

## Regarding a Regress

Yuri Cath

*Pacific Philosophical Quarterly* 94, 3 (2013) 358-388

Preprint, please cite the published version

**Abstract:** Is there a successful regress argument against intellectualism? In this paper I defend the negative answer. I begin by defending Stanley and Williamson's (2001) critique of *the contemplation regress* against Noë (2005). I then identify a new argument—*the employment regress*—that is designed to succeed where the contemplation regress fails, and which I take to be the most basic and plausible form of a regress argument against intellectualism. However, I argue that the employment regress still fails. Drawing on the previous discussion, I criticise further regress arguments given by Hetherington (2006) and Noë (2005).

What kind of knowledge does one have when one *knows how to do something*? One important view, often known as *intellectualism*, is that knowing how to do something (or “knowledge-how”) is a kind of propositional knowledge (or “knowledge-that”). The most famous objection to this view, which originates with Ryle (1946, 1949), is that intellectualism must be false of pain of avoiding an infinite and vicious regress. That is, Ryle introduced the idea that there is a successful *regress argument* against intellectualism. A regress argument against intellectualism is an argument that 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 inferred that intellectualism is false. I assume that for such an argument to succeed the inference steps that generate the regress must be valid, any required premises (other than the *reductio* premise) must be sound, and the regress has to be genuinely vicious rather than merely benign.

Despite its fame, Ryle's regress objection has proven to be the most elusive of existing objections to intellectualism. For one thing, there is little agreement over not only the status but also the very structure of the best version of a regress argument against intellectualism. Part of the problem is that Ryle's own regress argument is not an argument against the view that knowledge-how is a kind of knowledge-that. Rather, the target of Ryle's argument is what he called *the intellectualist legend*—which is a view about the nature of intelligent actions, not knowledge-how.<sup>1</sup> Ryle clearly thought that this argument somehow also supported the conclusion that one cannot define “knowing how” in terms of “knowing that” (1949: 32). But the irony is that while Ryle made the idea that there is a successful regress argument against intellectualism famous he never explicitly stated such an argument himself.

Stanley and Williamson (henceforth S&W), in their important (2001) defence of intellectualism, offered a reconstruction of Ryle's regress argument as an argument explicitly against intellectualism about knowledge-how, and they argued that this argument is unsound. But Ryle's idea that there is a successful regress argument against intellectualism still lingers. Some anti-intellectualists, like Hetherington (2006), have responded by claiming that S&W misrepresent Ryle's implicit regress argument against intellectualism and that, properly understood, this argument does succeed. Others, like Noë (2005), have suggested that there is a successful regress argument against intellectualism closely related to, albeit perhaps not identical with, Ryle's argument. Furthermore, Noë argues that S&W's critique fails to

undermine the argument that they offer as a reconstruction of Ryle's implicit regress argument.

The purpose of this paper is to clarify and advance these often rather confusing debates about whether there is a successful regress argument against intellectualism. I begin in §1 by offering a more precise characterisation of intellectualism. In §2 I consider the argument S&W examine—what I will call *the contemplation regress*—and I defend S&W's conclusion that this argument does not succeed against Noë's criticisms. In §3 I offer a new regress argument against intellectualism, which is also partly inspired by Ryle's work—what I call *the employment regress*. The employment regress is, I think, the most basic and plausible form of a regress argument against intellectualism. However, in §4 I shall argue that the employment regress is still not, in the end, a convincing argument against intellectualism. In §5 I show how related considerations undermine the regress arguments given by Hetherington (2006) and Noë (2005). On the basis of the previous discussion, I conclude that there is no reason to think that there is any successful regress argument against intellectualism.<sup>2</sup>

## 1. Intellectualism

Intellectualism (as I will use the term) is the view that knowledge-how is a kind of knowledge-that.<sup>3</sup> On this view then knowing how to F is a matter of standing in the knowledge-that relation to the right kind of proposition. But what proposition? Intellectualists typically hold that knowing how to 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 that which 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. More precisely, the simple intellectualist 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 identity 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).

One obvious objection to simple intellectualism is that merely standing in the knowledge-that relation to some proposition concerning a way to F does not look to be a sufficient condition for knowing how to F.<sup>4</sup> In response to this worry, intellectualists often adopt some version of what we can call *sophisticated intellectualism* which 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 do not think that this is a sufficient a condition 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.<sup>5</sup> The sophisticated intellectualist then is someone who is committed to the truth of some instance of the following equivalence claim:

**The Sophisticated Equivalence Thesis:** Necessarily, 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)<sup>6</sup>

It is important to note that the regress arguments against intellectualism that Ryle has inspired all target some version of the simple identity thesis. This suggests two possible worries about the scope of such arguments. First, even if they did succeed could these arguments be extended to all versions of sophisticated intellectualism? Furthermore, it is conceivable that a simple intellectualist might only endorse the simple equivalence thesis and not the simple identity thesis.<sup>7</sup> In which case, one might worry that arguments against the simple identity thesis will not automatically impugn the simple equivalence thesis. However, in what follows I will ignore such questions and focus only on the issue of whether there is a successful regress argument against the simple identity thesis. For the sake of argument, I will assume that any successful argument against the simple identity thesis would also be a successful argument against intellectualism in general.

## 2. The Contemplation Regress

The argument S&W 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, we also need to assume, for the sake of argument, that intellectualism is true or that some thesis entailed by intellectualism is true. S&W (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, so  $\Phi(F)$  is the proposition that is the value of  $\Phi$  when its input is the action F. S&W reconstruct Ryle's regress argument as an argument against the following identity claim 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, with the only major difference between the two theses being that (RA) does not include the condition that the proposition in question has to concern a way to F.

S&W 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. (S&W 2001, p. 414)

This is the regress argument against intellectualism that S&W attribute to Ryle.<sup>8</sup> S&W point out (2001, p. 414) that if this argument is to 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>9</sup>

One might always debate whether this argument can be legitimately attributed to Ryle, given that Ryle himself only explicitly states a regress argument against the intellectualist legend and not intellectualism. But such interpretative issues are not our concern here. Our concern is simply whether this argument succeeds or not. Hence, I will refer to this argument as *the contemplation regress*, so as to remain neutral on the interpretative question. Does the contemplation regress succeed? S&W argue that it fails, and they reach this conclusion in three steps.

### *Step 1: Restrict the Action Premise*

The first step 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 S&W (2001, p. 414) note, 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 she did not know how to win a fair lottery, since she only won by sheer chance.

One might worry that the digestion example is not a counterexample because digesting is not something that we do; rather it is a process that occurs inside our bodies (Noë 2005, p. 279). 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. Similarly, sweating is a mere bodily process—like the beating of our hearts—and not an action that we as agents perform. But grammatically speaking, sweating—unlike the beating of our hearts—is something that we do and so, in this limited sense of the word "action", it is an action.

The digesting and sweating cases are counterexamples to the action premise because these "doings" are just not the kind of thing that one knows how to do. And there are also counterexamples (not discussed by S&W) where one does know how to F but one 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 S&W (2001, p. 415), the lesson of their counterexamples is that the action premise is only plausible if we restrict the range of actions that can be values for "F" to *intentional actions*, that is, S&W 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

Indeed, this intentional action premise is an actually an important component of S&W's full account of knowledge-how. So, S&W do acknowledge a significant connection between knowledge-how and action, but only once we restrict our attention to intentional actions. And, as S&W point out, neither the digestion 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 S&W's critique, it is worth mentioning that one could appeal to alternative restrictions on the range of actions that can be values for "F". Indeed, as S&W point out, Ryle himself appears to endorse a version of the action premise that is restricted to *intelligent* rather than *intentional* actions. Ryle thought that intelligent actions should be analysed as actions that exercise our knowledge-how. In particular, he seemed 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.

S&W grant that restricting the action premise to intelligent actions also avoids the digestion and lottery counterexamples. However, S&W do not regard this way of saving the action premise as being significantly different from their own suggestion. For after mentioning Ryle's alternative restriction they simply say that 'the range of actions under consideration must be restricted to intentional actions, or perhaps a proper subset thereof' (2001, p. 415), the thought being that any intelligent action will also be an intentional action.<sup>10</sup>

### *Step 2: Deny the Contemplation Premise*

The first step in S&W's critique of the contemplation regress is to claim that the action premise is false unless we restrict it to intentional actions. The second step is to claim that the contemplation premise is false. In support of this claim S&W cite the following passage from Ginet (1975, p. 7):

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.

S&W 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, p. 415).

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

S&W (2001, p. 415-16) 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.

The third and final step in S&W's critique is to point out that while this response might save the contemplation premise it would not save the contemplation regress. 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 intentional actions or some proper subset thereof. And, they claim, the action premise is only plausible if it is so restricted.

S&W conclude then 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 (2001, p, 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?*

Noë (2005, p. 278-82) argues that S&W fail to show that the contemplation regress is unsound. Noë's main criticism is directed at S&W's 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 S&W 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 this conclusion, Noë points out, 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, p. 282)

Perhaps Noë is right is that a proponent of the contemplation regress would be best advised to respond to Ginet's case by claiming that whenever one employs one's knowledge that p one does intentionally contemplate the relevant proposition, its 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 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 my doing so. But why in the first place should we believe 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 a 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 by rejecting this assumption rather than (RA). I doubt that any such argument could be given and so I think we must agree with S&W that the contemplation regress fails.

### 3. The Employment Regress

The contemplation regress does not succeed. Is there a more promising regress argument against intellectualism? In this section I identify a regress argument that is related to the contemplation regress 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 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 the implicit target 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. I will discuss Hetherington's argument in §5. My aim now is to construct what I take to be the clearest and most plausible statement of this kind of argument. To help introduce 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 *F*s, one employs knowledge how to *F*

Unlike the contemplation premise, the action premise is *prima facie* plausible, as is suggested by the fact that both intellectualists, like S&W, and anti-intellectualists, like Ryle and Noë, all endorse restricted versions of this claim.<sup>11</sup> An interesting point about the action premise—that is not noted by any of its aforementioned proponents—is that the unrestricted version of this thesis by itself generates an infinite regress. For suppose that Hannah performs some action *F*<sub>1</sub>. By the action premise Hannah employs knowledge how to *F*<sub>1</sub>. But employing one's knowledge how to *F*<sub>1</sub> is a legitimate value for “*F*” in the *unrestricted* action premise for, at least grammatically speaking, employing one's knowledge-how is something that one *does*. In which case, we can reapply the action premise to conclude that Hannah employs knowledge how to employ her knowledge how to *F*<sub>1</sub>. 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 is distinct from one's state of knowledge how to employ one's knowledge how to F)

For to accept 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 this conclusion tells us that we never employ our knowledge-how.<sup>12</sup>

### *The employment regress*

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

To construct a regress argument 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<sub>1</sub>, and that Ari's state of knowing how to employ his knowledge that p<sub>1</sub> is distinct from his state of knowing that p<sub>1</sub>. By the identity premise, it follows that Ari thereby employs his knowledge that p<sub>2</sub> (for some proposition p<sub>2</sub> concerning a way to employ one's knowledge that p<sub>1</sub>). By the employment premise, it follows that Ari also employs knowledge how to employ his knowledge that p<sub>2</sub>, and that Ari's state of knowing how to employ his knowledge that p<sub>2</sub> is distinct from his state of knowing that p<sub>2</sub>. 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 we must also possess and employ an infinite number of further and distinct states of knowledge-that.<sup>13</sup> This conclusion is absurd for given that we are finite beings what this conclusion 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 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 a successful argument against intellectualism? One might think that it is at least a more promising argument than the contemplation regress. For the employment regress does not rely on the dubious idea that one must contemplate the proposition that  $p$  if one is to employ one's knowledge that  $p$ .

However, an intellectualist will no doubt suspect that they can reasonably reject 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. 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 can be direct does not by itself generate an infinite and vicious regress. But still why should we think that one can never employ one's knowledge that  $p$  without also employing 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 the unrestricted version of the action premise is subject to clear counterexamples. The intellectualist may well suspect then that we should expect to find analogous counterexamples to the employment premise.

Still, it does not follow from the fact 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 that still generates an infinite and vicious regress together with the simple identity thesis. I will examine such issues in §4. But first it will be useful to look at an attempt by Ryle to motivate something very close to the idea expressed by the employment premise.

### *Rylean motivations for the employment premise*

Consider the following passage from Ryle (1946):

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, p. 216-17)

There are numerous different ideas and arguments suggested by Ryle's brief discussion of this Carroll inspired example. For one thing, I think Ryle uses this example to support something like the insufficiency objection mentioned in §1: that intellectualism is false because no mere knowledge-that is sufficient for knowledge-how. Ryle appears to be making this kind of objection by claiming that the schoolboy could know any given proposition while still failing to know how to reason.<sup>14</sup>

However, Ryle also infers a slightly different insufficiency claim from this example, namely, that knowing that *p* does not suffice for knowing how to employ one's knowledge that *p*. Ryle often talks of a 'gap' or 'gulf' 'between...acknowledging principles in thought and intelligently applying them in action' (1946, p. 218). But, at the same time, Ryle also refers to a 'gulf' 'between having the postulated knowledge of those facts and knowing how to use or apply it' (1946: 218). In other words, Ryle thinks that one can know that *p* without knowing how to employ or apply one's knowledge that *p*.

How does the schoolboy 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 example as a case where someone stands in the knowledge-that relation to some proposition(s) but they do not know how to employ this knowledge-that.

Ryle claims then that merely knowing that *p* is not a sufficient condition for knowing how to employ one's knowledge that *p*. This claim does not establish the employment premise but it does support the claim made by the bracketed clause in 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 employment premise and the simple identity thesis together generate an infinite regress of employments of knowledge-that. The bracketed clause tells us that each new state of knowledge-that in this infinite regress is 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 stated by the bracketed clause, 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 this schoolboy case as supporting a claim that is even more closely related to the employment premise, namely, 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 the schoolboy 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 schoolboy, for Ryle, is that only the former '[knows] how to move from acknowledging some facts to acknowledging others.'

Similarly, Ryle concludes, on the basis of an example involving a stupid chess-player who knows all the same propositions concerning chess strategies that are known to the clever chess-player, but who is unable to intelligently apply them 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, p. 216). Ryle’s seems to suggest here not only that employing one’s knowledge—that is an action which 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. For Ryle will think that without such knowledge—how the stupid chess-player ‘might not see how to apply’ (Ryle 1946, p. 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 schoolboy 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 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*. For recall that Ryle only endorses a restricted version of the action premise according to which if one Fs *intelligently* then one employs knowledge how to F. And Ryle holds that employments of knowledge—that are actions that can be performed more or less intelligently. Presumably then 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 can be used to generate a plausible regress argument against intellectualism is an issue we will address in the next section. For now it will suffice to simply note Ryle’s attempt to motivate something like the employment premise.

#### 4. Intellectualist Responses

What reasons might an intellectualist offer for rejecting the employment premise? To begin with, let us consider Ryle’s attempts to motivate the employment premise. Do Ryle’s schoolboy and stupid chess-player cases 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 that failing to know how to employ one’s knowledge that p will *sometimes* result in one’s failing 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 such a failure. Suppose that all of the cases Ryle describes are in fact cases where someone’s fails to employ their knowledge that p because they do not know how to employ their knowledge that p. Even given this assumption, it does not follow that either possessing or employing

knowledge how to employ one's knowledge that p is a necessary condition of employing one's knowledge that p.

Furthermore, as suggested earlier, an intellectualist will likely suspect that there will be counterexamples to the employment premise analogous to those we find for the action premise. Before we consider whether there are such counterexamples note that *if* there are such counterexamples a proponent of the employment regress could not respond to them by merely restricting the employment premise like so:

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

For one cannot validly derive an infinite regress from this claim and the simple identity thesis. Rather, for the derivation to be valid one would have to endorse the following claim:

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

But 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 should only be interpreted as applying to intentional employments of knowledge-that. And if intelligent actions are simply a kind of intentional action then the same point applies to Ryle's version of the employment premise that is restricted to intelligent employments of knowledge-that.

Of course, it does not follow from the fact there are counterexamples to the action premise that there will be counterexamples to the employment premise. Indeed I think it is actually quite difficult to identify clear counterexamples to the employment premise. For one thing, it seems difficult to imagine a counterexample to the employment premise that parallels S&W's lottery case or the vase case, that is, a case where someone employs their knowledge that p but only by mere luck or accident and thereby without employing knowledge how to employ their knowledge that p.<sup>15</sup> And even if there are such cases I think it would be reasonable for a proponent of the employment regress to respond to them by restricting the employment premise in something like the following way:

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

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

Can we find counterexamples to the employment premise by looking for analogues of the digestion and sweating cases? The only kind of case 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 this tacit knowledge.

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 rather 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; that is, it is something that Mary does *directly*. This conclusion will seem 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 plausible that we employ them directly. But, nonetheless, I think there is an inherent problem with this strategy of trying to identify counterexamples to the employment premise. The problem is that it is not at all clear that such states should really be thought as being genuine states of propositional knowledge, rather than merely some other kind of informational state. For example, propositional knowledge requires justification or warrant but, arguably, such notions are not even applicable to the states of "tacit" or "implicit" knowledge appealed to in psychology; as Davies (2001, p. 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.' And Evans (1981) and Stich (1978, 1980) have argued that 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 contested this issue is, appealing to such states does not look to be a promising way of identifying counterexamples to the employment premise. Rather, if the intellectualist is to identify a convincing counterexample to the employment premise they need 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 informational state.

I think a better case for the intellectualist to appeal to would again be Ginet's door case. For this is a case where someone clearly employs a genuine state of knowledge-that. And it is also a case where this action of employing knowledge-that appears not to be performed either consciously or intentionally. The intellectualist might argue then that because Ginet's 'automatic' employment of his knowledge-that is not an intentional action then it is an action that he performs without employing knowledge how to perform it.

Is this diagnosis of Ginet's door case correct? That is, is it 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 door by turning the knob and pushing it. But nor do I consciously or intentionally employ my knowledge that one can open the door by turning the knob and yet—as Ginet's example makes clear—it still seems correct to say that, in some sense, I employ this knowledge-that in opening the door. A proponent of the employment regress might claim then that, analogously, in opening the door I do employ my knowledge how to employ my knowledge that one can open the door by turning the knob, even though I do not do so either consciously or intentionally.

It is not clear to me what the correct thing is to say here. 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 this knowledge-that, and that I employ this knowledge-how when I open the door. But, on the other hand, it seems far clearer to me 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 one might infer that I thereby 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 suppose one held the following view: Necessarily, one knows how to *F* if and only if one possesses the ability to *F*—a view that in the recent literature has been called *neo-Ryleanism*.<sup>16</sup> From the assumption that I have the ability to employ my knowledge that one can open the door by pushing the knob, a neo-Rylean would infer that I know how to employ this knowledge-that, because they hold that possessing the ability to *F* entails knowing how to *F*. Furthermore, neo-Ryleans typically make the further claim that to know how to *F* just is to possess the ability to *F*. And if this identity claim is true then if I employ an ability to employ my knowledge that *p* it follows that I thereby 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 it is likely that Ryle himself is implicitly assuming the truth of neo-Ryleanism when he appeals to his schoolboy and stupid chess-player cases as 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 schoolboy 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, p. 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 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 have offered numerous plausible counterexamples to the idea that possessing the ability to *F* is either identical with, or entails, knowing how to *F*.<sup>17</sup> The intellectualist then can happily agree 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 that if one employs one's knowledge that *p* then one employs *knowledge how* to employ one's knowledge that *p*.

Now, S&W are intellectualists who hold that if one has the ability to *F* *intentionally* then one knows how to *F*.<sup>18</sup> But this claim is consistent with the point being made here. For one thing, whilst S&W accept this entailment they do not think that knowing how to *F* entails having the ability to *F*, so they are clearly not committed to identifying knowing to *F* with possessing the ability to *F*. So, they are not committed to the idea that to employ an ability to *F* intentionally just is to employ one's knowledge how to *F*. And, anyway, such an identity claim could, at best, only be used to help motivate a version of the employment regress that relied on the claim that if one employs 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 then 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 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 decisive counterexample to it. 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 can 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 all cases of someone employing their knowledge that  $p$  are also cases 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 §2, digesting and sweating are also things that, in the grammatical sense, we do. But it would obviously be a mistake to infer from this 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? I think we must conclude that, at best, the employment regress presents a very inconclusive case against intellectualism. 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 relies on the idea that unlike employments of knowledge-how employments of knowledge-that cannot be direct. But it is not clear that the intellectualist must accept this asymmetry between knowledge-how and knowledge-that, for it is not obviously the case that we must *always* employ knowledge how to employ our knowledge that  $p$  whenever we employ our knowledge that  $p$ . And while it is difficult to describe a really clear counterexample to the employment premise the intellectualist can still offer powerful considerations for thinking that it is false.

## 5. Related Regress Arguments

We have seen that there are strong reasons for thinking that both the employment regress and the contemplation regress are unsound. In this section I want to show how broadly related problems arise for the regress arguments presented by Hetherington (2006) and Noë (2005).

### *Hetherington's regress*

Hetherington presents his regress argument 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, p. 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 the idea is that  $\Phi$  is some 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 (where p is some proposition concerning a way to F) which one knows how to, and one does, apply so as to do F.

And Hetherington's regress argument against (R) is meant to proceed like so. Suppose Ari Fs and he knows how to F. From (R) three things follow, namely: Ari knows that  $p_1$  (for some proposition  $p_1$  concerning a way to F); Ari applies his knowledge that  $p_1$ ; and Ari 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 knowledge that  $p_1$ ); Ari applies his knowledge that  $p_2$ ; and Ari 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 reply (R), and so on *ad infinitum*.

Hetherington thinks that this regress is vicious because he assumes that each new action of applying knowledge-that will be distinct from its predecessor, and that each new state of knowledge-that will be more complex than, and distinct from, its predecessor. I think that as (R) is formulated these assumptions do not clearly follow but, for the sake of argument, let us grant that the regress generated by (R) is vicious and that, therefore, we must reject (R).

What is the import of this conclusion? Hetherington (2006, p. 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.' Hetherington's talk of "intellectualism-as-applied-to-our-intelligently-performed-actions" is somewhat difficult to interpret. But Hetherington (2006. 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.' Hetherington thinks that it is a virtue of his reconstruction of what he takes to be Ryle's regress argument that, unlike the contemplation regress, the reasons S&W 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, so one never digests one's food whilst knowing how to digest one's food. And Hetherington's (R) does not include the claim if one employs or applies one's knowledge that p then one contemplates the proposition that p.

To properly assess Hetherington's argument it will help to notice 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 these claims:

**(R1)** 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).

**(R2)** 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.

**(R3)** 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.

Having 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 (R1), 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 (R1) they can still consistently deny (R2) and/or (R3), 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 (R2) and/or (R3), as to do so would be consistent with their commitment to (R1). So, if Hetherington's regress argument against intellectualism is to succeed it must be that (R2) and (R3) are independently plausible, and not that they are entailed by the intellectualist view of knowledge-how, as Hetherington suggests.

Furthermore, an intellectualist faced with Hetherington's argument can easily justify rejecting (R3). 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, (R3) 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.

But 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 (R3) whilst still endorsing (R1).

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 that one Fs and *in so doing one applies one's knowledge how to F*, and not merely that one Fs and one also knows how to F. To focus on the deeper issues facing Hetherington's argument then, I suggest that we replace his (R) with (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

The vase example is not a problem for (R\*) because one does not apply one's knowledge how to knock the vase off the mantelpiece when one accidentally knocks it off the mantelpiece. And we can generate essentially the same infinite regress from (R\*) that we did from (R).<sup>19</sup> In which case, if we assume again that this regress is vicious then we should reject (R\*).

I think this argument against (R\*) is basically the argument Hetherington has in mind when he gives his 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\*):

**(R1\*)** 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).

**(R2\*)** 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.

**(R3\*)** 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 (R1), any intellectualist must accept (R1\*), 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 p concerning a way to F). What about (R2\*)? An intellectualist committed to some version of the simple identity thesis must accept (R2\*). 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.

Does Hetherington's argument show us that the simple identity thesis is false? Given that an intellectualist who endorses the simple identity thesis must accept both (R1\*) and (R2\*) this issue boils down to whether such an intellectualist can offer reasons for rejecting (R3\*) that are not also reasons for rejecting (R1\*) or (R2\*).

The first thing to note is that any intellectualist can consistently deny (R3\*) whilst maintaining that (R1\*) and (R2\*) are true, for no form of intellectualism entails the conditional stated by (R3\*). Furthermore, note that an intellectualist who is committed to the simple identity thesis will regard (R3\*) 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.

And this claim is obviously false. For this claim 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 knowledge how to apply one's knowledge how to F. And, as we saw in §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 (R3) is that he assumes that something like the employment premise is correct. In particular, I think Hetherington implicitly relies on 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

For the basic idea behind Hetherington's argument appears to be that given the intellectualist assumption that to know how to F is to know that p (for some proposition p concerning a way to F), and the assumption 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 (let us grant) vicious regress and, therefore, one of these two assumptions must be false. Because Hetherington assumes that the application premise is true he thereby concludes that intellectualism is false.

This is, I think, 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 then, we must conclude that Hetherington's argument, at best, presents an inconclusive case against intellectualism.

### *Noë's possession regress argument*

All of the regress arguments against intellectualism that we have considered have been arguments that claim that if intellectualism is true then an infinite and vicious ensues whenever we *employ*, or *apply*, our knowledge-how. Noë (2005), however, has sketched a regress argument that is meant to show that if intellectualism is true then an infinite and vicious ensues whenever we *possess* knowledge-how. Noë presents his 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. (2005, p. 285–286)

I think the simplest way of representing the regress argument Noë appears to have in mind here is that it is an argument that generates an infinite regress from the following three premises:

- (N1) If one knows that p then one possesses the ability to F (for some action F)
- (N2) To possess the ability to F is to know how to F
- (N3) To know how to F is to know that p (for some proposition p concerning a way to F)

Now suppose that Hannah knows that  $p_1$ , for some proposition  $p_1$ . By (N1) it follows that Hannah possesses the ability to  $F_1$ , for some action  $F_1$ . By (N2) it follows that Hannah's ability to  $F_1$  consists in her knowing how to  $F_1$ . By (N3) it follows that Hannah's knowing how to  $F_1$  consists in her knowing that  $p_2$ , for some proposition  $p_2$ . But then we can reapply (N1) to conclude that Hannah possesses the ability to  $F_2$ , for some action  $F_2$ . By (N2) it follows that Hannah's ability to  $F_2$  consists in her knowing how to  $F_2$ . By (N3) Hannah's knowing how to  $F_2$  consists in her knowing that  $p_3$ , for some proposition  $p_3$ . But then we can

reapply (N1) to conclude that Hannah possesses the ability to  $F_3$ , for some action  $F_3$ , and so on *ad infinitum*.

Together (N1), (N2) and (N3) generate an infinite regress of states of knowledge-that. Noë holds that the right response to this regress is to reject (N3), which is just the simple identity thesis. It is not clear that this regress must be 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 is to reject (N3), which is just the simple identity thesis. How might an intellectualist respond to this argument? Presumably, they would not deny (N1), as Noë is surely right that it is a precondition of knowing that  $p$  (for any proposition  $p$ ) that one possess certain abilities or dispositions. But an intellectualist will deny (N2), as (N2) is just a statement of neo-Ryleanism and, as mentioned earlier, intellectualists have argued that neo-Ryleanism is false on the grounds that possessing the ability to  $F$  is neither necessary nor sufficient for knowing how to  $F$  (see fn. 15).

Noë acknowledges that intellectualists like S&W will deny (N2) but he contests the reasons they give for rejecting this premise in his argument. I do not wish to examine Noë's criticisms of S&W's case against neo-Ryleanism. For one thing, I do not think these criticisms 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 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 contraries. 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. But then the truth of neo-Ryleanism entails the falsity of any version of either simple or sophisticated intellectualism. In which case, the key premise that Noë's regress argument relies upon is a premise which by itself would establish that intellectualism is false. The real issue here then is the status of neo-Ryleanism, a view which intellectualists must reject and which they have offered numerous counterexamples to.

## 6. Conclusion

We have examined four regress arguments against intellectualism, all of which can (to some degree) be seen as drawing their inspiration from Ryle. Our discussion strongly suggests that these arguments all fail. The contemplation regress relies on the implausible claim that if one employs one's knowledge that  $p$  then one contemplates the proposition that  $p$ . The employment regress relies on the employment premise, and while this premise is *prima facie* more plausible than the contemplation premise, an intellectualist can still offer strong considerations for rejecting it. And the same kind of reasons can be offered for rejecting the application premise that Hetherington's regress argument implicitly relies on. Finally, Noë's regress argument is redundant as it crucially relies on the assumption that neo-Ryleanism is true, which is a view of knowledge-how that any intellectualist must reject. Furthermore, the standard counterexamples to neo-Ryleanism suggest that Noë's argument is also unsound. There may still be some other successful regress argument against intellectualism out there waiting to be found. However, the kinds of regress arguments that have been offered against intellectualism give us little reason to think that such an argument exists. I submit we must look elsewhere then if we are to find a successful argument against intellectualism.<sup>20</sup>

## University of East Anglia

### References

- Cath, Y. (2008). *A Practical Guide to Intellectualism*, PhD Thesis Australian National University.
- Cath, Y. (2009). ‘The Ability Hypothesis and the New Knowledge-how’, *Noûs* Volume 43:1, pp. 137–56.
- Cath, Y. (2011) “Knowing How without Knowing That”, in J. Bengson and M. Moffett (eds) *Knowing How: Essays on Knowledge, Mind, and Action*, Oxford: Oxford University Press, pp. 113–35.
- Bengson, J. and Moffett, M. (2011). “Two Conceptions of Mind and Action: Knowing How and the Philosophical Theory of Intelligence”, in J. Bengson and M. Moffett (eds) *Knowing How: Essays on Knowledge, Mind, and Action*, Oxford: Oxford University Press, pp. 3–58.
- Bengson, J. and Moffett, M. (2007). ‘Know How and Concept Possession’, *Philosophical Studies* 136, pp. 31–57.
- Bengson, J. and Moffett, M. (2008). “The Folk On Knowing How”, *Philosophical Studies* 142, pp. 387–401.
- Brogaard, B. (2008). “Knowledge-*The* and Propositional Attitude Ascriptions”, *Grazer Philosophische Studien* 77, pp. 147–190.
- Davies, M. (2001). “Knowledge (Explicit and Implicit)”, in N.J. Smelser and P.B. Baltes (eds), *International Encyclopedia of the Social and Behavioral Sciences*, Amsterdam: Elsevier Science Ltd: 12, pp. 8126–32.
- Evans, G. (1981). “Semantic Theory and Tacit Knowledge”, in S. Holtzman and C. Leitch (eds), *Wittgenstein: To Follow a Rule*, Routledge and Kegan Paul; reprinted in his *Collected Papers*. Oxford: Oxford University Press, pp. 322–42.
- Fodor, J. A. (1968). “The Appeal to Tacit Knowledge in Psychological Explanation”, in *The Journal of Philosophy* 65, pp. 627–40.
- Ginet, C. (1975). *Knowledge, Perception, and Memory*, Boston: Reidel.
- Hetherington, S. (2006). “How to Know (that Knowledge-that is Knowledge-how)”, in S. Hetherington (ed) *Epistemology Futures*, Oxford: Oxford University Press, pp. 71–94.
- Hunter, D. (1998). “Understanding and Belief”, *Philosophy and Phenomenological Research* 58, pp. 559–80
- Noë, A. (2005). “Against Intellectualism”, *Analysis* 65, pp. 278–90.
- Ryle, G. (1946). “Knowing How and Knowing That”, Reprinted in his (1972) *Gilbert Ryle: Collected Papers Volume Two*, New York: Barnes and Noble, pp. 212–25
- Ryle, G. (1949). *The Concept of Mind*, Chicago: Chicago University Press.
- Snowdon, P. (2003). “Knowing How and Knowing That: A Distinction Reconsidered”, *Proceedings of the Aristotelian Society*, 104, pp. 1–29.
- Stanley, J. (2011). *Know How*, Oxford: Oxford University Press.
- Stanley, J. and Williamson, T. (2001). “Knowing How”, *The Journal of Philosophy* 98, pp. 411–44.
- Stich, S. (1978). “Beliefs and Subdoxastic States”, *Philosophy of Science* 45, pp. 499–518.
- Stich, S. (1980). “What Every Speaker Knows”, *The Philosophical Review* 80, pp. 476–96.

---

<sup>1</sup> See Snowdon (2004) for discussion.

<sup>2</sup> Bengson and Moffett (2011) and Stanley (2011) have also recently defended the claim that there is no successful regress argument against intellectualism. This paper is broadly related in spirit to these works but was developed independently and is based on material in my PhD (Cath 2008). For reasons of space I will not discuss these works here.

<sup>3</sup> This use of “intellectualism” is common but sometimes the term is also used in a broader way. For example, Bengson, Moffett and Wright (2008, fn. 5) offer a broader definition which includes any view that analyses knowledge-how in terms of some kind of propositional attitude. And, as noted, Ryle uses the related term “the intellectualist legend” as a name for a view of the nature of intelligent actions. My concern in this paper is only with the question of whether there is a successful regress argument against the view that knowledge-how is a kind of knowledge-that. I examine Ryle’s regress argument against the intellectualist legend in Cath (2008) (where I argue that it is best understood as an argument by dilemma of which only one of the horns is a vicious regress).

<sup>4</sup> 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 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 the simple intellectualist equates knowledge-how with whilst failing to know how to F. See Brogaard (2008, pp. 183–85) for a version of simple intellectualism that is designed to avoid this insufficiency objection.

<sup>5</sup> So, for example, in response to the sufficiency objection, S&W (2001) endorse a view according to which knowing how to F is not only a matter of standing in the knowledge-that relation to the right kind of proposition concerning a way to F but one also has to entertain that proposition under a *practical mode of presentation*. Bengson and Moffett (2007) also endorse a version of sophisticated intellectualism as they offer a view on which knowing how to F is not only a matter of standing in the knowledge-that relation to the right kind of proposition concerning a way to F but one also has to *minimally understand* that way. See Cath (2011, fn. 1) for related discussion.

<sup>6</sup> As with simple intellectualism, the sophisticated intellectualist might also go a step further and endorse the corresponding identity claim.

<sup>7</sup> Take an intellectualist who endorses the thesis that S knows how to F iff there is some way  $w$  such that S knows that  $w$  is a way for S to F. One reason to think this intellectualist is not thereby committed to the corresponding identity thesis is that there is a multiple realizability issue lurking here, because for most actions there will be as many ways to perform that action as there are ways to skin a cat. And this version of the simple equivalence thesis 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 that there is some way to swim  $w_1$  such that Ari knows that  $w_1$  is a way for him to swim. Our intellectualist will think it thereby 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 for him to swim he will still know how to swim (according to our simple intellectualist) if he knows that  $w_2$  is a way to swim, for some other distinct way to swim  $w_2$ . See Bengson, Moffett and Wright (2008, fn. 2) for related discussion.

<sup>8</sup> S&W also note (2001, p. 414) that Ryle himself would have presumably endorsed an even stronger version of the contemplation premise, namely, that “employment of knowledge-that requires a *prior* action of contemplating a proposition.” However, as S&W 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” is 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>9</sup> One also has to assume the substitutivity of identicals within the scope of ‘employs’ (which seems reasonable), and 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 series. Suppose we have shown that some thesis, if true, generates an infinite regress of actions, A1, A2, A3, A4 and so on *ad infinitum*, where each action in the series is distinct from its predecessor so  $A1 \neq A2$ , and  $A2 \neq A3$ , and  $A3 \neq A4$  and so on. 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 because, for example, it could be that after A1 and A2 every ‘new’ action in the regress is identical to either A1 or A2 as follows:  $A3 = A1$ ,  $A4 = A2$ ,  $A5 = A1$ ,  $A6 = A2$ ,  $A7 = A1$  and so on. 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.

<sup>10</sup> If S&W are right that *all* intentional actions are employments of knowledge-how this 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 all still employments of knowledge-how.

<sup>11</sup> Indeed, for S&W (2001, p. 443) the plausibility of the intentional action premise explains why philosophers have been swayed by Ryle’s regress objection: ‘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.’

<sup>12</sup> Note that there are two ways one might deny this claim, one way would be to deny the claim that if we employ our knowledge how to F then we employ knowledge how to employ our knowledge how to F. This response would block the regress. Alternatively, one could accept the regress but claim that it is benign by denying the bracketed clause that states that these two states of knowledge-how are distinct. The idea 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 the whole claim if we are to avoid saying that an infinite and vicious regress ensues whenever anyone employs their knowledge how to do something.

<sup>13</sup> Although, strictly speaking, 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 (see fn. 9). I ignore the possibility that such a regress loops back on itself simply because it strikes me as being implausible. If one was worried about this one could ensure that the regress is vicious by reformulating the employment premise to get something like the following claim: If one employs knowledge that p, one employs knowledge how to employ one’s knowledge that p (where one’s employment of one’s knowledge how to employ one’s knowledge that p is distinct from and *prior to* one’s employment of one’s knowledge that p). But this seems to me to be a less plausible version of the employment premise, so I think it is more useful to focus on the employment premise as stated above.

<sup>14</sup> The idea that it is possible for one to stand in the knowledge-that relation to *any* proposition p whilst failing to know how to perform some action F is clearly false. For example, just let p be the proposition that one knows how to F (or any proposition which entails that one knows how to F). The claim made by the insufficiency objection then is best interpreted as something like this: 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.

<sup>15</sup> One might think that the following is an example of such a case: Suppose Mary is a contestant in a TV game show. The first question is ‘What is the capital of New Zealand?’ and Mary knows the answer. However, Mary was not listening when she was told the rules of this game and so she is not sure what she is meant to do. Mary knows that she has to shout the answer out loud some 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, the right thing to do at that point was to shout the answer 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. However, I think this is a mistaken diagnosis of the case. 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.

---

<sup>16</sup> The use of “neo-Ryleanism” originates with Bengson and Moffett (2007, fn. 25). One reason for using this term, rather than simply “Ryleanism”, is that it is not entirely clear that Ryle needed to be committed to this view. What is clear is that Ryle identified knowing how to F with the possession of a complex of dispositions. S&W (2001, p. 411) attribute both the disposition view and neo-Ryleanism to Ryle, as they claim 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 two 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 is why S&W take the counterexamples they offer to (i) to be counterexamples to Ryle’s account of knowledge-how. Brian Weatherson, on his blog *Thoughts Arguments and Rants* (see: <http://tar.weatherson.org/2006/07/22/ryle-on-knowing-how/#comments>), argues that such counterexamples do not apply to Ryle on the grounds that he is only committed to (ii) and not (i). I think Ryle does sometimes assume the truth of (i) when arguing against intellectualism (see the discussion above of his schoolboy and chess-player examples). But, with respect to simply his own account of knowledge-how, it seems to me that Ryle would lose little if, in response to the standard counterexamples to (i), he were to simply reject (i) whilst retaining (ii).

<sup>17</sup> For example, S&W present the following case: ‘[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).’ Cases like this appear to be counterexamples to neo-Ryleanism because they are cases where, intuitively, someone knows how to F even though they do not possess the ability to F. For putative counterexamples to the claim that possessing the ability to F is sufficient for knowing how to F see Snowdon’s man in a room case (2004, p. 11) and Bengson and Moffett’s Irina case (2007, p. 46).

<sup>18</sup> See S&W (2001: 442–43) and for discussion Cath (2009) and (2011).

<sup>19</sup> Suppose that Ari Fs and in so doing he applies his knowledge how to F. From (R\*) it follows that: Ari knows that p1 (for some proposition p1 concerning a way to F), Ari applies his knowledge that p1, and in so doing Ari applies knowledge how to apply his knowledge that p1. But then we can reapply (R\*) to conclude that: Ari knows that p2 (for some proposition p2 concerning a way to apply one’s knowledge that knowledge that p1), Ari applies his knowledge that p2, and in so doing he applies his knowledge how to apply his knowledge that p2, and so on ad infinitum.

<sup>20</sup> See Cath (2011) for what I take to be three of the most powerful arguments against intellectualism. This paper is based on material in my PhD thesis. I would like to thank Daniel Stoljar and David Chalmers, my supervisors at the Australian National University, for all of their guidance and support. Martin Davies was also my supervisor for a short time before he left the ANU for Oxford, and I would like to thank him for many stimulating hours of discussion connected to my earliest thoughts on the issues discussed here. I presented ancestors of this paper to audiences at the ANU, UNC Chapel Hill, the 2005 Meeting of the North Carolina Philosophical Society, and the 2005 Australasian Association of Philosophy Meeting. I would like to thank all of those in attendance for their helpful feedback. Finally, I would like to thank Jason Stanley for his support and encouragement of my work on these issues.