' $w$  is a way to F' that is only accidentally true.

that this is the case? There is nothing obviously special about propositions concerning ways to perform actions such that S could know that p, even though S's belief that p is merely accidentally true, whenever p happens to be a proposition of the form '*w* is a way for S to F.'

Perhaps Stanley and Williamson might argue that the relevant kind of knowledge—that that is not susceptible to Gettier cases is the knowledge-that which S has when S stands in the knowledge-that relation to some proposition of the form '*w* is a way for S to F' and, in standing in this relation, S entertains that proposition under a practical mode of presentation. That is, Stanley and Williamson could claim that the fact that Charlie's belief that *w*<sub>1</sub> is a way for him to change a light bulb is accidentally true is irrelevant to whether or not he knows that *w*<sub>1</sub> is a way for him to change a light bulb *under a practical mode of presentation*.

But note how odd this suggestion would be. No one has ever tried to defend the tripartite analysis of knowledge by claiming that while the subjects in Gettier cases do not come to know that p under such-and-such mode of presentation, they do come to know that p under some other mode of presentation. And there is a good reason why not. For the fact that someone's belief that p is merely accidentally true is surely a reason to think that they do not know that p *simpliciter*, regardless of what mode of presentation they happen to entertain that proposition under.

At the very least, if Stanley and Williamson were to adopt this response they would owe us an explanation of why knowledge of propositions of the form '*w* is a way for one to F' is resistant to Gettier influences in the special case where one entertains that proposition under a practical mode of presentation. And this explanation cannot simply consist in the claim that knowing that *w* is a way for one to F under a practical mode of presentation is knowledge-how, and knowledge-how is resistant to Gettier influences.

The problem is that modes of presentation look like the wrong kind of thing on which to base such an explanation. Consider the sorts of reasons that are typically offered to explain why S fails to know that p in a given Gettier scenario: that the

truth of S's belief that p is not appropriately related to S's reasons for holding that belief, or that the source of S's belief that p is unreliable, and so on. Such reasons for thinking that S fails to know that p do not seem even to be addressed—let alone outweighed or undermined—by the extra information that S happens to entertain p under such-and-such a mode of presentation.

The general point is that it is difficult to see how the intellectualist could motivate the claim that in denying NKT.1 they need only endorse a localized, rather than wholesale, rejection of the idea that knowledge-that is subject to an anti-luck condition. This is because the kind of knowledge-that that intellectualists' equate knowledge-how with has no distinctive features that would support such a claim. Denying NKT.1 is still an unattractive response then to the lucky light bulb case, given that it commits the intellectualist to such a major revision of the standard conception of knowledge-that.

And, if anything, the situation with regard to NKT.2 and NKT.3 is worse. Recall the reasons given in §3.1 for accepting these two claims. NKT.2 was supported by the claim that at  $t_2$  Lucy's belief that  $w_2$  is a way for her to juggle is defeated, and hence does not possess the justification or warrant necessary for it to constitute knowledge. NKT.3 was supported by the claim that at  $t_3$  Jodie does not believe that  $w_3$  is a way for her to juggle. If we accept the defeat and no-belief claims, the consequences of denying NKT.2 and NKT.3 are severe. If the defeat claim is true, to deny NKT.2 is also to deny that having justification or warrant for one's belief that p is a necessary condition for knowing that p. And if the belief claim is true, to deny NKT.3 is also to deny that believing that p is a necessary condition for knowing that p.

Faced with a choice between maintaining that knowledge-how is a species of knowledge-that, and denying that knowledge-that is subject to a justified or warranted belief condition, I take it that the right choice is clear. We should reject the intellectualist thesis. Faced with the parallel choice with regard to the belief condition, the right choice is just as clear. Again, we should reject the intellectualist thesis. To choose otherwise in either case would be to radically

revise our conception of knowledge-that, just to maintain the thesis that knowledge-how is a kind of knowledge-that.

If the intellectualist is to deny NKT.2 and NKT.3 then, they must establish that the defeat and belief claims are false. But can one plausibly deny either of these claims? Perhaps, against the defeat claim, the intellectualist might argue that when one entertains a proposition  $p$  under a practical mode of presentation, then one's belief that  $p$  can be justified even when one has a justified belief that their belief that  $p$  is unreliable. But again, I think the intellectualist would be hard pressed to justify this 'localized' rejection of what clearly looks like a necessary condition for knowledge-that in general. Namely, that if one knows that  $p$  then one does not have a justified belief that one's belief that  $p$  is unreliable, or epistemically inappropriate.<sup>40</sup> The fact that Lucy has a justified belief that her belief that  $w_2$  is a way for her to juggle is unreliable, is surely a reason to conclude that she does not know that  $w_2$  is a way for her to juggle *simpliciter*. It is not merely a reason to conclude that Lucy does not know that  $w_2$  is a way for her to juggle, if she happens to entertain this proposition under a non-practical mode of presentation.

What of the no-belief claim? Could not one argue that at  $t_3$  Jodie still *implicitly* or *tacitly* believes that  $w_3$  is a way for her to juggle? And, if so, could not one argue that Jodie still implicitly or tacitly knows that that  $w_3$  is a way for her to juggle? Undoubtedly, there is a good sense in which at  $t_3$  it will still *seem* to Jodie that  $w_3$  is a way for her to juggle. For example, if Jodie imagines  $w_3$ , this way will still strike her as being a way to juggle. But we should not confuse mere seemings with beliefs. Even if one knows that the two lines in a Müller-Lyer figure are of the same length, it will still seem to one that they differ in length. And as George Bealer (1993) has pointed out (amongst others), the same point applies not only

---

<sup>40</sup> This kind of condition is widely accepted as a necessary condition for knowledge-that by both internalists and externalists, for discussion see Michael Bergman (1997: 399–417). There is a debate about whether one's second-order belief that one's belief that  $p$  is not reliable must itself be justified in order for it to defeat one's first-order belief that  $p$ . But this debate is not relevant here given that Lucy's higher-order belief is justified.

to perceptual seemings but also to intellectual seemings. To use one of Bealer's examples, it can still seem to one that the naïve axiom of set theory is true, even though one does not believe that it is true because one knows that it leads to a contradiction. Likewise, while it seems to Jodie that  $w_3$  is a way for her to juggle, I think it is clear that she fails to believe that  $w_3$  is a way for her to juggle.

Furthermore, Jodie has consciously reflected on the question of whether or not  $w_3$  is a way for her to juggle, and she has concluded on the basis of her relevant evidence that  $w_3$  is not a way for her to juggle. If someone has consciously reflected on the question of whether or not  $p$  and concluded on the basis of their relevant evidence that not- $p$ , this is normally a strong indicator that they do not believe that  $p$ . There are difficult cases (including ones involving delusional beliefs) where one might think that someone has both the belief that  $p$  and the belief that not- $p$  at the same time. But I see no reason to regard the non-dogmatic hallucinator as being such a case. More importantly, even if there can be cases where one still believes that  $p$  after coming to believe that not- $p$  on the basis of the kind of conscious reflection Jodie engages in, these would clearly be cases where one fails to know that  $p$ .

Denying the relevant NKT claim does not look to be a plausible response for the intellectualist to any of our putative counterexamples. In each case, denying the NKT claim forces the intellectualist to reject a plausible and widely accepted assumption about the nature of knowledge-that. However, there is still another form of response to these counterexamples that needs to be evaluated.

### 3.3 The Knowledge-how Claims

The second possible form of response to our putative counterexamples is to contest the relevant KH claim. There is reason to think that Stanley and Williamson would at least reject KH.1. Consider what Stanley and Williamson say about the following example they offer as proof that there can be Gettier cases for knowledge-how:

Bob wants to learn how to fly in a flight simulator. He is instructed by Henry. Unknown to Bob, Henry is a malicious imposter who has inserted a randomising device in the simulator's controls and intends to give all kinds of incorrect advice. Fortunately, by sheer chance the randomising device causes exactly the same results in the simulator as would have occurred without it, and by incompetence Henry gives exactly the same advice as a proper instructor would have done. Bob passes the course with flying colors. He has still not flown a real plane. Bob has a justified true belief about how to fly. But there is a good sense in which he does not *know* how to fly (Stanley and Williamson 2001: 435).

So, Stanley and Williamson think that this example—I will call it *the flight simulator case*—is a case where someone fails to know how to F for the same kind of reason one fails to know that p in a Gettier scenario. Now, for the reasons discussed at the end of §3.1, if Stanley and Williamson's evaluation of this case is correct it does not follow that KH.1 is false. Nevertheless, given the obvious similarities between the flight simulator and lucky light bulb cases one might reasonably expect that our verdicts about whether Bob knows how to fly and whether Charlie knows how to change a light bulb should be the same. If Stanley and Williamson are right in claiming that Bob does not know how to fly, this would at least give us some reason to reconsider KH.1.

But are Stanley and Williamson right? Is there a good sense in which Bob does not know how to fly? Clearly, Bob has justified and true beliefs about flying that do not constitute knowledge-that, because they are only accidentally true. However, I think Stanley and Williamson are simply wrong that the intuitive thing to say of this case is that Bob does not know how to fly. As Poston (forthcoming) says: "As far as intuition goes this does not seem correct. There is a good sense in which Bob *does* know how to fly."

To make the intuition vivid, compare Bob with his near perfect counterpart Joe. The only salient difference between Bob and Joe is that in Joe's world his simulator not only operates correctly but it has not been interfered with, and his instructor not only gives him the correct advice but he intended to do so. So, when Joe exits his simulator, we can safely assume that he knows how to fly. But

on what grounds then, could we deny that Bob knows how to fly? The fact that Bob, unlike Joe, is extremely lucky to receive the very same feedback from his simulator/instructor does not seem to be a reason to conclude that only Joe comes to know how to fly on the basis of receiving this feedback.<sup>41</sup>

Someone might try to argue for the claim that there is both a good sense in which Bob knows how to fly and a good sense in which he does not know how to fly. I doubt that this is the case, but two points are worth mentioning about this claim. First, it is clear that Stanley and Williamson themselves do not take knowledge-how ascriptions to be ambiguous in this way. Second, as Stanley and Williamson acknowledge, Bob's relevant belief of the form '*w* is a way for Bob to fly' does not constitute knowledge-that in this scenario. If so, then if there is a good sense in which Bob knows how to fly it follows that there is a good sense in which knowledge-how comes apart from knowledge-that in the flight simulator case. In other words, it would follow that there is a good sense in which knowledge-how is not a kind of knowledge-that.<sup>42</sup>

---

<sup>41</sup> Note that we could have used a similar comparison to support the intuition for KH.1. Compare Charlie with his near perfect counterpart Jack. Jack's world is just like Charlie's in all but one salient respect, namely, in Jack's world *The Idiots Guide\** was written by a non-malicious author who intended to fill her book with helpful descriptions of ways to perform everyday jobs (and there were no errors during printing etc.). The text in the light bulb section of Jack's copy of *The Idiot's Guide\** is the same as the text in Charlie's copy of *The Idiot's Guide*. So Jack reads the exact same description of how to change a light bulb that Charlie reads. And like Charlie, Jack comprehends these instructions perfectly. Obviously, it is safe to assume that Jack knows how to change a light bulb after reading these instructions. This is an ordinary way of gaining knowledge-how. But how could we deny that Charlie comes to know how to change a light bulb after reading the very same instructions? The fact that Charlie, unlike Jack, is extremely lucky to read these instructions does not seem to be a reason to conclude that only Jack comes to know how to change a light bulb.

<sup>42</sup> Stanley (2005: 133) explicitly denies the claim that knowledge-how ascriptions are ambiguous between a sense in which they ascribe knowledge-how and a sense in which they do not, as do Bengson and Moffett (2007: 38–40). As discussed in Chapter 1, Brogaard does hold that 'S

Stanley and Williamson’s interpretation of this case is also strange given that their own account of knowledge-how tells us that Bob knows how to fly. Let me explain. The core of Stanley and Williamson’s account of knowledge-how was stated earlier in §3.1. Recall, however, that we saw in Chapter 2 that Stanley and Williamson also make two further, and important, claims about the nature of knowledge-how. First, Stanley and Williamson (2001: 442–3, see also 2001: 415–16) hold that all intentional actions “are employments of knowledge-how”. That is, they accept the following claim:

(1) If S Fs intentionally, S knows how to F.

Second, Stanley and Williamson infer from (1) a further claim concerning abilities, as their discussion of the ability hypothesis reply to the knowledge argument reveals:

For the ability to imagine an experience of red is clearly an ability to perform an intentional action. And we do find it very plausible that intentional actions are employments of knowledge-how ... But if intentional actions are employments of knowledge-how then Mary’s acquisition of an ability to imagine an experience of red brings with it knowledge how to imagine red [.] (*ibid.* 442–3).

So, Stanley and Williamson hold that if one has the ability to perform an action intentionally then one knows how to perform that action. That is, they accept the following claim:

(2) If S has the ability to F intentionally, S knows how to F.

---

knows how to F’ ascriptions are ambiguous, as she claims that ‘John knows how to play the piano’ can be interpreted either as “saying that there is a *w* such that John knows that *w* is how John may play the piano or as saying that there is a *w* such that John knows that *w* is how one may play the piano” (forthcoming b: 47). But clearly, on either disambiguation, knowing how to play the piano is still a kind of knowledge-that.

But then it is a necessary consequence of their full account of knowledge-how—and a plausible assumption—that Bob does know how to fly. The assumption is that Bob has the ability to fly a plane intentionally. And this is very plausible. After all, Bob passes the course that imparts this ability with “flying colours”. To emphasize the point, note that Joe has the ability to fly a plane intentionally as he exits his simulator. But then we must conclude that Bob also has this ability, for Joe and Bob are clearly equivalent with respect to their abilities to fly a plane.

The issue here can be illustrated by noting that the following three claims form an inconsistent triad:

- (3) If S has the ability to F intentionally, S knows how to F.
- (4) Bob has the ability to fly intentionally.
- (5) Bob does not know how to fly.

Stanley and Williamson claim both that Bob does not know how to fly, and that having the ability to F intentionally entails knowing how to F; that is, they endorse both (2) and (4). However, (3) is true. It must be the case then that either (2) or (4) or both (2) and (4) are false. So, to maintain that Bob does not know how to fly Stanley and Williamson would have to deny (2), thereby denying a key commitment of their full account of knowledge-how.

Furthermore, if Stanley and Williamson are right that having the ability to F intentionally entails knowing how to F, this is highly important in this context given that the following ability ascriptions are very plausible:

- (6) At  $t_1$  Charlie has the ability to change a light bulb intentionally
- (7) At  $t_2$  Lucy has the ability to juggle intentionally
- (8) At  $t_3$  Jodie has the ability to juggle intentionally

For if Stanley and Williamson are right that (2) is true then (6), (7) and (8) each entail the corresponding knowledge-how ascription, that is, they entail KH.1, KH.2 and KH.3 respectively.

Could Stanley and Williamson reply that while Charlie, Lucy and Jodie possess the ability to perform these actions they do not possess the ability to perform them intentionally? Perhaps with regard to (8), one might argue that to have the ability to juggle intentionally, Jodie would have to believe that  $w_3$  is a way for her to juggle. However, as mentioned earlier, at  $t_3$  it would still seem to Jodie that  $w_3$  is a way for her to juggle. So, one could convince Jodie to try to juggle that way that merely seems to her to be a way to juggle. And if she did try, she would likely succeed. In which case, I think the natural thing to say would be that Jodie not only juggled but that she did so intentionally.

In any case, even if one could resist (8) on these grounds, (6) and (7) seem straightforwardly true. It may be a necessary condition of S's having the ability to F intentionally that there be some way  $w$  that is a way for S to F such that S believes that  $w$  is a way for S to F. But it is surely not a necessary condition of S's having the ability to F intentionally that such a belief must also be non-accidentally true and/or justified.<sup>43</sup>

If the intellectualist is to deny any of these KH claims (or at least KH.1 and KH.2) then they must deny that having the ability to F intentionally entails knowing how to F. Luckily, other intellectualists have identified cases that are fairly plausible counterexamples to this entailment, as a means of arguing against neo-Ryleanism; that is, the view that to know how to F is to simply possess the ability to F. Recall, as I mentioned in Chapter 1, Bengson and Moffett present the following scenario—I will refer to it from now on as *the salchow* case—as an example where intuitively someone has the ability to F but does not know how to F:

---

<sup>43</sup> Bengson et al. (forthcoming: fn. 23) make a similar point in regard to Stanley and Williamson's claim that having the ability to F intentionally entails knowing how to F.

Suppose that Irina is seriously mistaken about how to perform a salchow. She believes incorrectly that the way to perform a salchow is to take off from the front outside of her skate, jump in the air, spin, and land on the front inside edge of her skate. (The correct sequence is to take off from the *back inside edge* and land on the *back outside edge* of the opposite foot after one or more rotations in the air.) However, Irina has a severe neurological abnormality that makes her act in ways that differ dramatically from how she actually thinks she is acting. Whenever she actually attempts to do a salchow (in accordance with her misconceptions) this abnormality causes her to reliably perform the correct sequence of moves. So, although she is seriously mistaken about how to perform a salchow, whenever she actually attempts to do a salchow (in accordance with her misconceptions) the abnormality causes Irina to perform the correct sequence of moves, and so she ends up successfully performing a salchow. Despite the fact that what she is doing and what she thinks she is doing come apart, she fails to notice the mismatch. In this case, it is clear that Irina is (reliably) able to do a salchow. However, due to her mistaken belief about how to perform the move, she cannot be said to know how to do a salchow (Bengson and Moffett 2007: 46).

And Paul Snowdon (2004) presents the following scenario—I will refer to it from now on as *the man in a room* case—as a counterexample to the claim that if one has the ability to F then one knows how to F:

A man is in a room, which, because he has not explored it in the least, he does, as yet, not know how to get out of. In fact, there is an obvious exit which he can easily open. He is perfectly able to get out, he can get out, but does not know how to (as yet) (Snowdon 2004: 11).

Both of these examples appear to be plausible counterexamples not only to the claim that if one has the ability to F then one knows how to F, but also to the stronger claim that if one has the ability to F intentionally then one knows how to F. And if (2) is false then one can *consistently* deny KH.1, KH.2 and KH.3 whilst accepting (6), (7) and (8). However, this is not yet a reason to think that any of the KH claims are false.

It is clear that in practice, many subjects would have the intuition that the KH claims are correct. One line of response available to the intellectualist would be to claim that while the KH claims are intuitive they are nonetheless false. But if they

are to deny these intuitive claims, the intellectualist owes us some explanation of why our intuitions about these cases are so systematically misleading.

Probably the most obvious explanation would be to claim that we somehow confuse the fact that the subjects in our putative counterexamples possess the relevant ability with their possessing the corresponding knowledge-how. Appealing to the idea that ability ascriptions implicate—but do not entail—the corresponding knowledge-how ascription, would be one way to develop such an argument. The explanation then of our intuitions regarding KH.1, KH.2 and KH.3, would be that we confuse a conversational implicature with an entailment. For example, our intuition that Charlie knows how to change a light bulb is explained by the fact that we know that Charlie has the ability to change a light bulb, and we mistakenly think that ‘S has the ability to F’ entails ‘S knows how to F’.

This strategy for explaining away our intuitions regarding the KH claims may appear promising.<sup>44</sup> The *salchow* and *man in a room* cases do seem to show that there is no entailment from ‘S has the ability to F’ to ‘S knows how to F’, even when S’s ability is an ability to F intentionally. But presumably, in stereotypical or paradigmatic cases of someone’s having the ability to F, they will also know how to F, in which case it seems reasonable to suppose that there is a sense in which ‘S has the ability to F’ implicates ‘S knows how to F’.

However, note that there is an inherent tension in this kind of response to our putative counterexamples. To establish that having the ability to F does not entail knowing how to F, the intellectualist needs there to be clear cases where intuitively someone has the ability to F but does not know how to F. And there are such cases. But then why does our familiarity with the relevant implicature lead us to mistakenly have the intuition that KH.1, KH.2 and KH.3 are true,

---

<sup>44</sup> Bengson and Moffett think that there is a stereotypical implicature in the other direction, from knowing how to F to having the ability to F. For further discussion of the notion of a stereotypical implicature see Bengson and Moffett (2007: 35).

when it obviously does not lead us to make the parallel mistake with regard to *the salchow* and *man in a room* cases? In both sets of cases the relevant subject has the ability to F intentionally and, according to the intellectualist, does not know how to F. The intellectualist then would have to provide a plausible explanation of this asymmetry that is also consistent with their interpretation of these cases. Perhaps there is some such explanation, but I am not sure what it would be.

On the other hand, we can offer a natural explanation of this asymmetry in our intuitions, namely, that the subjects in the lucky light bulb, dogmatic hallucinator and non-dogmatic hallucinator cases know how to perform the relevant actions, whereas the subjects in *the salchow* and *man in a room* cases do not.

I also doubt that it is an essential feature of the counterexamples offered here that the subjects in these scenarios possess the ability to perform the relevant actions. As intellectualists often point out, one can know how to F without possessing the ability to F, as Stanley and Williamson's (2001: 416) case of the master pianist who loses her arms in a tragic car accident illustrates. Intuitively, after such an accident the master pianist would still know how to play the piano even though she has lost her ability to do so. Again, such examples are cited by intellectualists as evidence against neo-Ryleanism, for they suggest that having the ability to F is not a necessary condition for knowing how to F.

Bearing this point in mind, let us add an unfortunate twist to the lucky light bulb case. Namely, just after Charlie grasps the instructions in *The Idiot's Guide* at  $t_1$  his arms are removed (I will leave the details of how to your imagination). Otherwise, the case remains exactly the same. Does Charlie still know how to change a light bulb? As with the pianist case, I take it that the intuitive answer is yes. In which case, we still have a scenario where intuitively Charlie knows how to F, and the same reasons are still present for thinking that Charlie does not possess the kind of knowledge-that that such knowledge-how might be plausibly equated with. However, in this modified scenario, Charlie also lacks the ability to change a light bulb. So, the intellectualist cannot dismiss the knowledge-how intuition here by claiming that we are merely confusing the fact that Charlie has the ability to

change a light bulb with his knowing how to change a light bulb. And one could modify the dogmatic hallucinator and non-dogmatic hallucinator cases to achieve the same kind of result.

However, I think the more important point is simply that there are good reasons to be suspicious of such attempts to dismiss our intuitions regarding the KH claims. Consider the very examples intellectualists appeal to when arguing against neo-Ryleanism—the *salchow* and *man in a room* cases or Stanley and Williamson’s pianist case. As counterexamples to neo-Ryleanism, these cases are compelling. But the intuitive force of such examples suggests that we are quite capable of discerning the difference between knowing how to F and possessing the ability to F. It seems implausible then to suppose that our intuitions about the KH claims are merely the result of our confusing the fact that a subject has the ability to F with their knowing how to F.

There is no simple way to dismiss our intuitions about the KH claims. But intellectualism requires that we deny the KH claims, for we saw in §3.2 that denying the NKT claims is not a plausible response to our putative counterexamples. In the absence of some good way to dismiss our intuitions regarding the KH claims, I submit that we should reject intellectualism.

### 3.4 Toward a New Theory

I have argued in this chapter that intellectualism is false. But what is knowledge-how if it is not a kind of knowledge-that? The most prominent alternative to intellectualism is neo-Ryleanism. However, neo-Ryleanism does not appear to be a viable alternative to intellectualism. The *salchow* and *man in a room* cases suggest that having the ability to F does not suffice for knowing how to F, and examples like the pianist case suggest that having the ability to F is not necessary for knowing how to F.

However, I think that by reflecting on our three counterexamples we can find a promising alternative to both intellectualism and neo-Ryleanism. According to this view, knowing how to F is a matter of standing in a relation to a proposition

other than the knowledge-that relation. The relevant relation is the one that S stands in to a proposition p, when it *seems* to S that p is the case.

Importantly, this relation is not the belief relation. As mentioned earlier, it can seem to one that p even when one fails to believe that p. That is, believing that p is not a necessary condition for it seeming to one that p. In which case, seemings cannot be understood as merely some species of belief.<sup>45</sup>

Also, the kind of seemings that are relevant to this new account of knowledge-how are not perceptual (or sensory) seemings. Paradigmatic cases of perceptual seemings include visual seemings, as when one looks at the Müller-Lyer figure and one's visual experience presents the two horizontal lines as being unequal in length. Paradigmatic examples of non-perceptual seemings include intellectual seemings, as when one entertains the proposition that if p then not not p, and then one "sees" that it is true.

Intellectual seemings are often contrasted not only with perceptual seemings, but also with seemings that are grounded in introspection, imagination, or memory.<sup>46</sup> In which case, the class of non-perceptual seemings includes, but is not exhausted by, the class of intellectual seemings. In claiming that the kind of seemings relevant to our new account of knowledge-how are non-perceptual, I intend to remain neutral on the question of whether the seemings in question should be characterised as purely intellectual seemings, or as non-perceptual seemings

---

<sup>45</sup> Of course, this is not to say that there aren't important connections between seemings and belief. As is often noted, seemings typically incline one to believe their propositional content. Although, as William Tolhurst (1998: 297) points out, seemings do more than just incline one to believe, for when it seems to one that p "one experiences believing [that p] to be demanded or required." Furthermore, "seemings have the feel of truth, the feel of a state whose content reveals how things really are" (*ibid.* 298–9).

<sup>46</sup> As Bealer (1993: 102) says, intellectual seemings are seemings that one can have "in the absence of any particular sensory (imaginative) or introspective experiences." For a similar construal of intellectual seemings see also Michael Huemer (2005) and Joel Pust (2000).

(because they might be seen as essentially relying on introspection, imagination or memory).

Bearing these points in mind—that seemings are not beliefs and that the kind of seemings at issue here are non-perceptual seemings—we can introduce a new analysis of knowledge-how:

#### The Seeming Analysis

S knows how to F if, and only if, there is some way  $w$  that is a way to F such that:

- (c) S stands in a (non-perceptual) seeming relation to the proposition that  $w$  is a way to F, and
- (d) In standing in this relation S entertains  $w$  under a practical mode of presentation.

The motivation for the seeming analysis is that in all of our counterexamples to intellectualism there is intuitively still some way  $w$  for the subject to perform the relevant action F such that it *seems* to the subject that  $w$  is a way to F. It seems to Charlie that  $w_1$  is a way to change a light bulb even though his belief that  $w_1$  is a way for him to change a light bulb is only accidentally true. It seems to Lucy that  $w_2$  is a way to juggle even though her belief that  $w_2$  is a way for her to juggle is defeated. And, as noted earlier, it still seems to Jodie that  $w_3$  is a way to juggle even though she does not believe that  $w_3$  is a way for her to juggle.

Of course, the point here is not that if Charlie (for example) were to simply observe someone else changing a light bulb at  $t_1$  then it would still seem to him that that way is a way to change a light bulb. This is true, but this would only be a perceptual seeming. Rather, the point is that if at  $t_1$  Charlie were to entertain the proposition that  $w_1$  is a way to F, or if he were to simply think of  $w_1$ , it would still seem to him that  $w_1$  is a way to F. This is the sense then in which this seeming is

non-perceptual, for it is a seeming that Charlie can have in the absence of his having any particular kind of perceptual experience.<sup>47</sup>

The seeming analysis accords then with our intuitions that the KH claims are correct. Furthermore, the seeming analysis also accords with our intuitions about *the salchow* and *man in a room* cases. There is a series of actions such that it seems to Irina that that series of actions is a way to perform the salchow. But this series of actions is not in fact a way to perform the salchow. The seeming analysis rightly predicts then that Irina does not know how to perform the salchow. And while there is a way for the man in a room to exit the room, it does not seem to the man that that way is a way to exit the room, as he is not even aware yet of this way to exit the room. Furthermore, the seeming analysis accords with our intuitions about the pianist case. For even after her accident, it will still seem to the pianist that that way she used to play the piano is a way to play the piano.

It appears that across a diverse range of cases the seeming analysis accords with our intuitions better than both intellectualism and neo-Ryleanism. Unlike intellectualism, it accords with our intuitions about the lucky light bulb, dogmatic hallucinator and non-dogmatic hallucinator cases. And unlike neo-Ryleanism, the seeming analysis accords with our intuitions about *the salchow*, *man in a room* and unfortunate pianist cases.

Both conditions (c) and (d) of our seeming analysis require some explanation. Starting with (d), why include the parallel of Stanley and Williamson's condition (b) in our new analysis of knowledge-how? As we saw in Chapter 1, Stanley and Williamson include (b) because without this condition their analysis would clearly not describe a sufficient condition for knowing how to F. Intuitively, there can be contexts in which one fails to know how to F, even though there is some way *w* such that one knows that *w* is a way for oneself to F. Likewise, one could

---

<sup>47</sup> As mentioned above, whether this kind of seeming can be had in the absence of Charlie having any particular kind of introspective, imaginative, or mnemonic experience is an issue I wish to remain neutral on here.

presumably fail to know how to F, even though there is some way *w* such that it non-perceptually seems to one that *w* is a way for oneself to F. Stanley and Williamson's condition (b) is intended to be a solution to this problem. Insofar as this fix works for their intellectualist account of knowledge-how, the same fix will work for our seeming analysis of knowledge-how. If practical modes of presentation cannot be used to solve this problem then we could appeal to other intellectualist strategies for addressing the same issue, like Bengson and Moffet's strategy of claiming that one knows how to F only if one minimally understands some way to F.

Obviously, with regard to (c), one can know how to F even when it does not *occurrently* seem to one that some way *w* is a way to F—for example, when one is asleep. The seeming analysis will not be plausible then unless one can satisfy (c), even when it does not occurrently seem to one that some way *w* is a way to F. But the idea that an ascription of the form 'It seems to S that p' could be true of S even when it does not occurrently seem to S that p, might initially appear strange. Typically, philosophers are interested in such ascriptions only insofar as they refer to occurrent and conscious seemings. For example, philosophers concerned with the nature of philosophical intuition are often concerned with the kind of occurrent state one is in when, on considering a Gettier scenario (say), it seems to one that the subject in this scenario does not know that the relevant proposition is true.

But there is also a natural interpretation of 'It seems to S that p' ascriptions, whereby they can be satisfied by non-occurrent states. Suppose that during a conversation about the intuitions of our friends I assert, "It still seems to Bill that the naïve axiom of set theory is true even though he knows it to be false". In such a context, it is no objection to my claim to point out that Bill is currently in a deep, dreamless sleep. For my claim is naturally interpreted as being satisfied by some standing, or non-occurrent, state of Bill, rather than some occurrent state of it seeming to Bill that the naïve axiom of set theory is true. And presumably, this non-occurrent state is one that consists (at least partly) in the disposition for it to

occurently seem to Bill that naïve axiom of set theory is true, in certain relevant conditions.<sup>48</sup>

Likewise, condition (c) should be understood in such a way that in order to satisfy (c) it suffices that it seem to one that  $w$  is a way to F, in this non-occurrent sense of ‘It seems to S that p’. That is, where for it to non-occurrently seem to one that some way  $w$  is a way to F is to be in a state that consists (at least partly) in the disposition for it to occurently (and non-perceptually) seem to one that  $w$  is a way to F.

The seeming analysis is something of an intermediate position between Ryle’s own account of knowledge-how and intellectualism. The seeming account is related to intellectualism because it claims that knowing how to do something is a matter of either standing in, or being disposed to stand in, an intentional relation to a true proposition. Of course, unlike intellectualism, the relevant relation is not the knowledge-that relation. Rather, it is the relation of it occurently seeming to one that some proposition is true.

The seeming account is related to Ryle’s account of knowledge-how because both accounts appeal to dispositional states. Neo-Ryleanism is often attributed to Ryle,

---

<sup>48</sup> Philosophers who identify philosophical intuitions with seemings typically deny that intuitions are dispositional states. It is worth emphasizing that this denial is consistent with the idea of dispositional seemings appealed to here. I am not claiming that occurrent seemings can be analysed as some kind of dispositional state (indeed, I think it is clear that they cannot be so analysed). Rather, I am claiming that as well as occurrent seemings we can quite naturally talk of dispositional seemings, where these are understood to be dispositions to have an occurrent seeming. A similar distinction applies to the related notion of understanding. As David Hunter (1998: 559–80) notes, there is a natural and useful distinction to be made between states of occurrent understanding and dispositions to be in such occurrent states—what he calls “dispositions to understand.” As with seeming ascriptions, I think it is clear that ascriptions of understanding can be satisfied not only by states of occurrent understanding, but also by non-occurrent states that consist (at least partly) in dispositions to be in occurrent states of understanding.

and indeed some things he says could be seen as supporting such an ascription. Still, it is not clear that Ryle did identify knowing how to F with the corresponding ability to F. What is clear, however, is that Ryle identified knowing how to F with the possession of a complex of dispositions.<sup>49</sup> Similarly, the seeming account allows that one can know how to F in virtue of being in the relevant kind of dispositional state.

In setting out his account of knowledge-how, Ryle mainly appealed to dispositions to perform various kinds of observable actions. In contrast, the seeming account appeals to a disposition to be in a certain kind of conscious and intentional state. In this respect, the seeming account parts ways with at least the letter of Ryle's account of knowledge-how. Exactly how far it departs from the spirit of Ryle's account is difficult to say. While Ryle mainly appeals to behavioural dispositions when setting out his account of knowledge-how, sometimes he also appeals to dispositions to be in certain cognitive or

---

<sup>49</sup> Or, as Ryle (1949: 44) would say, it is a single disposition "the exercises of which are indefinitely heterogenous." Brian Weatherson, in an entry on his blog "Thoughts, Arguments and Rants", claims that Ryle is only committed to this view, and not neo-Ryleanism and, therefore, the standard counterexamples to neo-Ryleanism simply do not apply to Ryle (see, <<http://tar.weatherson.org/2006/07/22/ryle-on-knowing-how/#comments>>). As we saw in Chapter 1 (§1.2), Stanley and Williamson (2001: 411) attribute both this view and neo-Ryleanism to Ryle, claiming that according to Ryle "knowledge-how is ability, which is in turn a complex of dispositions." That is, they take Ryle to be committed to both of the following identity claims: (i) to know how to F is to possess the ability to F; and (ii) to know how to F is to possess a complex of dispositions. This explains why Stanley and Williamson take their counterexamples to (i) (like *the pianist* case) to be counterexamples to Ryle's account of knowledge-how. Like Weatherson, I am not convinced that Ryle is committed to (i), but even if he is, it seems to me that he would lose little if in response to the standard counterexamples to (i) he were to simply reject (i) whilst retaining (ii).

phenomenal states. Ryle's account of knowledge-how then is more complex than his reputation as a philosophical behaviourist would suggest.<sup>50</sup>

Obviously, much work remains to be done to develop this alternative account of knowledge-how. The role of the seeming account here is only to illustrate the possibility of promising alternatives to both intellectualism and neo-Ryleanism. In the literature, intellectualism and neo-Ryleanism are normally the only accounts of knowledge-how that are discussed. As I mentioned in Chapter 1, this situation can lead to a tendency to regard arguments against either account as being arguments, by default, for the other. The seeming account emphasizes the fact that we should not regard the arguments against intellectualism given here as being arguments for neo-Ryleanism. Furthermore, it shows us that even if knowledge-how is not a kind of knowledge-that, as I have argued, it could still be the case that knowledge-how is propositional in nature. In looking beyond the standard dichotomy of intellectualism and neo-Ryleanism, we may just find a more adequate account of knowledge-how.

However, as the main purpose of this dissertation is to examine the intellectualist view of knowledge-how, I will not consider the possibility of alternative views of knowledge-how in any further detail here. Rather, in the final chapter I want to consider the most famous objection to the intellectualist view of knowledge-how; namely, the objection that intellectualism must be false because assuming it to be true leads to an infinite and vicious regress.

---

<sup>50</sup> For discussion of these complexities, in particular with regard to Ryle's dispositional account of belief, see Eric Schitzwgebel (2002: 259–60), and also Brian Weatherson (forthcoming). Whether these complexities in Ryle's views are consistent with his claims to have undermined "the dogma of the Ghost in the Machine" (Ryle 1949: 15–16) is an interesting question that is obviously beyond the scope of our discussion here.

## Chapter 4 Regarding a Regress

---

In the previous chapter I presented three new arguments against intellectualism. In this final chapter I want to examine the most prominent of existing objections against intellectualism, namely, the objection that there is a successful regress argument against this view of knowledge-how. Such an argument asks us to assume, for the purposes of *reductio*, that intellectualism, or some thesis entailed by intellectualism, is true. It then purports to show that from this assumption, and (presumably) certain other premises, one can derive an infinite regress. From this initial conclusion it is then inferred that intellectualism is false. I assume that for this type of argument to succeed the inference steps that generate the infinite regress must be valid; any required premises (other than the *reductio* premise) must be sound; and the regress has to be intuitively vicious rather than merely benign.

Our search for such an argument begins in §4.1 with an examination of Gilbert Ryle's (1946, 1949) regress argument against what he called *the intellectualist legend*. This argument is the source of the idea that there is some successful regress argument against intellectualism. All existing attempts to explicate this kind of argument are either inspired by Ryle's argument or are offered as reconstructions of his argument. In §4.2 I examine a well-known regress argument against intellectualism that Stanley and Williamson (2001) offer as a reconstruction of what they take to be Ryle's argument against intellectualism—what I will call *the contemplation regress argument*. In §4.3 I construct another regress argument against intellectualism that draws its inspiration from Ryle's argument against the intellectualist legend—what I call *the employment regress argument*. And in §4.4 I consider what reasons an intellectualist might offer for rejecting the key premise that this argument relies upon. In §4.5 I examine regress arguments against intellectualism offered by Stephen Hetherington (2006) and Noë (2005). In §4.6 I discuss some general features of the regress arguments considered here and, in

light of the previous discussion, I answer the question of whether there is some successful regress argument against intellectualism.

#### 4.1 Ryle's Regress

Gilbert Ryle discusses the nature of knowledge-how in two works titled "Knowledge How": the famous second chapter of *The Concept of Mind* (Ryle 1949) and a less well-known paper (Ryle 1946). In each of these works Ryle's primary concern is to examine the nature of intelligence and intelligent actions. For Ryle this examination is closely tied to an examination of the 'mental-conduct concepts' or 'intelligence epithets' that we apply to both persons and actions, which include adjectives like: "clever', 'sensible', 'careful', 'methodical', 'inventive', 'mental', 'quick-witted'" (Ryle 1949: 26). And, on the other hand, there are adjectives that are used to indicate that a person is deficient in intelligence or that an action is not performed intelligently, like: "stupid, 'dull', 'silly', 'careless', 'unmethodical', 'uninventive'" (*ibid.*).

Ryle describes the intellectualist legend as a view concerning these intelligence epithets that is widely accepted but also fundamentally in error. As Snowdon (2004: 15–16) points out, Ryle's discussion is made rather confusing by the way he moves back forth between questions about what it is for *persons* to be intelligent (or clever, or skilful etc.), and what it is for *actions* to be intelligent (or clever, or skilful etc.). Sometimes Ryle describes the intellectualist legend as a view about what it is for a person to be intelligent, or as a view about what the 'faculty' of intelligence is. But primarily, Ryle construes it as a view about the nature of intelligent actions, and it is this view that his regress argument is directed at.

Ryle gave many different glosses of the intellectualist legend all of which unfortunately are somewhat vague, with different terminology often being used to express seemingly related but still importantly distinct concepts. One of Ryle's (1949: 29) most well-known statements of the intellectualist legend occurs after the following key passage:

What is involved in our descriptions of people as knowing how to make and appreciate jokes, to talk grammatically, to play chess, to fish or to argue? Part of what is meant is that, when they perform these operations, they tend to perform them well, i.e. correctly or efficiently or successfully. Their performances come up to certain standards, or satisfy certain criteria. But this is not enough. The well-regulated clock keeps good time and the well-drilled circus seal performs its tricks flawlessly, yet we do not call them ‘intelligent’. We reserve this title for the persons responsible for their performances. To be intelligent is not merely to satisfy criteria, but to apply them; to regulate one’s actions and not merely to be well-regulated. A performance is described as careful or skilful, if in his operations he is ready to direct and correct lapses, to repeat and improve upon successes, to profit from the examples of others and so forth. He applies criteria in performing critically, that is, in trying to get things right (Ryle 1949: 29).

I take it that Ryle’s main aim in this passage is to motivate a distinction between merely successful actions and intelligent or skilful actions. And Ryle uses this distinction to characterize both what it is to know how to perform some action and what it is for a person to be intelligent. Ryle suggests that part of what we mean when we say that someone knows how to F, is not only that they tend to F successfully, but that they tend to F intelligently or skilfully. Similarly, it appears that Ryle thinks of an intelligent person as someone who not only tends to perform successful actions, but someone who tends to perform actions that are intelligent, or careful, or skilful etc.

The examples Ryle appeals to in support of this distinction—between merely F-ing successfully and F-ing intelligently—are perhaps not that helpful.<sup>51</sup> However, the distinction itself is a perfectly fine one. Consider the action of hitting the bull’s eye in darts. One can obviously succeed in hitting the bull’s eye and yet fail

---

<sup>51</sup> For one thing, I think it would be both natural and correct to use adjectives like ‘intelligent’ ‘skilful’ or ‘clever’ to describe either the seal’s actions or the seal itself. The fact that the seal has been ‘well-drilled’ so as to perform its tricks does not seem to be a reason to deny that its actions of performing those tricks are skilful or clever. An Olympic gymnast has also been well-drilled to perform her floor routine, but her action of performing this routine is skilful nonetheless.

to do so intelligently or skilfully; for example, when one hits the bull's eye when one was trying to hit the twenty.

What more is required for one to not only successfully perform some action but to also perform that action intelligently? Ryle's suggestion appears to be that to F intelligently one not only has to satisfy criteria—by which I take it he means criteria for F-ing successfully—but one has to also apply those criteria in one's action of F-ing, or one has to regulate one's actions of F-ing so that it satisfies these criteria. Ryle's use of this notion of applying criteria is somewhat elusive, but it appears that the general idea he has in mind here is just the obvious one that to F intelligently one not only has to succeed in F-ing, but one has to also—in some sense—be *responsible* for this success. Ryle claims that this point is commonly expressed 'in the vernacular' by saying:

[T]hat an action exhibits intelligence if, and only if, the agent is thinking what he is doing while he is doing it, and thinking what he is doing in such a manner that he would not do the action so well if he were not thinking what he is doing. This popular idiom is sometimes appealed to as evidence in favour of the intellectualist legend (Ryle 1946: 30).

Ryle's claim that this analysis of intelligent action is a "popular idiom" seems rather far-fetched—for one thing the construction 'if and only if' is rarely used by anyone other than philosophers and logicians. But, more importantly, what does it mean to say that an agent is 'thinking what he is doing while doing it'? Should we interpret this as shorthand for the claim that the agent is thinking *about* what he is doing while doing it? Ryle does not say. It is true that when someone fails to perform an action skilfully or carefully we will sometimes explain this failure by saying that they were 'acting without thinking' or that 'they were not thinking about what they were doing'. But it is not at all clear that we are thereby expressing a commitment to Ryle's popular idiom.

These difficulties in interpreting Ryle's claims about this popular idiom are not important for our purposes, however. What is important is the structure of this popular idiom, for Ryle goes on to present the intellectualist legend as a view that

offers a reinterpretation of the right-hand side of this biconditional claim. As Ryle continues:

This popular idiom is sometimes appealed to as evidence in favour of the intellectualist legend. Champions of this legend are apt to try to reassimilate knowing *how* to knowing *that* by arguing that intelligent performance involves the observance of rules, or the application of criteria. It follows that the operation which is characterized as intelligent must be preceded by an intellectual acknowledgement of these rules or criteria; that is, the agent must first go through the internal process of avowing to himself certain propositions about what is to be done ('maxims', 'imperatives', or 'regulative propositions' as they are sometimes called); only then can he execute his performance in accordance with those dictates. He must preach to himself before he can practise. The chef must recite his recipes to himself before he can cook according to them; the hero must lend his inner ear to some moral imperative before swimming out to save the drowning man; the chess-player must run over in his head all the relevant rules and tactical maxims of the game before he can make correct and skilful moves. **To do something thinking what one is doing is, according to this legend, always to do two things; namely, to consider certain appropriate propositions, or prescriptions, and to put into practice what these propositions or prescriptions enjoin.** It is to do a bit of theory and then to do a bit of practice. ... I shall argue that the intellectualist legend is false and that when we describe a performance as intelligent, this does not entail the double operation of considering and executing (Ryle 1949: 29–30; bold emphasis added).

According to Ryle's popular idiom, one's action F-ing is performed intelligently if and only if one is thinking what one is doing while F-ing, and thinking what one is doing in such a manner that one would not F so well if one were not thinking what one is doing. The passage above suggests that, according to the intellectualist legend, we can replace the condition that one's action of F-ing is performed intelligently only if one is thinking what one is doing while F-ing, with the condition that one's action F-ing is performed intelligently only if prior to F-ing one contemplates some relevant proposition(s) concerning something like a way or procedure for F-ing.

What of the condition in the popular idiom, that one's action F-ing is performed intelligently only if one is thinking what one is doing *in such a manner that one would*

*not F so well if one were not thinking what one is doing?* The passage above suggests that proponents of the intellectualist legend hold that one's act of F-ing has to be the *execution* or *application* of the way described by the proposition(s) that one contemplates. Elsewhere, Ryle indicates that proponents of the intellectualist legend hold that one's action of F-ing has to be both 'introduced' and 'steered' by one's contemplating. For example, he describes the intellectualist legend as the view that:

[P]ractical activities merit their titles 'intelligent', 'clever', and the rest only because they are accompanied by some such internal acts of considering propositions (and particularly 'regulative' propositions). That is to say, doing things is never itself an exercise of intelligence, but is, at best, a process **introduced and somehow steered by some ulterior act of theorising** (Ryle 1946: 212; bold emphasis added).

I think the intellectualist legend that Ryle opposes then can be represented as a view of intelligent actions that is either equivalent to, or at least entails, the following equivalence claim, where the square parentheses indicate some of the variations in the terminology Ryle uses to express this view:

One performs an action F intelligently [or skilfully, or cleverly, carefully etc.] if and only if one Fs and one's action of F-ing is introduced and guided [or steered] by one's contemplating [or considering, or acknowledging] of some relevant proposition(s) concerning a way to F [or a rule, or procedure for F-ing] (and this contemplating is distinct from and prior to one's F-ing).

We cannot avoid altogether the variations in the terminology Ryle uses when describing the intellectualist legend. But it will be useful to focus on just one version of the biconditional stated above. And for the sake of brevity, I will also replace 'introduced and guided' with just 'guided' (so in the following discussion 'guided' should be read as shorthand for 'introduced and guided'). I will also drop the condition, stated in the bracketed clause, that the relevant contemplating must occur *prior* to one's action of F-ing. Ryle often describes the intellectualist legend

as including this condition, but sometimes he acknowledges that a proponent of the intellectualist legend could hold that this contemplating and one's action of F-ing occur at the same time. And Ryle's regress argument against the intellectualist legend does not need to appeal to this condition. I will take the intellectualist legend then that Ryle opposes to be the following equivalence claim that I will refer to as simply *the legend*:

### The Legend

One performs an action F intelligently if and only if one Fs and one's action of F-ing is guided by one's contemplating of some relevant proposition(s) concerning a way to F and one's action of F-ing (and this contemplating is distinct from one's F-ing).

As I will show, Ryle rejects both directions of this biconditional. But he was particularly concerned to establish that the left to right direction of the legend is false, as he wanted to maintain that some actions are intelligent even though they are not preceded by any 'shadow act' of contemplating propositions. That is, Ryle wanted to show that having one's action of F-ing be guided by one's prior contemplation of certain propositions is not a *necessary* condition for one's action of F-ing to be an intelligent action.

It is important to note that Ryle does not reject the idea that some actions are intelligent as a result of their being guided by some distinct action of contemplating a proposition; as Ryle (1946: 30) says, "Certainly we often do not only reflect before we act but reflect in order to act properly." What Ryle rejects is the idea that *all* intelligent actions are guided by the contemplating of certain relevant propositions, or that "all intelligent performance requires to be prefaced by the consideration of appropriate propositions" (*ibid*).

Indeed, I think Ryle would even allow that actions that are successful as a result of one's prior reflecting on how to perform them are paradigmatic instances of intelligent or skilful actions. For example, consider a kind of scenario familiar from many an action film; namely, a scene where the hero is faced by a corridor

in which a series of mechanical guillotines are going up and down at different rates. The hero, who has to get through the corridor, pauses to take careful note of the sequence in which the different blades fall, and then he makes his way through the corridor, stopping and starting at just the right points and for just the right amount of time, so that he is not chopped in two. The hero's action here is a prime example of a skilful or intelligent action, and his performance of this action is, in some sense, guided by his prior reflection on, or contemplation of, propositions concerning a way to get through the corridor in one piece.

Now Ryle need not deny the description of this case. I take it that Ryle's point is simply that it would be a mistake to infer from the fact that our hero's action is a paradigmatic example of an intelligent action that it is an essential feature of *any* intelligent action that it be guided by the agent's prior contemplation of relevant propositions. And Ryle's point is surely reasonable, especially given that there are other equally clear cases of intelligent or skilful actions where it does not seem very plausible that the action is guided by the agent's contemplation of a proposition concerning a way to perform it. Consider, for example, a skilful behind-the-back pass by Chris Paul in a basketball game. Unlike the hero's action, the idea that in making this pass Chris Paul is guided by his contemplation of a proposition concerning a way to perform a behind-the-back pass is *prima facie* implausible. But, of course, it is one thing to point out that a view is implausible in some respect and it is another to show that assuming it to be true leads to an infinite and vicious regress. So, what is Ryle's regress argument against this legend?

*Ryle's regress argument: version 1*

The following passage from Ryle is the one that is typically used to illustrate his regress argument against the legend:

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. (Ryle 1946: 31)

If we consider this passage, on its own, it seems that the regress argument Ryle envisages here is generated from something like the following two premises:

(1) If one performs an action F intelligently then one contemplates some proposition p intelligently (and this contemplating is distinct from one's action of F-ing).

(2) To contemplate a proposition is to perform an action.

An infinite regress is then generated like so. Imagine that Ari performs some action F intelligently. By (1) it follows that prior to his act of F-ing Ari contemplated some proposition  $p_1$  intelligently. By (2) it follows that Ari's intelligent contemplating of  $p_1$  was itself an action. But then we can reapply (1) to conclude that prior to his act of contemplating  $p_1$  Ari contemplated some proposition  $p_2$  intelligently. By (2) it follows that Ari's intelligent contemplating of  $p_2$  was itself an action. By (1) it follows that prior to his act of contemplating  $p_2$  Ari contemplated some proposition  $p_3$  intelligently. By (2) it follows that Ari's contemplating of  $p_3$  was itself an action. And so on *ad infinitum*.

So, if (1) and (2) are true then, to perform any action intelligently one would have to also perform an infinite number of distinct actions of intelligently contemplating propositions.<sup>52</sup> Now this conclusion looks to be absurd, for given that we are finite beings, what this conclusion tells us essentially is that we never perform intelligent actions. But while this conclusion is absurd, the inferences that generate this regress are valid, therefore, we must reject either (1) or (2). Ryle

---

<sup>52</sup> Note that it does not follow from (1) and (2) that these distinct acts of contemplating propositions would have to involve the contemplation of distinct propositions, that is, it is consistent with (1) and (2) that  $p_1 = p_2$  and  $p_2 = p_3$  and so on.

assumes that the right response is to reject (1), as he seems to regard (2) as being obviously true.<sup>53</sup>

However, even granting to Ryle that the right response to this regress is to reject (1), his regress argument appears vulnerable to a simple objection; namely, that the conclusion that (1) is false is consistent with the legend being true. For (1) is not entailed by the left-to-right direction of the biconditional expressed by the legend; that is, (1) is not entailed by the claim that if one performs an action F intelligently then one's action of F-ing is guided by one's contemplation of some relevant proposition concerning a way to F. As (1) says, if one performs an action intelligently then one contemplates some proposition *intelligently*. But from the legend, it only follows that if one performs an action intelligently then one contemplates some proposition. A proponent of the legend could then reply to Ryle's regress argument by pointing out that it only shows us that if the legend is true it must be the case some actions are intelligent, even though they are guided by actions of contemplating propositions that are not themselves intelligently

---

<sup>53</sup> Ryle does sometimes gesture at an argument for (2), one example is the following passage:

That thinking-operations can themselves be stupidly or intelligently performed is a notorious truth which by itself upsets the assumed equation of 'exercising intelligence' with 'thinking'. Else 'stupid thinking' would be a self-contradictory expression and 'intelligent thinking' would be a tautology. It also helps to upset the assumed type-difference between thinking and doing, since only subjects belonging to the same type can share predicates. But thinking and doing do share lots of predicates, such as 'clever', and 'stupid', 'careful', 'strenuous', 'attentive', etc (Ryle 1946: 213).

The argument suggested by this passage is something like this: (i) if the same predicates apply both to practical operations and thinking operations then they are operations of the same type; (ii) the same predicates apply to both practical and thinking operations; and therefore, (iii) practical and thinking operations are operations of the same type. From this initial conclusion, Ryle infers that theoretical operations are a kind of action. It is also worth noting that Ryle often suggests that proponents of the intellectualist legend deny (2): "it is also assumed that theorising is not a sort of doing, as if 'internal doing' contained some contradiction" (Ryle 1946: 212).

performed. And, as Ryle's regress argument gives us no reason to think that this could not be the case, his argument fails to establish that the legend is false.

Snowdon (2004: 15) and Scott Soames (2005: 102) both raise variations of this simple objection to Ryle's regress argument. For example, Soames claims that all that Ryle's argument establishes is the following disjunction:

[e]ither some performances are intelligent by virtue of intellectual preplanning (or the application of theoretical knowledge) that is not itself intelligent, **or** some performances are intelligent but not by virtue of intellectual preplanning (or the application of theoretical knowledge) at all (Soames 2005: 102; bold emphasis is original).

However, as Soames notes, what Ryle wants to conclude is that the second disjunct in the above disjunction is true—that some performances are intelligent but not by virtue of any intellectual preplanning. To paraphrase Soames' point (so as to reflect the statement of the legend earlier), it seems that Ryle's regress argument merely establishes that either some actions are intelligent even though they are guided by a non-intelligent action of contemplating a proposition, or some actions are intelligent even though they are not guided by any action of contemplating a proposition.

*Ryle's regress argument: version 2*

Does this simple objection to Ryle's regress argument succeed? I think this objection actually misrepresents Ryle's reasoning when he presents his regress argument. To see why, let us replace (1) with our statement of the legend, whilst retaining (2):

### The Legend

One performs an action F intelligently if and only if one Fs and one's action of F-ing is guided by one's contemplating of some relevant proposition(s) concerning a way to F and one's action of F-ing (and this contemplating is distinct from and prior to one's F-ing).

(2) To contemplate a proposition is to perform an action.

I think that Ryle's regress argument against the legend is really an argument by dilemma, where only one of the horns of the dilemma is an infinite and vicious regress. The argument proceeds like so: either the contemplating referred to in the legend must be intelligently performed or it need not be. If it must be intelligently performed, then (as we have seen) the left-to-right direction of the legend, together with (2), will generate an infinite and vicious regress. Assuming that (2) is true then, we must reject the legend. On the other hand, Ryle would claim that if the contemplating referred to in the legend need not be intelligently performed then there is no reason to believe that the right-to-left direction of the legend is true. For if one's action of F-ing is guided by, for example, one's stupid or careless contemplating of relevant propositions concerning a way to F, then one will in all likelihood fail to F intelligently.

To support this interpretation of Ryle, consider this passage which follows shortly after the famous statement of his regress argument, where Ryle is now considering "some salient points at which this regress would arise":

Next, supposing still that to act reasonably I must first perpend the reason for so acting, how am I led to make a suitable application of the reason to the particular situation which my action is to meet? For the reason, or maxim, is inevitably a proposition of some generality. It cannot embody specification to fit every detail of the particular state of affairs. Clearly, once more I must be sensible and not stupid, and this good sense cannot itself be a product of the intellectual acknowledgement of any general principle. A soldier does not become a shrewd general merely by endorsing the strategic principles of Clausewitz; he must also be competent to apply them (Ryle 1949: 31–32).

I take it that Ryle's point here is that if one is to F intelligently as a result of one's contemplating propositions concerning a way F, then one had better exercise intelligence both when contemplating these propositions, and when applying their content in action. For otherwise, one might still fail to F intelligently even though one's action of F-ing is guided by one's contemplating of relevant propositions concerning a way to F.

Ryle's writings on knowledge-how are full of cases—like the soldier case—that purport to illustrate the existence of what he calls a 'gap' or 'gulf' "between acknowledging principles in thought and intelligently applying them in action" (1946: 218). Ryle's soldier can make the wrong tactical decisions even though he always contemplates the relevant strategic principles of Clausewitz before making these decisions. For if the soldier does not exercise intelligence when he contemplates these principles, he might not see how to apply them to his particular situation. Similarly, Ryle (1946: 215–16) uses the example of a 'stupid chess-player' who knows all the same maxims or propositions concerning chess strategies that are known to the 'clever chess-player', but who is unable to intelligently apply them:

What facts or what sorts of facts are known to the sensible which are not known to the silly? For example, what truths does the clever chess-player know which would be news to his stupid opponent? Obviously there is no truth or set of truths of which we could say, 'If only the stupid player had been informed of them, he would be a clever player,' or 'When once he had been apprised of these truths he would be play well.' We can imagine a clever player generously imparting to his stupid opponent so many rules, tactical maxims, 'wrinkles', etc. that he could think of no more to tell him; his opponent might accept and memorise all of them, and be able and ready to recite them correctly on demand. Yet he might still play chess stupidly, that is, be unable intelligently to apply the maxims, etc (Ryle 1946: 215).

Ryle (1946: 216) claims that his stupid chess-player, who knows all the right propositions concerning chess strategies, could contemplate the right proposition concerning a chess strategy whenever he tries to make a good move in a game of chess and yet still fail to make a good move because "he might not see that it was the appropriate maxim or if he did he might not see how to apply it."

What is the import of such examples? At the very least, I think Ryle takes such cases to support something like the following claim:

#### The Gap Premise

One can fail to perform an action F intelligently even though one Fs and

one's action of F-ing is guided by one's contemplation of some relevant proposition concerning a way to F.

The gap premise simply tells us that one can fail to perform an action F intelligently even when one Fs and one's action of F-ing is guided by one's contemplation of true and relevant propositions concerning a way to F. In particular, Ryle thinks this can happen when the action of contemplating is not itself performed intelligently. If the gap premise is true, then the legend is false because the gap premise entails that the right-to-left-direction of the legend is false. And, as we have seen, a proponent of the legend cannot respond to this objection by simply stipulating that the contemplating of propositions referred to in the legend must be intelligently performed. For this response, together with the assumption that to contemplate a proposition is to perform an action, leads to an infinite and vicious regress.

*Responses to Ryle's regress argument*

I think it is clear that, once properly understood, Ryle's regress argument against the legend succeeds. But there are related positions that a proponent of the legend could still retreat to. For example, they could accept the gap premise, whilst still avoiding an infinite and vicious regress, by separating the two conditionals embedded in the legend and instead endorsing the following two claims:

(LR) If one performs an action F intelligently then one Fs and one's action of F-ing is guided by one's contemplating of some relevant proposition concerning a way to F.

(RL\*) If one Fs and one's action of F-ing is guided by one's *intelligent* contemplating of some relevant proposition concerning a way to F then one Fs intelligently.

In other words, the proponent of the intellectualist legend could respond to Ryle's gap premise by inserting 'intelligently' into the antecedent of the right-to-

left conditional expressed by the legend—giving us (RL\*)—whilst avoiding an infinite and vicious regress by not inserting ‘intelligently’ into the consequent of the left-to-right conditional—i.e. (LR).

I think this response reveals the grain of truth in the simple objection to Ryle’s argument. For what this response shows is that Ryle’s regress argument does not demonstrate that the key idea expressed by (LR) is false, that is, the idea that all intelligent actions are guided by the contemplating of propositions. If Ryle’s gap premise is true it follows that *sometimes* one will fail to F intelligently when one’s act of F-ing is guided by one’s non-intelligent contemplating of some proposition. But it does not follow that one will *always* fail to F intelligently when one’s action of F-ing is guided by one’s non-intelligent contemplating of some proposition. In which case, it could still be true that some intelligent actions are guided by non-intelligent actions and, therefore, one could hold that all intelligent actions are guided by the contemplation of propositions whilst avoiding Ryle’s regress argument.

I suspect that Ryle would find the idea that an action could be intelligent if it were guided by a non-intelligent action of contemplating a proposition to be deeply implausible. It seems that Ryle holds that if one’s action of F-ing is guided by one’s non-intelligent contemplating of propositions, it is not only the case that one *could* but that one *must* fail to F intelligently. That is, I think Ryle holds that it is a precondition of one’s F-ing intelligently when one’s action of F-ing is guided by one’s contemplating of propositions, that this contemplating be performed intelligently. For example, Ryle (1946: 216) appears to endorse this idea when he concludes on the basis of the stupid chess-player case that: “it *requires* intelligence not only to discover truths, but also to apply them” (emphasis added).

However, I think one might reasonably object here that all that Ryle’s examples clearly demonstrate is that if one’s action of F-ing is guided by some *stupid* or *careless* action of contemplating a proposition, then one will not F intelligently. But these examples do not establish that if one’s action of F-ing is guided by some *non-intelligent* action of contemplating a proposition, then one must fail to F

intelligently. To claim otherwise, according to this objection, would be to conflate an action's being non-intelligent with it being stupid or careless etc.

Ryle does not demonstrate then that (LR) is false. But Ryle's regress argument does undermine the legend. And, in showing that the legend is false, Ryle demonstrates that one cannot simply analyse intelligent actions as actions that are guided by one's contemplating of certain relevant propositions.

### *The legend and intellectualism*

Ryle's regress argument against the legend succeeds. But there is still a serious problem with this argument. The problem, given our interests, is simply that Ryle's regress argument is not an argument against intellectualism, where 'intellectualism' is the view that knowledge-how is a kind, sort, or species of knowledge-that. For the legend that Ryle's regress argument targets is a view about the nature of intelligent actions, not the nature of knowledge-how.

While Ryle made famous the idea that there is a successful regress argument against intellectualism, the irony is that he never explicitly stated such an argument himself. It is true that Ryle frequently intimates that his regress argument against the legend somehow also supports the conclusion that knowledge-how is not a kind of knowledge-that, or that one cannot define "knowing how" in terms of "knowing that" (1949: 32). But it is not clear exactly why he thought this. One explanation would be that Ryle thought that intellectualism entailed the legend, in which case any argument that showed that the legend is false would thereby entail that intellectualism is false. But clearly, the view that knowledge-how is a kind of knowledge-that does not entail the rather odd view of intelligent actions expressed by the legend.

Perhaps a more charitable explanation is that a regress argument against intellectualism is implicitly suggested by Ryle's critique of the legend. This would also explain why so many philosophers have credited him with providing such an argument, when he only ever explicitly states a regress argument against the legend. One might hope then that we can reconstruct Ryle's implicit regress

argument against intellectualism, and that this argument will turn out to be a successful argument against intellectualism. Stanley and Williamson (2001) have provided the most well known reconstruction of Ryle's supposed regress argument against intellectualism. I will now examine the argument they present before going on to identify another form of regress argument against intellectualism, which is also related to Ryle's argument against the legend.

## 4.2 The Contemplation Regress

The argument Stanley and Williamson offer as a reconstruction of what they take to be Ryle's regress argument against intellectualism relies on two premises, what I will call *the action premise* and *the contemplation premise*:

### The Action Premise

If one Fs, one employs knowledge how to F.

### The Contemplation Premise

If one employs knowledge that p, one contemplates the proposition that p.

Of course, for the sake of the argument, we also need to assume that intellectualism is true or that some thesis entailed by intellectualism is true. Stanley and Williamson (2001: 413–14) claim that: “If knowledge-how is a species of knowledge-that, the content of knowledge-how to F is, for some  $\Phi$ , the proposition that  $\Phi(F)$ ”, where  $\Phi$  is some function that maps acts to propositions, and so  $\Phi(F)$  is the proposition that is the value of the function  $\Phi$  when the input to that function is the action F. They reconstruct Ryle's regress argument against intellectualism as an argument against the following identity claim, which they call *the reductio assumption*:

### The Reductio Assumption (RA)

Knowledge how to F is knowledge that  $\Phi(F)$ .

Note that (RA) is essentially a version of the simple identity thesis that we saw in Chapter 1:

The Simple Identity Thesis

To know how to F is to know that p (for some relevant proposition p concerning a way to F).

The only difference between (RA) and the simple identity thesis is that (RA) does not include the condition that the proposition in question has to concern a way to F. Stanley and Williamson describe how Ryle's regress argument against (RA) is meant to proceed like so, where 'C(p)' stands for the act of contemplating some proposition p:

Suppose that Hannah Fs. By [the action premise], Hannah employs the knowledge how to F. By RA, Hannah employs the knowledge that  $\Phi(F)$ . So, by [the contemplation premise], Hannah C( $\Phi(F)$ )s. Since C( $\Phi(F)$ ) is an act, we can reapply [the action premise], to obtain the conclusion that Hannah knows how to C( $\Phi(F)$ ). By RA, it then follows that Hannah employs the knowledge that  $\Phi(C(\Phi(F)))$ . By [the contemplation premise], it follows that Hannah C( $\Phi(C(\Phi(F))))$ s. And so on.

Ryle's argument is intended to show, that, if [the action premise] and [the contemplation premise] are true, then, if knowledge-how is a species of knowledge-that, doing anything would require contemplating an infinite number of propositions of ever increasing complexity (Stanley and Williamson 2001: 414).

This then is the regress argument against intellectualism that Stanley and Williamson attribute to Ryle.<sup>54</sup> They point out (*ibid*) that if this argument is to

---

<sup>54</sup> Stanley and Williamson (*ibid*) also note that Ryle himself would have presumably endorsed an even stronger version of the contemplation premise; that the "employment of knowledge-that requires a *prior* action of contemplating a proposition." For recall that Ryle characterises the intellectualist legend as saying that an action is intelligent just in case it is guided by a *prior* act of contemplating a relevant proposition. However, as Stanley and Williamson point out, this stronger version of the contemplation premise is not needed as the conclusion that to engage in any action "it is necessary to contemplate an infinite number of distinct propositions" (*ibid*) is

succeed at least two further premises are also required: (i) that the function  $\Phi$  maps distinct acts to distinct propositions; and (ii) that  $C(p)$  is a distinct act from  $C(\Phi(C(p)))$ , which is a distinct act from  $C(\Phi(C(\Phi(C(p)))))$ , and so on.<sup>55</sup>

Now, one could always debate whether or not this argument can be legitimately attributed to Ryle, given that Ryle himself only explicitly states a regress argument against the legend and not intellectualism. But such interpretative issues are not my concern here. My concern is simply whether this argument succeeds or not. Accordingly, from here on I will refer to this argument as *the contemplation regress argument* against intellectualism, so as to remain neutral on the issue of whether it should be attributed to Ryle. The question we have to address then is this: is the contemplation regress argument a successful regress argument against intellectualism?

---

itself surely false. Ryle does not need the stronger conclusion, that to engage in any action one would have to perform an infinite number of distinct actions of contemplating propositions performed over an infinite time span.

<sup>55</sup> One also has to assume that the infinite regress of act of contemplating propositions does not 'loop' back on itself. For even if every member of an infinite series is distinct from its immediate predecessor in that series, it could still be the case that every member of that infinite series is identical to some other member of the infinite series. To illustrate the point, suppose we have shown that some thesis, if true, generates an infinite regress of actions:  $A_1, A_2, A_3, A_4$  and so on *ad infinitum*, where each action in the series is distinct from its predecessor so  $A_1 \neq A_2$ , and  $A_2 \neq A_3$ , and  $A_3 \neq A_4$  and so on *ad infinitum*. Now it may appear that a commitment to such a regress commits one to the existence of an infinite number of distinct actions. But this is not quite right. For example, it could be that after  $A_1$  and  $A_2$  every 'new' action in the regress is identical to either  $A_1$  or  $A_2$  as follows:  $A_3 = A_1, A_4 = A_2, A_5 = A_1, A_6 = A_2, A_7 = A_1$  and so on *ad infinitum*. This is of course possible, because unlike identity, non-identity is not a transitive relation. And if this were the case, our regress of actions would be intuitively benign rather than vicious, as it would only commit us to the existence of exactly two distinct actions, rather than an infinite number of distinct actions.

*Stanley and Williamson's Critique of the Contemplation Regress*

The contemplation regress argument relies on two key premises. The action premise tells us that one Fs only if one employs knowledge how to F. The contemplation premise tells us that one employs knowledge that p only if one contemplates p. Stanley and Williamson argue that the contemplation regress argument is unsound on the grounds that there is no interpretation of these two premises such that they are both plausibly true, and we can derive an infinite regress from these two premises and RA. They reach this conclusion in three steps.

*Step 1: Restrict the Action Premise*

The first step in Stanley and Williamson's critique is to point out that the action premise is clearly false for many values that we could give to 'F'. For example, consider the following claim:

If Hannah digests food, she knows how to digest food.

As Stanley and Williamson (2001: 414) point out, this claim is false because: "Digesting food is not the kind of thing that one knows how to." They also offer the example of Hannah, who wins a fair lottery, but who did not know how to win a fair lottery, since she only won the lottery by sheer chance.

One might worry that the digestion example is not a counterexample to the action premise because it is not something that we do; rather, digesting is an involuntary process that occurs inside our bodies.<sup>56</sup> But I take it that Stanley and Williamson's point is simply that *grammatically* speaking, digesting is something that we *do*. In which case, 'Hannah digests food' is a legitimate value for 'F' in the action premise, and this is why the case is a counterexample to the claim that if one Fs then one employs knowledge how to F.

---

<sup>56</sup> Noë (2005: 279) raises this kind of concern about this case.

Similarly, sweating is presumably something we should classify as a mere involuntary bodily process or activity—like the beating of our hearts—and not as an action that we as agents perform. But grammatically speaking, sweating—unlike the beating of our hearts—is something that we do, and so at least in this very limited sense of the word ‘action’, it is an action. Hence, the claim that if David sweats then he employs knowledge how to sweat, is a counterexample to the action premise because the claim is clearly false, and yet ‘David sweats’ is a legitimate value for ‘F’ in the action premise.

The digesting and sweating cases are counterexamples to the action premise because these ‘doings’ or ‘actions’ are just not the kind of thing that one knows how to do. But there will also be cases that are counterexamples to the action premise where one does know how to F but one still Fs without employing this knowledge-how. For example, I know how to knock the vase off the mantelpiece, but when I do so accidentally I do not employ this knowledge-how.

According to Stanley and Williamson (2001: 415), the lesson of such counterexamples to the action premise is that this premise is only correct if we restrict the range of actions that can be values for ‘F’ to *intentional actions*. That is, they claim that the action premise is false but that the following claim is true:

The Intentional Action Premise

If one Fs *intentionally*, one employs knowledge how to F.

As discussed in Chapters 2 and 3, this intentional action premise is actually an important component of Stanley and Williamson’s full account of knowledge-how. So, they do acknowledge that there is an important constitutive or conceptual connection between knowledge-how and action, but only once we restrict our attention to intentional actions. And, as they point out, neither the digestion case nor the lottery case is a counterexample to the intentional action premise, as neither of these actions are actions that Hannah does intentionally. And the same point obviously applies to David’s sweating or one’s action of accidentally knocking the vase off the mantelpiece.

Before proceeding to the second step in Stanley and Williamson's critique, it is worth mentioning that one could presumably appeal to alternative restrictions on the range of actions that can be values for 'F' in the action premise, as a means of avoiding such counterexamples. Indeed, as Stanley and Williamson themselves point out, Ryle himself appears to endorse the following version of the action premise that is restricted to *intelligent* rather than *intentional* actions:

The Intelligent Action Premise

If one Fs *intelligently* [or cleverly, or carefully etc.], one employs [or applies] knowledge how to F.

Ryle appeared to think that only intelligent actions should be analysed as actions that exercise our knowledge-how. In particular, he seems to be committed to some claim of the form: one Fs intelligently if, and only if, one Fs and in F-ing one employs (or as Ryle would say, exercises or applies) one's knowledge how to F. In which case, Ryle would hold that if one Fs intelligently then one employs knowledge how to F.

Stanley and Williamson note that Ryle would endorse the intelligent action premise, and they grant that this restricted version of the action premise also avoids the digestion and lottery counterexamples; for neither Hannah's digesting of her food nor her winning the lottery are actions that she performs intelligently. However, they do not regard this alternative strategy for revising the action premise as being significantly different from their own suggestion that one restrict the action premise to intentional actions. For after mentioning Ryle's alternative restriction Stanley and Williamson (*ibid.* 415) say that: "the range of actions under consideration must be restricted to intentional actions, or perhaps a proper subset thereof." The thought here being that any intelligent action will also be an intentional action.

Note that if Stanley and Williamson are right that all intentional actions are employments of knowledge-how then this strongly suggests that Ryle was wrong to think that there was some special connection between intelligent actions and

knowledge-how. For, sadly, many of our actions that are ‘stupid’, ‘dull’, ‘silly’, ‘careless’, ‘unmethodical’ or ‘uninventive’ etc., are nonetheless actions that we perform intentionally. But if the intentional action premise is true these non-intelligent actions are still all employments of knowledge-how.

*Step 2: Deny the Contemplation Premise*

The first step in Stanley and Williamson’s critique of the contemplation regress argument is to claim the action premise is false unless we stipulate that it only applies to intentional actions. The second step is to claim that the contemplation premise—the claim that if one employs knowledge that *p*, one contemplates the proposition that *p*—is false. In support of this claim Stanley and Williamson (2001: 415) cite the following passage from Carl Ginet:

I exercise (or manifest) my knowledge *that* one can get the door open by turning the knob and pushing it (as well as my knowledge *that* there is a door there) by performing that operation quite automatically as I leave the room; and I may do this, of course, without formulating (in my mind or out loud) that proposition or any other relevant proposition (Ginet 1975: 7).

Stanley and Williamson think that what Ginet’s door example illustrates is that we often exercise or employ our knowledge that *p* without contemplating the proposition that *p*. In which case, it is a mistake to assume that employments “of knowledge-that must be accompanied by distinct acts of contemplating propositions” (2001: 415).

*Step 3: Block a Bad Reply to Step 2 by Appealing to Step 1*

Stanley and Williamson do imagine a way in which someone might try to accommodate Ginet’s door example whilst maintaining that the contemplation premise is correct:

Ginet clearly construes “contemplating a proposition” as referring to an intentional act of contemplating a proposition, which is one natural sense of the phrase. If “contemplating a proposition” is construed in its intentional action sense, then [the

contemplation premise] is false. But we can rescue [the contemplation premise] from Ginet's objection by denying that "contemplating a proposition" should be taken in its intentional action sense in [the contemplation premise]. Perhaps there is a sense of "contemplating a proposition" in which it refers to an action that is no more intentional than is the action of digesting food. Or perhaps it can also be construed as denoting an action in some deflationary sense of "action". If "contemplating a proposition" is taken in such a sense, then [the contemplation premise] can be salvaged after all (Stanley and Williamson 2001: 415–16).

The third and final step then in Stanley and Williamson's critique is to point out that while this response might save the contemplation premise it would not save the contemplation regress argument, given what they have said about the action premise. For if 'contemplates the proposition that p' is interpreted so that it refers to a non-intentional action, then it is not a legitimate substitution for 'F' if the action premise is interpreted so that the range of actions that can be values for 'F' is restricted to just intentional actions or some proper subset thereof. And, they claim, the action premise is only plausible if it is so restricted.

Stanley and Williamson conclude that there is no interpretation of the action premise and the contemplation premise such that both premises are plausibly true and we can derive a regress from these premises and RA. Their diagnosis of the contemplation regress argument (*ibid.* 416) is that this argument is "unsound" and so it "fails to establish any difficulty for the thesis that knowledge-how is a species of knowledge-that."

#### *A Defence of the Contemplation Regress?*

Is there some way of defending the contemplation regress argument against this critique? Noë (2005: 278–82) argues that Stanley and Williamson fail to show that the contemplation regress argument is unsound. Noë's main criticism of their critique is directed at their claim that the contemplation premise is false if we assume that 'contemplates the proposition that p' in the contemplation premise refers to an intentional action. Noë claims that Stanley and Williamson do not provide an argument for this claim. Presumably, this is because they take the

claim to be intuitively obvious once one has considered examples like Ginet's door case. However, Noë argues that all that the door case establishes is that when we perform actions that exercise our knowledge-that we need not be *consciously aware* of contemplating the relevant proposition. But as Noë points out, this conclusion is at least consistent with the possibility that we always contemplate the relevant proposition when we exercise our knowledge-that, and that we do so *intentionally*:

Ryle can accommodate Ginet's observation by countenancing the possibility that not every act of contemplating a proposition is performed consciously. To say that it is or could be performed unconsciously is not to say that it is not the sort of thing that could be performed intentionally (Noë 2005: 282).

Perhaps Noë is right that a proponent of the contemplation regress argument would be best advised to respond to Ginet's door cases by claiming that whenever one employs one's knowledge that *p*, one does intentionally contemplate the relevant proposition, it is just that one need not be consciously aware of performing this intentional action. The problem is that Noë does not provide us with any reason to think that the contemplation premise is true.

The contemplation premise is deeply implausible if we interpret it as claiming that if one employs one's knowledge that *p* then one *consciously* contemplates the proposition that *p*—as Ginet's case clearly establishes. But the contemplation premise is, at best, only marginally less implausible if we interpret it as saying that if one employs one's knowledge that *p* then one *intentionally* (but not necessarily consciously) contemplates the proposition that *p*. Consider my everyday action of opening my office door in the morning. As Ginet points out, it is natural to say that in performing such actions I exercise or employ various kinds of knowledge-that, including my knowledge that one can open my office door by turning the knob and pushing it. But why think that in employing this knowledge-that I must also intentionally perform the action of contemplating the proposition that one can open my office door by turning the knob and pushing it? Noë is right that it is at least *possible* that I perform such an intentional action even though I am not

consciously aware of doing so. But why should we believe in the first place that performing an intentional action of contemplating a proposition is a precondition of employing one's knowledge-that?

If we follow Noë and assume that if one employs one's knowledge that *p* then one *intentionally* (but not necessarily consciously) contemplates the proposition that *p*, then we can derive an infinite and vicious regress from this assumption, the intentional action premise, and RA. But in the absence of some argument for this strange assumption, it is perfectly reasonable for intellectualists to respond to the contemplation regress argument by rejecting this assumption rather than RA. I doubt that any such argument could be given, and so I think we must agree with Stanley and Williamson that the contemplation regress argument fails.

### 4.3 The Employment Regress

We have seen that the contemplation regress argument is not a successful regress argument against intellectualism. Is there a more promising regress argument against intellectualism? In this section I want to identify another form of regress argument against intellectualism that is related to the contemplation regress argument, but which does not rely on the premise that if one employs one's knowledge that *p* then one contemplates the proposition that *p*. Rather, this form of regress argument relies on some premise of the form: if one employs [or applies, or exercises] one's knowledge that *p* then one employs [or applies, or exercises] knowledge how to employ one's knowledge that *p*.

This form of argument is, I believe, often implicit in discussions of Ryle's supposed regress argument against intellectualism. In particular, Hetherington (2006) presents an argument that I think implicitly relies on this kind of premise, and he claims both that it is Ryle's regress argument and that it is a successful regress argument against intellectualism. I will discuss Hetherington's argument separately in §4.5. In this section I will construct what I take to be the clearest statement of this general form of argument, and then identify the connections between this argument and Ryle's critique of the legend. However, before

introducing this argument it will be useful to first make note of a certain fact about employments of knowledge-how.

*Direct knowledge-how*

Consider the action premise again:

The Action Premise

If one Fs, one employs knowledge how to F.

Unlike the contemplation premise, one might think that the action premise is at least *prima facie* plausible. This is suggested by the fact that both proponents of intellectualism—like Stanley and Williamson—and opponents of intellectualism—like Ryle and Noë—have endorsed certain restricted versions of this claim. In fact, for Stanley and Williamson (2001: 443), the plausibility of the intentional action premise explains why philosophers have (albeit mistakenly in their view) found Ryle’s supposed regress argument against intellectualism so plausible: “the thesis that intentional actions are in fact employments of knowledge-how is precisely what accounts for the initial plausibility of Ryle’s original argument against the claim that knowledge-how is a species of knowledge-that.”

An interesting point about the action premise—that is not noted by any of its aforementioned proponents—is that the unrestricted version of the action premise by itself generates an infinite regress. For suppose that Hannah performs some action  $F_1$ . By the action premise, Hannah employs knowledge how to  $F_1$ . But employing one’s knowledge how to  $F_1$  is a legitimate value for ‘F’ in the *unrestricted* action premise because, at least grammatically speaking, employing one’s knowledge-how is something that one *does*. But then, given that employing knowledge-how is something that we do, we can reapply the action premise to conclude that Hannah employs knowledge how to employ her knowledge how to  $F_1$ . But then by the action premise, Hannah also employs knowledge how to employ her knowledge how to  $F_1$ . And so on *ad infinitum*.

The unrestricted action premise can be used then to generate an infinite regress of employments of knowledge-how. To avoid an infinite and vicious regress one must allow that sometimes we can employ our knowledge-how *directly*, in the sense that sometimes we employ our knowledge how to F without also employing some distinct state of knowledge how to employ our knowledge how to F. In other words, the following claim must be false:

If one employs knowledge how to F, one employs knowledge how to employ one's knowledge how to F (and one's state of knowledge how to F and one's state of knowledge how to employ one's knowledge how to F).

To deny this claim is to commit oneself to the conclusion that whenever we employ our knowledge how to F we also employ an infinite number of further and distinct states of knowledge-how. This conclusion is surely absurd, for given that we are finite beings, it tells us that we never employ our knowledge-how.<sup>57</sup>

---

<sup>57</sup> Note that there are two ways one might deny this claim. One way would be to deny the claim that if we can employ our knowledge how to F without employing knowledge how to employ our knowledge how to F. This response would block the regress of employments of knowledge-how. Alternatively, one could accept the regress of employments of knowledge-how but claim that it is benign, rather than vicious, by denying the bracketed clause that states that these two states of knowledge-how are distinct. The idea then would be that when we employ our knowledge how to F we always employ knowledge how to employ our knowledge how to F, but these two states of knowledge-how need not be distinct. The first proposal seems by far the more natural to my mind, and I suspect that the second proposal just collapses into the first. But the important point is simply that we must deny this claim if we are to avoid saying that an infinite and vicious regress ensues whenever anyone employs their knowledge how to do something.

*The employment regress argument*

Some employments of knowledge-how must be direct. I take it that once pointed out, this claim is obvious, even trivial. The employment regress argument relies on the idea that employments of knowledge-that—unlike employments of knowledge-how—cannot be direct. That is, the employment regress argument relies on the following premise that I will refer to as *the employment premise*:

The Employment Premise

If one employs knowledge that p, one employs knowledge how to employ one's knowledge that p (and one's state of knowledge that p is distinct from one's state of knowing how to employ one's knowledge that p).

The contemplation premise tells us that if one employs one's knowledge that p then one contemplates the proposition that p. The employment premise tells us that if one employs one's knowledge that p then one employs knowledge how to employ one's knowledge that p. Furthermore, it says that the relevant state of knowledge that p and the relevant state of knowledge how to employ one's knowledge that p are distinct.

Now, to construct a regress argument here we also need to assume that some intellectualist thesis is true. Let us assume then, for the purposes of *reductio*, that the simple identity thesis is true:

The Simple Identity Thesis

To know how to F is to know that p (for some relevant proposition p concerning a way to F).

Together, the employment premise and the simple identity thesis generate an infinite regress like so: imagine that Ari Fs and in so doing he employs his knowledge how to F. By the identity thesis, it follows that Ari thereby employs his knowledge that p<sub>1</sub> (for some proposition p<sub>1</sub> concerning a way to F). By the employment premise, it follows that Ari also employs knowledge how to employ

his knowledge that  $p_1$ , and that Ari's state of knowing how to employ his knowledge that  $p_1$  is distinct from his state of knowing that  $p_1$ . By the identity premise, it follows that Ari thereby employs his knowledge that  $p_2$  (for some proposition  $p_2$  concerning a way to employ one's knowledge that  $p_1$ ). By the employment premise, it follows that Ari also employs knowledge how to employ his knowledge that  $p_2$ , and that Ari's state of knowing how to employ his knowledge that  $p_2$  is distinct from his state of knowing that  $p_2$ , and so on *ad infinitum*.

We have an infinite regress then of employments of knowledge-that; and every state of knowledge-that in this infinite series is distinct from the state of knowledge-that that immediately precedes it. Such a regress certainly seems vicious. For it is a consequence of this regress that whenever we employ our knowledge how to do something we must also possess and employ an infinite number of further and distinct states of knowledge-that.<sup>58</sup> This conclusion is absurd, for given that we are finite beings what it tells us is that we never employ our knowledge-how. But the inferences required to generate this absurd conclusion are valid, therefore, we must reject either the employment premise or the simple identity thesis. The employment regress argument tells us that we should reject the simple identity thesis, on the basis of the assumption that the employment premise is true.

#### *Challenges to the employment premise*

Is the employment regress argument a successful argument against intellectualism? One might think that it is at least more promising than the contemplation regress argument. For unlike the contemplation regress argument,

---

<sup>58</sup> Strictly speaking, however, the existence of this regress only entails that whenever we employ knowledge how to F we also have to possess and employ an infinite number of distinct states of knowledge-that if we assume that this regress does not 'loop' back on itself (as explained earlier, see fn. 54). I ignore here the possibility that such a regress loops back on itself simply because it strikes me as being highly implausible.

the employment regress argument does not rely on the dubious idea that one employs one's knowledge that  $p$  only if one contemplates the proposition that  $p$ .

However, an intellectualist faced with this argument will no doubt suspect that they can offer good reasons for rejecting the employment premise. For one thing, as we saw earlier, to avoid an infinite and vicious regress it must be the case that some employments of knowledge-how are direct. That is, sometimes we must employ our knowledge how to  $F$  without also employing some further state of knowledge how to employ our knowledge how to  $F$ . But then why should the same not be true of employments of knowledge-that? It is true that, unlike knowledge-how, merely denying that employments of knowledge-that can be direct does not by itself generate an infinite and vicious regress. But still why should we think that one employ one's knowledge that  $p$  only if one employs knowledge how to employ one's knowledge that  $p$ ?

Furthermore, as already indicated, the employment premise is just an instance of the more general claim made by the action premise:

#### The Action Premise

If one  $F$ s, one employs knowledge how to  $F$ .

That is, the employment premise makes the same claim as the action premise except that the range of actions that can be values for ' $F$ ' is restricted to a particular kind of action, namely, employments of knowledge-that. But, as discussed earlier, the unrestricted version of the action premise is subject to clear counterexamples. The intellectualist may well suspect then that we should find analogous counterexamples to the employment premise.

Still, it does not follow from the fact that there are counterexamples to the action premise that there will be counterexamples to the employment premise. And even if there are such counterexamples, it may be that we can offer some revised version of the employment premise that avoids them, and which still generates an infinite and vicious regress together with the simple identity thesis. I will examine such issues in §4.4. But first I will return briefly to Ryle, for I think that within his

critique of the legend we can perceive an attempt to motivate something like the claim made by the employment premise.

*Rylean motivations for the employment premise*

Consider the following passages where Ryle discusses an example modelled on Lewis Carroll's (1895) famous dialogue between Achilles and the Tortoise:

A pupil fails to follow an argument. He understands the premises and he understands the conclusion. But he fails to see that the conclusion follows from the premises. The teacher thinks him rather dull but tries to help. So he tells him that there is an ulterior proposition which he has not considered, namely, that *if these premises are true, the conclusion is true*. The pupil understands this and dutifully recites it alongside the premises, and still fails to see that the conclusion follows from the premises even when accompanied by the assertion that these premises entail this conclusion. So a second hypothetical proposition is added to his store; namely, that the conclusion is true if the premises are true as well as the first hypothetical proposition that if the premises are true the conclusion is true. And still the pupil fails to see. And so on for ever. He accepts rules in theory but this does not force him to apply them in practice. He considers reasons, but he fails to reason. (This is Lewis Carroll's puzzle in 'What the Tortoise said to Achilles'. I have met no successful attempt to solve it.)

What has gone wrong? Just this, that knowing how to reason was assumed to be analysable into the knowledge or supposal of some propositions, namely, (1) the special premises, (2) the conclusion, plus (3) some extra propositions about the implication of the conclusion by the premises, etc., etc., *ad infinitum*. 'Well but surely the intelligent reasoner *is* knowing rules of inference whenever he reasons intelligently.' Yes, of course he is, but knowing such a rule is not a case of knowing an extra fact or truth; it is knowing how to move from acknowledging some facts to acknowledging others. Knowing a rule of inference is not possessing a bit of extra information but being able to perform an intelligent operation. Knowing a rule is knowing how (Ryle 1946: 216-17).

There are obviously numerous different ideas and arguments suggested by Ryle's brief discussion of this example. For one thing, I think Ryle uses this example to support a version of the insufficiency objection to intellectualism that was discussed in Chapter 1; that is, the objection that intellectualism is false because knowledge-that is not sufficient for knowledge-how. He appears to be using the

pupil case to support the claim that someone could know any given proposition while still failing to know how to do something, for example, according to Ryle the pupil could know any given proposition whilst still failing to know how to reason.<sup>59</sup>

However, I think Ryle also infers a more specific insufficiency claim from this example; namely, that knowing that *p* does not suffice for knowing how to employ one's knowledge that *p*. As we saw in §4.1, Ryle (1946: 218) often refers to a 'gap' or 'gulf' "between ... acknowledging principles in thought and intelligently applying them in action." But, at the same time, Ryle (*ibid.*) also refers to a 'gulf' "between having the postulated knowledge of those facts and knowing how to use or apply it". In other words, Ryle often points to a gap between knowing that *p* and knowing how to employ or apply one's knowledge that *p*.

How does the pupil case support this idea that one can know that *p* without knowing how to employ one's knowledge that *p*? If we let *r* be the proposition that *p* and (if *p* then *q*), Ryle appears to take one of the morals of this case to be that one could know that *r* whilst failing to know how to employ one's knowledge that *r* so as to perform some action, like (say) the action of inferring *q* from *r*. In other words, Ryle characterizes this case as one where someone stands in the knowledge-that relation to some proposition(s), but they do not know how to employ or apply this knowledge-that.

Ryle's claim that merely knowing that *p* is not a sufficient condition for knowing how to employ one's knowledge that *p* does not establish the employment premise, but it does at least support the claim made by the bracketed clause in the employment premise:

---

<sup>59</sup> As noted in Chapter 1, there is at least one clear kind of counterexample to the claim that it is possible to stand in the knowledge-that relation to *any* given proposition *p* whilst failing to know how to perform some action *F*; namely, the case where *p* is the proposition that one knows how to *F*, or where *p* is some other proposition the truth of which entails that one knows how to *F*.

### The Employment Premise

If one employs knowledge that  $p$ , one employs knowledge how to employ one's knowledge that  $p$  (and one's state of knowledge that  $p$  is distinct from one's state of knowing how to employ one's knowledge that  $p$ ).

The employment premise and the simple identity thesis together generate an infinite regress of employments of knowledge-that. The bracketed clause in the employment premise ensures that each new state of knowledge-that in this infinite regress will be distinct from its predecessor, and so it ensures that this regress is vicious. Ryle's claim that it is always possible that one know that  $p$  but fail to know how to employ one's knowledge that  $p$ , supports the assumption that this bracketed clause expresses; namely, that knowing that  $p$  and knowing how to employ one's knowledge that  $p$  are always two distinct states of knowledge.

Furthermore, I think Ryle sees the pupil case as supporting a claim that is even more closely related to the employment premise: the claim that if one employs one's knowledge that  $p$  then one knows how to employ one's knowledge that  $p$ . For example, it appears that Ryle infers from this case that if one is to infer  $q$  from one's knowledge that  $r$ , then one must know how to employ one's knowledge that  $r$  so as to reach this conclusion; as the difference between the intelligent reasoner and the pupil for Ryle (1946: 217) is that only the former "[knows] how to move from acknowledging some facts to acknowledging others."

Similarly, on the basis of the stupid chess player example discussed in §4.1, Ryle concludes that "it requires intelligence not only to discover truths, but also to apply them, and knowing how to apply truths cannot, without setting up an infinite process, be reduced to knowledge of some extra bridge truths" (Ryle 1946: 216). Ryle seems to suggest here not only that employing one's knowledge-that is an action that requires intelligence, but also that it is an action the performance of which requires one to *know how* to perform that action. Presumably, for Ryle, the stupid chess player supports this claim because his repeated failures to make a good move in a game of chess are naturally explained

by his failure to know how to apply, or employ, his propositional knowledge of the relevant chess strategies. Ryle will think that without such knowledge-how the stupid chess-player “might not see how to apply” (1946: 216) this propositional knowledge, so as to make a good move in a game of chess.

It seems then that one of the many morals Ryle infers from examples like his pupil and stupid chess player cases is that it is a precondition of employing one’s knowledge that *p* that one know how to employ one’s knowledge that *p*. And from this idea it is a fairly short step to the claim that it is a precondition of employing one’s knowledge that *p* that one *employ* knowledge how to employ one’s knowledge that *p*. Indeed, I think it is reasonable to attribute both of these claims to Ryle, with one important qualification.

The qualification is that Ryle would presumably hold that these principles only apply to *intelligent* employments of knowledge-that. For recall that Ryle does not endorse the unrestricted version of the action premise; rather, he endorses the intelligent action premise:

The Intelligent Action Premise

If one Fs *intelligently* [or cleverly, or carefully etc.], one employs knowledge how to F.

Presumably then, Ryle would only hold that knowing how to employ one’s knowledge that *p* is a precondition of employing one’s knowledge that *p*, when one employs one’s knowledge that *p* *intelligently*, or *cleverly*, or *carefully* etc. For as we saw in §4.1, Ryle frequently stresses that employing one’s knowledge-that—or ‘applying truths’—is itself an action that can be performed more or less intelligently. We can assume then that Ryle would only endorse the following restricted version of the employment premise:

If one employs knowledge that *p* *intelligently*, one employs knowledge how to employ one’s knowledge that *p* (and one’s state of knowledge that *p* is

distinct from one's state of knowing how to employ one's knowledge that p).

Whether this restricted version of the employment premise could be used to generate a plausible regress argument against intellectualism is an issue I will address in the next section. For now it will suffice to simply note Ryle's attempt to motivate a version of the employment premise.

#### 4.4 Intellectualist Responses to the Employment Premise

In this section I consider what reasons an intellectualist might offer for rejecting the employment premise. To begin with, consider Ryle's attempts to motivate the claim made by the employment premise (or something in the neighbourhood of this claim) succeed. Do examples like Ryle's pupil and stupid chess player cases really support the idea that if one employs one's knowledge that p then one employs knowledge how to employ one's knowledge that p? Ryle is surely right in proposing that failing to know how to employ one's knowledge that p will *sometimes* result in one's failure to successfully employ one's knowledge that p. But the employment premise tell us that failing to know how to employ one's knowledge that p will *always* result in one's failing to successfully employ one's knowledge that p.

Suppose that all of the cases Ryle describes are in fact cases where someone fails to employ their knowledge that p because they do not know how to employ it. Even given this assumption, the existence of such individual cases by itself does not demonstrate that knowing how to employ one's knowledge that p is a necessary condition of employing one's knowledge that p, nor does it demonstrate that *employing* knowledge how to employ one's knowledge that p is a necessary condition of employing one's knowledge that p.

Furthermore, as I suggested earlier, an intellectualist will suspect that there will be counterexamples to the employment premise that are analogous to counterexamples we can give to the action premise. Before I consider whether

there are such analogous counterexamples, note that *if* there are such counterexamples a proponent of the employment regress argument could not respond to them by merely restricting the employment premise to *intentional* employments of knowledge-that, like so:

If one employs knowledge that p *intentionally*, one employs knowledge how to employ one's knowledge that p.

Even if this restricted version of the employment premise is correct, one cannot validly derive an infinite regress from this claim and the assumption that the simple identity thesis is true. Rather, for the derivation to be valid, one would have to instead endorse the following claim:

If one employs knowledge that p *intentionally*, one employs knowledge how to employ one's knowledge that p *intentionally*.

But this claim looks deeply implausible. Surely we can intentionally employ our knowledge-that without *intentionally* employing knowledge how to employ this knowledge-that. So, if there are counterexamples to the employment premise one could not plausibly respond to them by stipulating that the claim it makes only applies to intentional employments of knowledge-that. And if all intelligent actions are intentional actions then the same point applies to Ryle's version of the employment premise that is restricted to intelligent employments of knowledge-that.

Of course, we have not yet established that there are clear counterexamples to the employment premise. Indeed, I think this is actually quite a difficult endeavour. For one thing, it seems difficult to imagine a counterexample to the employment premise that parallels the case where Hannah wins the lottery, but she does not know how to win the lottery as she only wins the lottery due to sheer chance or

mere luck.<sup>60</sup> But suppose that we could describe a case where someone employs their knowledge that *p* but only by sheer chance or mere luck, and so they do not thereby employ knowledge how to employ their knowledge that *p*. Even if there was such a case, I think it would be reasonable for a proponent of the employment regress argument to respond to it by simply replacing the employment premise with something like the following claim:

If one employs one's knowledge that *p non-accidentally*, one employs knowledge how to employ one's knowledge that *p non-accidentally*.

---

<sup>60</sup> One might think that the following is an example of such a case: Suppose that Mary is a contestant in a TV game show. The first question is 'What is the capital of New Zealand?' Mary knows the answer, but she was not listening when she was told the rules of this game and so she is not sure what she is meant to do when she knows the answer to a question. Mary knows that she has to shout the answer out loud a specific number of times depending on how much prize money is at stake at that point in the game, but she has no idea how one is meant to calculate that number. Given her predicament, Mary shouts out 'Wellington!' nine times, simply because nine is her lucky number. Luckily for Mary, given the prize money on offer at that point in the game, the right thing to do to win the prize was to shout the correct answer exactly nine times, and so Mary wins the prize. One might think that this is a case where someone employs their knowledge that *p* but does not employ knowledge how to employ their knowledge that *p*, given the luck involved in Mary's winning the prize. However, I think this is a mistaken diagnosis. The right diagnosis is that it is a case where someone employs their knowledge that *p so as to F* without employing knowledge how to employ their knowledge that *p so as to F*, and such a case is not a counterexample to the employment premise. It is true that Mary did not employ knowledge how to employ her knowledge that Wellington is the capital of New Zealand so as to win the prize, for this is something Mary did not know how to do. But Mary did know how to employ her knowledge that Wellington is the capital of New Zealand so as to shout out the correct answer nine times, and she did employ this knowledge-how in shouting out 'Wellington!' nine times.

As this claim does not obviously commit one to the implausible claim that these non-accidental employments of knowledge-that and knowledge-how must be intentionally performed.

Perhaps we can find telling counterexamples to the employment premise by looking for analogues of the digestion and sweating counterexamples to the action premise, rather than the lottery counterexample? The only kind of cases that I can imagine that might be roughly analogous to such examples are employments of the kind of *tacit* knowledge appealed to in cognitive psychology. For example, suppose that our best theory of the linguistic competence of native English speakers attributes to them the tacit knowledge that ' $NP \rightarrow Det + Adj + N$ ' is a rule of English; such that sometimes when an English speaker exercises that competence (by producing grammatical utterances or detecting ungrammatical sentences etc.) they do so (in part) by employing their tacit knowledge that ' $NP \rightarrow Det + Adj + N$ ' is a rule of English.

Now imagine that Mary, a native English speaker, exercises her linguistic competence in some way and, in so doing, she employs her tacit knowledge that ' $NP \rightarrow Det + Adj + N$ ' is a rule of English. Would Mary thereby also employ knowledge how to employ her knowledge that ' $NP \rightarrow Det + Adj + N$ ' is rule of English? Arguably not, as it seems odd to say that Mary knows how to employ her knowledge that ' $NP \rightarrow Det + Adj + N$ ' is a rule of English. Employing her tacit knowledge that ' $NP \rightarrow Det + Adj + N$ ' is a rule of English, one might think, is something that Mary does but it is not something she knows how to do. In which case, when Mary employs her tacit knowledge that ' $NP \rightarrow Det + Adj + N$ ' is rule of English, she does not also employ knowledge how to employ this knowledge-that. This conclusion that Mary employs this tacit knowledge *directly* seems particularly plausible given that the content of such states of tacit knowledge is often thought to be inaccessible to conscious reflection and inferentially isolated from our belief forming mechanisms etc.

Insofar as we can be said to possess and employ such states of tacit knowledge, it seems to me that we must employ them directly. But, nonetheless, I think there is

an inherent problem with this strategy of appealing to employments of tacit knowledge as a way of trying to identify counterexamples to the employment premise. The problem is that even if we assume that we do need to posit such states of tacit knowledge to explain (say) our linguistic capacities, it is not at all clear that such states should really be thought of as being genuine states of propositional knowledge, rather than some other kind of intentional or informational state. For example, propositional knowledge requires justification or warrant, but such notions do not even seem to be applicable to the states of ‘tacit knowledge’ or ‘implicit knowledge’ appealed to in psychology. As Martin Davies (2001: 8127) writes: “the notion of justification does not seem to be applicable in cases where the subject is unaware of the presence or influence of the information.” Furthermore, Gareth Evans (1981) and Steven Stich (1978, 1980), amongst others, have argued that the states of ‘tacit knowledge’ are not even genuine belief states, let alone states of knowledge-that.

The intellectualist could always try to argue that the states of tacit knowledge appealed to in the cognitive sciences really are genuine states of propositional knowledge. But given how controversial this issue is, appealing to such states of tacit knowledge does not look to be a promising way of identifying clear counterexamples to the employment premise. Rather, the intellectualist needs to appeal to a case where it is clear that someone employs a genuine state of knowledge-that, and not merely some other kind of intentional or informational state.

Perhaps a better kind of case for the intellectualist to appeal to in the attempt to find a compelling counterexample to the employment premise would again be Ginet’s door case:

I exercise (or manifest) my knowledge *that* one can get the door open by turning the knob and pushing it (as well as my knowledge *that* there is a door there) by performing that operation quite automatically as I leave the room; and I may do this, of course, without formulating (in my mind or out loud) that proposition or any other relevant proposition (Ginet 1975: 7).

This is a case where someone clearly employs or exercises a genuine state of knowledge-that. And it is also a case where the employing of knowledge-that is not performed either consciously or intentionally. The intellectualist might argue then that because Ginet's 'automatic' employment of his knowledge (that one can get the door open by turning the knob and pushing it) is not an intentional action, then it is an action that he performs without employing knowledge how to perform it.

Is such a diagnosis of Ginet's door case correct? That is, is this case an example of someone employing their knowledge that *p* without employing knowledge how to employ their knowledge that *p*? What *is* clear is that when I open my office door in the morning I do not consciously or intentionally employ knowledge how to employ my knowledge that one can open my office door by turning the knob and pushing it. But nor do I consciously or intentionally employ my knowledge that one can open my office door by turning the knob and pushing it, yet it still seems correct to say that, in some sense, I employ or exercise this knowledge-that when I open my office door—as Ginet's example illustrates. A proponent of the employment regress argument might claim then, analogously, that in opening the door I do employ my knowledge how to employ my knowledge that one can open the door by turning the knob and pushing it, even though I do not do so either consciously or intentionally.

It seems to me that it is not entirely clear what the correct thing is to say about such examples. I can feel some pull towards the claim that I not only know that one can open my office door by turning the knob and pushing it; but that I also know how to employ my knowledge that one can open my office door by turning the knob and pushing it; and that I employ or exercise this knowledge-how when I open the door. But, on the other hand, it seems to me that it is far clearer that the knowledge-that attribution is correct than that the knowledge-how attribution is.

Of course, if I employ my knowledge that *p* then I *can* employ my knowledge that *p*, from which it follows that, in some sense, I have the *ability* to employ my

knowledge that *p* (at least if my employing of my knowledge that *p* was not a mere accident or fluke). Now, from this initial conclusion a neo-Rylean would infer that I thereby know how to employ my knowledge that *p*, because they hold that possessing the ability entails knowing how to *F*. In which case, having the ability to employ my knowledge that one can open the door by turning the knob and pushing it, entails that I know how to employ my knowledge that one can open the door by turning the knob and pushing it. Furthermore, neo-Ryleans claim that to know how to *F* just is to possess the ability to *F*. And if this identity claim is true, to employ an ability to employ my knowledge that *p* just is to employ knowledge how to employ my knowledge that *p*. In other words, if one accepts neo-Ryleanism it is a fairly easy matter to motivate the employment premise.

Indeed, I think that it is likely that Ryle himself is implicitly assuming the truth of neo-Ryleanism when he appeals to his pupil and stupid chess player cases in supporting the claim that if one employs one's knowledge that *p* then one knows how to employ one's knowledge that *p*. For recall that in his discussion of the pupil case Ryle moves freely from the claim that "[k]nowing a rule of inference is not possessing a bit of extra information but being *able to* perform an intelligent operation" to the claim that "Knowing a rule is *knowing how*" (Ryle 1946: 217; emphasis added). This suggests that for Ryle the claim that if one employs one's knowledge that *p* then one employs *knowledge how* to employ one's knowledge that *p* is simply equivalent to, or at least entails, the claim that if one employs one's knowledge that *p* then one employs one's *ability* to employ one's knowledge that *p*.

But the problem here is that intellectualists offer strong reasons for denying the idea that possessing the ability to *F* is either identical with or entails knowing how to *F*, as we saw in Chapter 1. In which case, the intellectualist will feel that they can happily agree, on the one hand, that Ryle's examples show us that if one employs one's knowledge that *p* then one employs an *ability* to employ one's knowledge that *p*, whilst denying, on the other hand, that if one employs one's

knowledge that p then one employs *knowledge how* to employ one's knowledge that p.

As I have illustrated in Chapters 2 and 3, Stanley and Williamson hold that if one has the ability to F *intentionally* then one knows how to F. But this claim is consistent with the point being made here. For one thing, whilst they accept this entailment, they do not think that knowing how to F entails having the ability to F intentionally (because of examples like the ski instructor case). So, they are clearly not committed to identifying knowing to F with possessing the ability to F intentionally. Furthermore, such an identity claim could only be used to help motivate a version of the employment regress argument that relied on the claim that if one employs one's knowledge that p *intentionally* then one employs knowledge how to employ one's knowledge that p *intentionally*. And this claim is clearly false.

The moral here is that the intellectualist will not want to contest the obvious truth that to employ one's knowledge that p one has to, in some sense, possess the ability or capacity to do so. What they *will* deny is that this ability must be identical to some state of knowledge-how.

Such considerations suggest that the intellectualist can offer good general reasons for rejecting the employment premise, even if it is difficult to produce a really decisive counterexample to this claim. For the intellectualist can argue that the counterexamples we find to the action premise at least suggest that when one non-intentionally employs one's knowledge that p, one need not also employ knowledge how to employ their knowledge that p. And they could also point out that even if our intuitions about examples like Ginet's door case are not entirely clear either way, this at least shows us that it is not obviously true that any case of someone employing their knowledge that p will also be a case where they employ knowledge how to employ their knowledge that p.

Furthermore, insofar as we are inclined to think that the employment premise is correct, the intellectualist could argue that we are being misled by the fact that grammatically speaking, 'employing', 'applying', or 'exercising', our knowledge-

that are things that we do. As discussed in §4.2, digesting and sweating are also things that, at least in a grammatical sense, we do. But it would obviously be a mistake to infer from this grammatical fact that such ‘actions’ are the kind of actions that we perform as agents, or that we know how to perform. Similarly, one might argue that at least some employments of knowledge—that are simply not the kind of ‘actions’ one knows how to perform. In opening my office door I employ my knowledge that one can open the door by turning the knob and pushing it, and my employing of this knowledge—that is, at least grammatically speaking, something that I *do*. But is this action of employing my knowledge really of a kind with my action of opening the door? That is, is this employment of knowledge—that also an action that I perform as an agent, and that I know how to perform? At the very least, I think an intellectualist could make a strong case for thinking that it is not.

Where does this leave our assessment of the employment regress argument? I think we must conclude that—at the very best—the employment regress argument presents an inconclusive case against the simple identity thesis. As we saw earlier, it must be possible to employ one’s knowledge-how directly, for to assume otherwise would lead to an infinite and vicious regress. The employment regress argument relies on the idea that unlike employments of knowledge-how, employments of knowledge-that cannot be direct. But it is not at all clear that the intellectualist must accept this asymmetry. And while it is difficult to describe a decisive counterexample to the employment premise, the intellectualist can offer a number of strong considerations for thinking that the principle it expresses is nonetheless false.

#### **4.5 Related Regress Arguments**

In the previous two sections, we have seen that there are good reasons for thinking that both the employment and contemplation regress arguments are unsound. In this section I will show how similar problems arise for the regress arguments presented by Hetherington (2006) (whose argument is closely related to the employment regress argument) and Noë (2005).

*Hetherington's regress argument*

Hetherington presents his regress argument against the intellectualist view of knowledge-how as a *reductio* of the following claim he calls simply R:

R For any action F, and for some content  $\Phi$  describing a sufficient criterion of how to do F: If (when doing F) one knows how to F; then (1) one already has knowledge that  $\Phi(F)$ , which (2) one knows how to, and one does, apply so as to do F (Hetherington 2006: 73).

Now, Hetherington's talk here of 'a content  $\Phi$  describing a sufficient criterion of how to do F' is somewhat opaque, but I take it that his idea is that  $\Phi$  is a proposition concerning something like a way to F. If we leave the quantifiers in Hetherington's original statement of R implicit, his regress argument then is meant to be a *reductio* of the following claim:

(R) If (when doing F) one knows how to F, then one already has knowledge that p (for some proposition p concerning a way to F), which one knows how to, and one does, apply so as to do F.

An infinite regress can be generated from R like so. Suppose Ari Fs and he knows how to F. From R three things follow: Ari knows that  $p_1$  (for some proposition  $p_1$  concerning a way to F); he applies his knowledge that  $p_1$ ; and he knows how to apply his knowledge that  $p_1$ . But applying his knowledge that  $p_1$  is something that Ari does, and it is something that Ari knows how to do, in which case we can reapply R to conclude that: Ari knows that  $p_2$  (for some proposition  $p_2$  concerning a way to apply one's knowledge that  $p_1$ ); he applies his knowledge that  $p_2$ ; and he knows how to apply his knowledge that  $p_2$ . But, again, applying his knowledge that  $p_2$  is something that Ari does, and it is something that he knows how to do, so we can reapply R, and so on *ad infinitum*.

Hetherington (2006: 73–4) clearly thinks that the infinite regress that follows from R is vicious, and that therefore we must reject R. This is clear in his own

description of how his “Rylean anti-intellectualist argument” against R is meant to proceed:

If one knows how to F, then one does F only if (for some content  $\Phi$  describing a sufficient criterion of how to do F): one already has knowledge that  $\Phi(F)$ , which one knows how to, and one does, apply so as to do F. But if one already knows how to—and one does—apply one’s knowledge that  $\Phi(F)$  so as to do F, then this is a *fresh* instance of both performing and knowing how to perform a specific action. At which point, R is again applicable; and so the foregoing form of reasoning recurs. We thereby begin a regress (an infinite vicious one) of more and more instances of increasingly complex regulative knowledge-that being needed and applications of them being performed—all of this, before one can perform even one action which manifests knowledge-how. Given R, therefore, we are unable to perform even one such action in the first place. Yet we *can* do so. Hence, R is false (*ibid.* 73–4).

Evidently Hetherington thinks that the regress that follows from R will be vicious because he assumes that each new action of applying knowledge-that in this infinite regress must be distinct from its predecessor in the series, and that each new state of knowledge-that must be more complex than, and distinct from, its predecessor. I think that as R is formulated, it is not obvious that these assumptions about the regress that follows from it are correct. But, for the sake of argument, let us grant that the regress generated by R is vicious, and that therefore we must reject it.

What is the import of this conclusion? Hetherington (*ibid.* 74) claims that in establishing that R is false his argument establishes that ‘intellectualism’ is false on the grounds that “R is intellectualism-as-applied-to-our-intelligently-performed-actions, which is to say that it *is* intellectualism.” Now, Hetherington’s talk of ‘intellectualism-as-applied-to-our-intelligently-performed-actions’ is difficult to interpret. But the important point is that Hetherington (*ibid.* 74) clearly holds that his regress argument establishes that the intellectualist view of knowledge-how is false, or as he says, that: “Knowledge how is not simply, or even complicatedly, knowledge-that.” And Hetherington (*ibid.* 74) thinks that it is a virtue of his regress argument that, unlike the contemplation regress, the reasons Stanley and

Williamson offer for rejecting the action premise and the contemplation premise do not apply to R. For digestion is not the kind of thing that one knows how to do, and so one never digests one's food whilst knowing how to digest one's food. And R does not include the dubious claim that if one employs or applies one's knowledge that p, then one contemplates the proposition that p.

To properly assess what Hetherington's argument establishes it will help to note that the claim made by R is of the form: if p then q, r and s. In which case, there are actually three conditional claims expressed by R of the form: if p then q; if p then q and r; and if p then q, r and s. To see more clearly what follows if R is false, let us separate out these three claims expressed by R as R.1, R.2 and R.3:

(R.1) If (when doing F) one knows how to F, then one already has knowledge that p (for some proposition p concerning a way to F).

(R.2) If (when doing F) one knows how to F, then one already has knowledge that p (for some proposition p concerning a way to F) which one knows how to apply so as to do F.

(R.3) If (when doing F) one knows how to F, then one already has knowledge that p (for some proposition p concerning a way to F) which one knows how to, and one does, apply so as to do F.

Now that we have distinguished these three claims we can see that an intellectualist can happily grant Hetherington's claim that his argument shows us that R is false. Any intellectualist must endorse R.1, for if any form of intellectualism is true it follows that if one knows how to F then there is some proposition p concerning a way to F such that one knows that p. But while an intellectualist must endorse R.1, they can still consistently deny R.2 and/or R.3, for neither of these claims is entailed by intellectualism. Hence, Hetherington's claim that R "is intellectualism" is simply false if 'intellectualism' is used to refer to the view of knowledge-how called 'intellectualism'. And in response to Hetherington's regress argument, an intellectualist can accommodate the

conclusion that R is false by denying R.2 and/or R.3, as to do so would be consistent with their commitment to R.1. So, if Hetherington's regress argument against intellectualism is to succeed it must be that R.2 and R.3 are independently plausible, and not that they are entailed by the intellectualist view of knowledge-how, as Hetherington himself appears to suggest.

An intellectualist faced with Hetherington's argument can easily justify rejecting R.3. To see why, consider an intellectualist who accepts the simple identity thesis, that is, an intellectualist who thinks that to know how to F is to know that p (for some proposition p concerning a way to F). According to such an intellectualist, R.3 is equivalent to the following claim:

If (when doing F) one knows how to F then one already knows how to F, one knows how to apply one's knowledge how to F, and one does *apply one's knowledge how to F* so as to do F.

To reuse an earlier example, I can know how to knock the vase off the mantelpiece but, when I do so accidentally, I do not apply this knowledge-how. What this shows us is that it is a simple task to describe cases where someone Fs and knows how to F, but does not apply their knowledge how to F when they F. In which case, an intellectualist faced with Hetherington's regress argument can easily justify rejecting R.3 whilst still endorsing R.1.

Rather than the idea that one merely Fs and one also knows how to F, I think what this problem reveals is that when Hetherington talks of '(when doing F) one knows how to F' the real idea he is aiming at is something like the idea that one Fs and *in so doing one applies one's knowledge how to F*. To focus on the deeper issues facing Hetherington's argument then, I suggest that we replace his R with something like R\*:

(R\*) If one Fs and in so doing one applies knowledge how to F, then one already has knowledge that p (for some proposition p concerning a way to

F), one applies one's knowledge that  $p$  and in so doing one applies knowledge how to apply one's knowledge that  $p$ .

R tells us that if one Fs and one knows how to F then: one knows that  $p$  (for some proposition  $p$ ); one knows how to apply one's knowledge that  $p$ ; and one does apply one's knowledge that  $p$  so as to F. The intellectualist can obviously reject this claim because one can clearly know how to perform some action and yet not apply that knowledge-how when one performs that action, as the vase example demonstrates. What  $R^*$  tells us, on the other hand, is that if one Fs and *in so doing one applies one's knowledge how to F* then: one knows that  $p$  (for some proposition  $p$ ); one applies one's knowledge that  $p$ ; and in so doing one applies knowledge how to apply one's knowledge that  $p$ . The vase example is not a problem for this claim because one does not apply one's knowledge how to knock the vase off the mantelpiece when one accidentally knocks it off the mantelpiece.

To see how we can generate an infinite regress from  $R^*$  suppose that Ari Fs and in so doing he applies his knowledge how to F. From  $R^*$  it follows that: Ari knows that  $p_1$  (for some proposition  $p_1$  concerning a way to F); he applies his knowledge that  $p_1$ ; and in so doing he applies knowledge how to apply his knowledge that  $p_1$ . But then we can reapply  $R^*$  to conclude that: Ari knows that  $p_2$  (for some proposition  $p_2$  concerning a way to apply one's knowledge that knowledge that  $p_1$ ); he applies his knowledge that  $p_2$ ; and in so doing he applies his knowledge how to apply his knowledge that  $p_2$ ; and so on *ad infinitum*. And, as with R, let us grant for the sake of argument that the infinite regress that follows from  $R^*$  is vicious, and therefore we must reject  $R^*$ .

I think this regress argument against  $R^*$  is essentially the one Hetherington has in mind when he presents his own regress argument against R, so from now on I will simply refer to it as Hetherington's argument. But again, as with R, to properly assess the import of this argument we need to distinguish the three conditionals expressed by  $R^*$ :

(R.1\*) If one Fs and in so doing one applies knowledge how to F, then one already has knowledge that p (for some proposition p concerning a way to F).

(R.2\*) If one Fs and in so doing one applies knowledge how to F, then one already has knowledge that p (for some proposition p concerning a way to F), and one applies one's knowledge that p.

(R.3\*) If one Fs and in so doing one applies knowledge how to F, then one already has knowledge that p (for some proposition p concerning a way to F), one applies one's knowledge that p, and in so doing one applies knowledge how to apply one's knowledge that p.

Now, as with R.1, any intellectualist must accept R.1\*, for if any form of intellectualism is true it follows that if one Fs, and in so doing one applies knowledge how to F, then one already has knowledge that p (for some proposition concerning a way to F).

What about R.2\*? An intellectualist committed to some instance of the simple identity thesis must accept R.2\*. For if to know how to F is to know that p (for some proposition p concerning a way to F) then it follows that if one applies one's knowledge how to F then one applies one's knowledge that p (for some proposition p concerning a way to F). But not all intellectualists need accept the simple identity thesis, so it is not clear that any intellectualist need accept R.2\*. This is a serious concern with this argument, but similar concerns arise for the contemplation and employment regress arguments, so I will discuss this general issue separately in §4.6. For now, let us simply consider whether Hetherington's argument can, at the very least, show us that the simple identity thesis is false. Given that an intellectualist who endorses the simple identity thesis must accept both R.1\* and R.2\*, the issue is whether such an intellectualist can deny R.3\* whilst still maintaining that R.1\* and R.2\* are true.

The first thing to note is that any intellectualist can consistently deny R.3\* whilst maintaining that R.1\* and R.2\* are true, for no form of intellectualism entails the conditional stated by R.3\*. Furthermore, note that an intellectualist who is committed to the simple identity thesis will regard R.3\* as being equivalent to the following claim:

If one Fs and in so doing one applies knowledge how to F, then one already has knowledge how to F, one applies one's knowledge how to F, and in so doing one applies knowledge how to apply one's knowledge how to F.

This claim is obviously false, for it tells us that applications of knowledge-how can never be *direct*; that is, that one can never apply one's knowledge how to F without also applying some distinct state of knowing how to apply one's knowledge how to F. And, as we saw in §4.3, the assumption that applications or employments of knowledge-how cannot be direct leads to an infinite and vicious regress.

I take it that the implicit reason that Hetherington thinks that *an intellectualist* could not justifiably reject R.3\* is that he assumes that something like the employment premise is correct. In particular, I think that it is clear that in giving his regress argument Hetherington is implicitly appealing to something like the following claim:

The Application Premise

If one applies one's knowledge that p (for some proposition p concerning a way to F) then one applies knowledge how to apply one's knowledge that p so as to F

The basic idea behind Hetherington's argument appears to be that if we assume that to know how to F is to know that p (for some proposition p concerning a way to F) and that the application premise is correct, then it follows that: if one Fs, and in so doing one applies one's knowledge how to F, then one knows that p

(for some proposition  $p$  concerning a way to  $F$ ); one applies one's knowledge that  $p$ ; and in so doing one also applies knowledge how to apply one's knowledge that  $p$ . But together these two assumptions generate an infinite and vicious regress, and therefore one of them must be false. Because Hetherington holds that the application premise is true, he thereby concludes that the assumption that to know how to  $F$  is to know that  $p$  (for some proposition  $p$  concerning a way to  $F$ ) is false. Hetherington appears to assume that in establishing this conclusion he thereby establishes that intellectualism is false, but as mentioned above, this conclusion would only clearly establish that the simple identity thesis is false.

This is, I think, is the best interpretation of how Hetherington's regress argument is meant to work. But, as we have seen, intellectualists can offer good reasons for rejecting the employment premise, and these same reasons could obviously be redeployed against the application premise. As with the employment regress argument then, I think we must conclude that Hetherington's argument, at best, only presents an inconclusive case against the simple identity thesis.

#### *Noë's possession regress argument*

All of the regress arguments against intellectualism that I have considered so far have been arguments that claim that if intellectualism is true then an infinite and vicious regress ensues whenever we *employ* or *apply* our knowledge-how. Noë (2005), however, has sketched a different kind of regress argument against intellectualism. Roughly, Noë claims that if intellectualism is true then an infinite and vicious regress ensues whenever we *possess* knowledge-how. He states his regress argument like so:

[g]rasping propositions itself depends on know-how; but if know-how consists in the grasp of further propositions, then one might wonder whether one could ever grasp a proposition. One way this argument might be fleshed out is in terms of concepts: to grasp a proposition, you need to understand the concepts deployed in it; to understand some concepts may be to grasp propositions; but this can't be true for all concepts, on pain of infinite regress. At some point, therefore, it must be possible to give possession-conditions for concepts in non-conceptual, and so non-propositional terms. For

example, my grasp on the concept *red* probably does not consist in my knowledge of propositions about redness. Indeed, one can reasonably wonder whether there could be such propositions. My grasp of *red* consists, it is more likely, in my disposition to apply *red* to an object when it exhibits a certain quality (Peacocke 1992). This regress argument remains unanswered. (Noë: 285–286)

I think the simplest way of representing the regress argument that Noë has in mind is that it is one in which an infinite regress is generated from the following three premises:

(N.1) If one knows that *p* then one possesses the ability to *F* (for some action *F*).

(N.2) To possess the ability to *F* is to know how to *F*.

(N.3) To know how to *F* is to know that *p* (for some proposition *p* concerning a way to *F*).

To see how these premises generate a regress, suppose that Hannah knows that *p*<sub>1</sub>, for some proposition *p*<sub>1</sub>. By N.1 it follows that Hannah possesses the ability to *F*<sub>1</sub>, for some action *F*<sub>1</sub>. By N.2 it follows that Hannah's ability to *F*<sub>1</sub> consists in her knowing how to *F*<sub>1</sub>. By N.3 it follows that Hannah's knowing how to *F*<sub>1</sub> consists in her knowing that *p*<sub>2</sub>, for some proposition *p*<sub>2</sub>. But then we can reapply N.1 to conclude that Hannah possesses the ability to *F*<sub>2</sub>, for some action *F*<sub>2</sub>. By N.2 it follows that Hannah's ability to *F*<sub>2</sub> consists in her knowing how to *F*<sub>2</sub>. By N.3 Hannah's knowing how to *F*<sub>2</sub> consists in her knowing that *p*<sub>3</sub>, for some proposition *p*<sub>3</sub>. But then we can reapply N.1 to conclude that Hannah possesses the ability to *F*<sub>3</sub>, for some action *F*<sub>3</sub>, and so on *ad infinitum*.

Together, N.1, N.2 and N.3 generate an infinite regress of states of knowledge—that. It is not clear that this regress must be a vicious, but let us grant for the sake of argument that it is. What is the import of this conclusion? Noë holds that the right response to this regress is to reject N.3, which is just the simple identity thesis. How might an intellectualist who is committed to the simple identity thesis

respond to this argument? Presumably, they would not deny N.1, as Noë is presumably correct in thinking that it is a precondition of knowing that  $p$  (for any proposition  $p$ ), that one possesses certain abilities or dispositions. But, of course, an intellectualist will deny N.2, as N.2 is just a statement of the neo-Rylean idea that to possess the ability to  $F$  is to know how to  $F$ . And, as we have seen, intellectualists argue that neo-Ryleanism is false on the grounds that possessing the ability to  $F$  is neither a necessary nor a sufficient condition for knowing how to  $F$ .

Now, Noë does acknowledge that intellectualists like Stanley and Williamson will reject N.2 but he does not believe that they would be justified in doing so:

Stanley and Williamson can perhaps evade these difficulties if they can show that having the ability to do something does not consist in knowing how to do it (for then they could admit that grasping propositions depends on basic practical abilities without admitting that it thereby depends on *knowledge-that*). As we have seen, they do not give us reason to follow them in making this separation. If, as I remain convinced, the possession of abilities is a matter of knowledge-how, then we are led to consider the possibility that the truth is exactly the opposite of what Stanley and Williamson maintain: All knowledge-that depends on and must be analysed in terms of a more basic knowledge-how. Intellectualism over-intellectualizes the mind (*ibid.* 286).

I do not wish to examine Noë's criticisms of the reasons Stanley and Williamson give for rejecting neo-Ryleanism. For one thing, I do not think they are very convincing. But, more importantly, the fact that Noë's regress argument relies upon the truth of neo-Ryleanism reveals that his regress argument against intellectualism is, in a sense, redundant. For it is not merely the case that, as a matter of fact, most proponents of intellectualism reject neo-Ryleanism. Rather, intellectualists *must* reject neo-Ryleanism, for these two views of knowledge-how are plausibly contraries, as I noted in Chapter 1. An easy way of illustrating that these views are contraries is to note that it is surely possible that one possess the ability to  $F$  even though there is no way  $w$  such that one knows that  $w$  is a way to  $F$ , or that  $w$  is a way for oneself to  $F$ , etc (at the very least, this is surely true in cases where this ability is *not* an ability to  $F$  *intentionally*). But then the truth of neo-

Ryleanism entails that any form of either simple or sophisticated intellectualism must be false. In which case, the key premise that Noë's regress argument relies upon is a premise the truth of which entails that intellectualism is false. The real issue raised here then is whether or not neo-Ryleanism is true, and as we saw in Chapter 1, intellectualists have offered examples that strongly suggest that possessing the ability to F is neither necessary nor sufficient for knowing how to F. In which case, neo-Ryleanism is false and Noë's argument is unsound.

#### 4.6 Assessing the Regress Objection

I have now examined four regress arguments against intellectualism, all of which can be seen as drawing inspiration from Ryle's critique of the intellectualist legend, to varying degrees. My examination of these arguments suggests that none of them is sound. But even if they are sound, there is another kind of problem with them, which I have alluded to but not as yet examined. Namely, the problem that even if these arguments are sound it is not clear that they could establish the conclusion that all forms of intellectualism are false. To introduce this problem, note that the contemplation regress, the employment regress, and Noë's regress argument are all arguments against some version of the simple identity thesis:

##### The Simple Identity Thesis

To know how to F is to know that p (for some relevant proposition p concerning a way to F).

All three arguments assume that the simple identity thesis is true and then attempt to show that this assumption leads to an infinite and vicious regress. But, as mentioned in passing in Chapter 1, it is actually not clear that even a simple intellectualist need accept the simple identity thesis. Recall that, as defined in Chapter 1, a simple intellectualist is someone who is committed to the truth of some instance of the simple equivalence thesis:

### The Simple Equivalence Thesis

Necessarily, S knows how to F if and only if there is some relevant proposition p concerning a way to F such that S stands in the knowledge-that relation to p.

To see why someone committed to the truth of some instance of this equivalence thesis might still reject the corresponding version of the simple identity thesis consider the following version of the simple equivalence thesis:

Necessarily, S knows how to F if and only if there is some way to F  $w$  such that S knows that  $w$  is a way to F.

In endorsing this view of knowledge-how, is one thereby committed to the claim that to know how to F is to know that p, for some proposition p of the form ‘ $w$  is a way to F’? One reason to think not, is that there is a multiple realizability issue lurking here, because for most (if not all) actions there will be as many ways to perform that action as there are ways to skin the proverbial cat. And simple intellectualism tells us that there will be as many different states of knowledge-that that one could be in when one knows how to F as there are ways to F. Suppose, for example, that there is some way to swim  $w_1$  such that Ari knows that  $w_1$  is a way to swim. A proponent of the equivalence thesis stated above will claim that it follows that Ari knows how to swim. But they might reasonably deny that it follows that Ari’s state of knowing that  $w_1$  is a way to swim is *identical* to his state of knowing how to swim. For Ari could have failed to possess this knowledge-that and still known how to swim, because even if Ari does not know that  $w_1$  is a way to swim he can still know how to swim if there is some other way to swim  $w_2$  such that Ari knows that  $w_2$  is a way to swim.

Even if the regress arguments we have considered are sound, it does not obviously follow that simple intellectualism is false, because it is not obvious that a simple intellectualist must endorse the simple identity thesis. Now, this issue is not a serious problem for Noë’s regress argument, as one could restate his

argument using an equivalence claim of the form ‘S has the ability to F if and only if S knows how to F’ instead of the identity claim that to possess the ability to F is to know how to F.

This issue does look to be a serious problem for the contemplation regress, the employment regress, and Hetherington’s regress argument. For, as we have seen, all of these arguments rely on the idea that when one employs (or applies) one’s knowledge how to F one thereby employs (or applies) one’s knowledge that p, for some proposition p. But if the relationship between Ari’s state of knowing how to swim and Ari’s state of knowing that  $w_1$  is a way to swim is not the identity relationship, then it does not obviously follow that if Ari employs his knowledge how to swim that he thereby employs his knowledge that  $w_1$  is a way to swim.

It appears then that even if they are sound the contemplation, employment and Hetherington’s regress arguments only clearly undermine the simple identity thesis, and not the simple equivalence thesis. In which case, it is not clear that these arguments could even undermine simple intellectualism if they are sound. What about sophisticated intellectualism? Recall that a sophisticated intellectualist is someone who is committed to the truth of some instance of the following claim:

The Sophisticated Equivalence Thesis

S knows how to F if and only if, for some relevant proposition p concerning a way to F:

- (i) S stands in the knowledge-that relation to p, and
- (ii) S satisfies X (for some further condition X).

As with the simple equivalence thesis, it is not at all clear that someone committed to some instance of the sophisticated equivalence thesis is thereby also committed to the corresponding version of the sophisticated identity thesis; that is, the claim that to know how to F is to know that p and to satisfy X (for some proposition p, and some condition X). The same kind of issues concerning

multiple realizability will obviously apply here as well. In which case, the contemplation, employment and Hetherington's regress arguments, at best, would only clearly establish that the sophisticated identity thesis is false and not that the sophisticated equivalence thesis is false.

But there is also a further reason to think that these arguments cannot establish that sophisticated intellectualism is false. For suppose that the truth of the sophisticated equivalence thesis did entail the truth of the corresponding identity claim that to know how to F is to know that p and to satisfy X (for some proposition p, and some condition X). Does it follow that if one employs (or applies) one's knowledge how to F one employs (or applies) one's knowledge that p, for some proposition p? I think not, for suppose one endorsed the following version of the sophisticated identity thesis:

The Proxy Identity Thesis

To know how to F is to:

- (i) Stand in the knowledge-that relation to p (for some relevant proposition p concerning a way to F), and
- (ii) To possess the ability to F.

Now imagine that Ari Fs, and in doing so he employs or applies his knowledge how to F. From the proxy identity thesis it follows that there is some proposition p such that Ari knows that p, but does it follow that Ari thereby *employs* or *applies* his knowledge that p? The answer is clearly no; the proxy identity thesis does not identify knowing how to F with knowing that p. Rather, the proxy identity thesis identifies knowing how to F with knowing that p *and* possessing the ability to F. In which case, all that clearly follows from the proxy identity thesis and the fact that Ari employed his knowledge how to F, is that *either* Ari employed his knowledge that p *or* he employed his ability to F.

There are good reasons then to think that even if they were sound the contemplation, employment and Hetherington's regress arguments could not establish that intellectualism in general is false. The same concerns do not apply

to Noë's regress argument, but we have seen that Noë's argument crucially relies on the assumption that neo-Ryleanism is true, which is a view of knowledge-how that any intellectualist must reject. And intellectualists have presented numerous cases that indicate that neo-Ryleanism is false. In which case, Noë's argument is unsound.

Furthermore, our discussion has shown that there are good reasons to think that the contemplation, employment and Hetherington's regress arguments are also unsound. The contemplation regress argument relies on the implausible claim that if one employs one's knowledge that  $p$  then one contemplates the proposition that  $p$ . The employment premise relies on the employment premise, and while this premise is perhaps slightly more plausible than the contemplation premise, an intellectualist can still offer strong considerations in favour of rejecting the idea that whenever we employ our knowledge that  $p$  we also employ a distinct state of knowledge how to employ our knowledge that  $p$ . And the same kind of considerations can be offered for rejecting the application premise that Hetherington's regress argument implicitly relies on.

Ryle's idea that there is a successful regress argument against intellectualism is the most famous objection to the intellectualist view of knowledge-how. However, in this chapter I have shown that it is an unconvincing objection. There are good reasons to conclude that all of the regress arguments I have considered are unsound. And even if the contemplation, employment and Hetherington's regress arguments were sound, they would not establish that all forms of intellectualism are false.

In conclusion, it could be that there is still some successful regress argument against intellectualism out there waiting for us to find it. However, on the basis of the discussion in this chapter, it is safe to say that the kinds of regress arguments that are currently offered against intellectualism give us no reason to be confident that such an argument exists. And, as demonstrated in Chapter 1, the insufficiency objection fails to establish that sophisticated intellectualism is false. In which case, the two most prominent objections to intellectualism both fail to

undermine the intellectualist view of knowledge-how. To find a successful argument against all forms of intellectualism I submit we must turn to the new arguments I presented in Chapter 3 for the conclusion that knowledge-that is not necessary for knowledge-how.

## Bibliography

---

Alter, Torin. 2001. "Know-how, Ability, and The Ability Hypothesis", *Theoria* 7: 229–39.

Bealer, George. 1993. "The Incoherence of Empiricism", in Steve Wagner and Richard Warner (eds), *Naturalism: A Critical Appraisal*: 99–138, Notre Dame: University of Notre Dame Press.

Bengson, John and Moffett, Marc A. 2007. "Know How and Concept Possession", *Philosophical Studies* 136: 31–57.

Bengson, John, Moffett, Marc A. and Wright, Jennifer C. Forthcoming. "The Folk On Knowing How", *Philosophy and Phenomenological Research*.

Bergman, Michael. 1997. "Internalism, Externalism and The No-Defeater Condition", *Synthese* 110: 399–417.

Brogaard, Berit. Forthcoming a. "What Mary Did Yesterday", *Philosophy and Phenomenological Research*.

Brogaard, Berit. Forthcoming b. "Knowledge-*The* and Propositional Attitude Ascriptions", in Franck Lihoreau (ed.), *Knowledge and Questions*.

Brown, D.G. 1970. "Knowing How and Knowing That, What", in Oscar P. Wood and George Pitcher (eds), *Ryle: Modern Studies in Philosophy*: 213–48, London: Macmillan & Co.

Cath, Yuri. Forthcoming. "The Ability Hypothesis and The New Knowledge-how", *Noûs*.

Chalmers, David J. 1996. *The Conscious Mind*, New York: Oxford University Press.

Chalmers, David J. 2004. "Phenomenal Concepts and the Knowledge Argument", in P. Ludlow, Y. Nagasawa and D. Stoljar (eds), *There's Something About Mary*: 269–98, Cambridge, Mass.: MIT Press.

Chomsky, Noam. 1988. *Language and Problems of Knowledge*, Cambridge, Mass.: MIT Press.

Davies, Martin. 2001. "Knowledge (Explicit and Implicit)", in N.J. Smelser and P.B. Baltes (eds), *International Encyclopedia of the Social and Behavioral Sciences*: 12, 8126–32, Amsterdam: Elsevier Science Ltd.

Evans, Gareth. 1981. "Semantic Theory and Tacit Knowledge", in S. Holtzman and C. Leitch (eds), *Wittgenstein: To Follow a Rule*, London: Routledge & Kegan Paul (Reprinted in Evans, G. 1985. *Collected Papers*: 322–42, Oxford: Oxford University Press).

Fodor, Jerry A. 1968. "The Appeal to Tacit Knowledge in Psychological Explanation", *The Journal of Philosophy* 65: 627–40.

Ginet, Carl. 1975. *Knowledge, Perception, and Memory*, Boston: Reidel.

Hetherington Stephen. 2006. "How to Know (that Knowledge-that is Knowledge-how)", in S. Hetherington (ed.), *Epistemology Futures*: 71–94, Oxford: Oxford University Press.

Horgan, Terence. 1984. "Jackson on Physical Information and Qualia", *Philosophical Quarterly* 34: 147–83.

Huemer, Michael. 2005. *Ethical Intuitionism*, New York: Palgrave.

Hunter, David. 1998. "Understanding and Belief", *Philosophy and Phenomenological Research* 58: 559–80.

Jackson, Frank. 1982. "Epiphenomenal Qualia", *Philosophical Quarterly* 32: 127–36.

Jackson, Frank. 1986. "What Mary Didn't Know", *The Journal of Philosophy* 83: 291–95.

- Jackson, Frank. 2004. "Mind and Illusion", in P. Ludlow, Y. Nagasawa and D. Stoljar (eds), *There's Something About Mary*: 421–42, Cambridge, Mass.: MIT Press.
- Jackson, Frank. 2005. "Consciousness", in F. Jackson and M. Smith (eds), *The Oxford Handbook of Contemporary Philosophy*: 310–33, Oxford: Oxford University Press.
- Johnson, Kent. 2006. "Externalist Thoughts and the Scope of Linguistics", in G. Preyer (ed.), *Protosociology: Contemporary Problems in Cognitive Science* 22: 19–35.
- Koethke, John. 2002. "Stanley and Williamson on Knowing How", *The Journal of Philosophy* 98: 325–8.
- Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean", *Linguistics and Philosophy* 1: 337–55.
- Lewis, David. 1979. "Attitudes De Dicto and De Se", *The Philosophical Review* 88: 513–43.
- Lewis, David. 1998. "What Experience Teaches", in W.G. Lycan (ed.), *Mind and Cognition: An Anthology*: 447–46, Second Edition, Oxford: Blackwell.
- Loar, Brian. 1990. "Phenomenal States", in J. Tomberlin (ed.), *Philosophical Perspectives* 4, Northridge, Calif.: Ridgeview.
- Lockwood, Michael. 1989. *Mind, Brain and the Quantum*, Oxford: Blackwell.
- Lycan, William G. 1996. *Consciousness and Experience*, Cambridge, Mass.: MIT Press.
- Nemirow, Laurence. 1980. "Review of *Mortal Questions* by Thomas Nagel", *Philosophical Review* 89: 473–7.
- Nemirow, Laurence. 1990. "Physicalism and the Cognitive Role of Acquaintance", in W.G. Lycan (ed.), *Mind and Cognition: A Reader*: 490–9, First Edition, Oxford: Blackwell.

- Nemirow, Laurence. 1995. "Understanding Rules", *The Journal of Philosophy* 92: 28–43.
- Noë, Alva. 2005. "Against Intellectualism", *Analysis* 65(4): 278–90.
- Nolan, Daniel. 2001. "What's Wrong with Infinite Regresses?" *Metaphilosophy* 32: 523–38.
- Perry, John. 2001. *Knowledge, Possibility and Consciousness*, Cambridge, Mass.: MIT Press.
- Pettit, Dean. 2002. "Why Knowledge is Unnecessary for Understanding Language", *Mind* 111: 519–50.
- Pettit, Phillip. 2004. "Motion Blindness and the Knowledge Argument", in P. Ludlow, Y. Nagasawa and D. Stoljar (eds), *There's Something About Mary*: 105–42, Cambridge, Mass.: MIT Press.
- Poston, Ted. Forthcoming. "Know how to be Gettiered?" *Philosophy and Phenomenological Research*.
- Pritchard, Duncan. 2005. *Epistemic Luck*, New York: Oxford University Press.
- Pust, Joel. 2000. *Intuitions as Evidence*, New York: Garland.
- Rosefeldt, Tobias. 2004. "Is Knowing-how Simply a case of Knowing-that?" *Philosophical Investigations* 27(4): 370–9.
- Ryle, Gilbert. 1946. "Knowing How and Knowing That", Reprinted in *Gilbert Ryle: Collected Papers Volume Two*: 212–25, 1971, New York: Barnes & Noble.
- Ryle, Gilbert. 1949. *The Concept of Mind*, Chicago: Chicago University Press.
- Schitzwzgebel, Eric. 2002. "A Phenomenal Dispositional Account of Belief", *Noûs* 36(2): 259–60.
- Snowdon, Paul. 2003. "Knowing How and Knowing That: A Distinction Reconsidered", *Proceedings of the Aristotelian Society* 104: 1–29.

- Soames, Scott. 2003. *Philosophical Analysis in the Twentieth Century: Volume 2 The Age of Meaning*, Princeton: Princeton University Press.
- Stanley, Jason. 2005. "Semantic Knowledge and Practical Knowledge", *The Aristotelian Society Supplementary Volume* 79: 131–46.
- Stanley, Jason and Williamson, Timothy. 2001. "Knowing How", *The Journal of Philosophy* 98: 441–4.
- Stich, Stephen. 1978. "Beliefs and Subdoxastic States", *Philosophy of Science* 45: 499–518.
- Stich, Stephen. 1980. "What Every Speaker Knows", *The Philosophical Review* 80: 476–96.
- Thau, Michael. 2002. *Consciousness and Cognition*, Oxford: Oxford University Press.
- Tolhurst, William. 1998. "Seemings", *American Philosophical Quarterly* 35: 293–302.
- Tye, Michael. 2000. "Knowing What It Is Like", in *Consciousness, Color and Content*, Cambridge, Mass.: MIT Press.
- Van Gulick, Robert. 2004. "So Many Ways of Saying No to Mary", in P. Ludlow, Y. Nagasawa and D. Stoljar (eds), *There's Something About Mary*: 365–405, Cambridge, Mass.: MIT Press.
- Weatherson, Brian. Forthcoming. "Doing Philosophy with Words", *Philosophical Studies*.