Episteme, 15, 1 (2018) 119–140 © Cambridge University Press
doi:10.1017/epi.2016.45

SELF-REGULATION AND KNOWLEDGE HOW

BENJAMIN ELZINGA
be72@georgetown.edu

ABSTRACT

In the 1940s, Gilbert Ryle argued for anti-intellectualism about know how. More recently, new intellectualists have challenged the canonical status of Ryle's arguments, and in the ensuing debate Ryleans appear to be on their back foot. However, contributors on both sides of the debate tend to ignore or misconstrue Ryle's own positive account of know how. In this paper, I develop two aspects of Ryle's positive account that have been overlooked. For Ryle, S knows how to Φ iff (1) S is able to reliably live up to the norms for Φ-ing, and (2) S Φ's responsibly. In the first half of the paper, I argue that the two conditions rule out the various counter-examples to the simple ability view attributed to Ryle. In the second half of the paper, I argue that Ryle's second condition provides us with an account of warrant related to know how.

Most contemporary epistemologists recognize a distinction between knowing how and knowing that. The sentence "James *knows how* to ride a bike" suggests that he is able to ride a bike. The sentence "James *knows that* the sky is blue" does not obviously suggest a corresponding ability, but a relationship between James and a true proposition. If we take these suggestions at face value, knowledge how and knowledge that turn out to be distinct and independent kinds of knowledge. To have knowledge that is to believe a true proposition, and to have knowledge how is to have an ability. This reflects the received *anti-intellectualist* view of know how originally defended by Gilbert Ryle in the 1940s. Ryle argued that anti-intellectualism not only appeals to common sense but also that the alternative *intellectualism*, the view that knowing how either reduces to or at least requires prior propositional knowledge, leads to a vicious infinite regress and therefore must be rejected (Ryle 2009a).

The canonical status of Ryle's regress argument has recently come up for critical review. Jason Stanley and Timothy Williamson have argued that Ryle assumed an implausible conception of propositional guidance and so failed to undermine intellectualism properly understood (Stanley and Williamson 2011; Stanley 2011). This has opened the door to new intellectualist accounts of know how, but doesn't tell conclusively in favor of intellectualism. If Ryleans can specify conditions invoking an agent's abilities that are necessary and sufficient for knowing how, anti-intellectualism will still be the presumptive favorite (Fantl 2008). Unfortunately, new intellectualists also put forward numerous counter-examples which show that having an ability is neither necessary nor sufficient for knowing how, and in the ensuing debate Ryleans appear to be on their back foot.

Contributors on both sides of the recent debate, however, tend to relegate Ryle's own positive account of know how to the background. Jennifer Hornsby goes so far as to claim that Ryle didn't even have a positive account of knowing how (2011). Others suggest in passing that his view was actually quite sophisticated, but most treat him as subscribing to

a *simple ability view* subject to glaring counter-example.[1] On this view, S knows how to Φ iff S has an ability to Φ. In this paper, I develop two aspects of Ryle's positive account of knowing how that have been overlooked. For Ryle, S knows-how to Φ iff

(1) S is able to reliably live up to the normative standards governing Φ-ing, and
(2) S Φ's responsibly.

In the first half of the paper, I argue that these conditions rule out the various counter-examples to the simple ability view. Ryle didn't only have a positive account of know how, his view is both sophisticated and relevant to contemporary interests.

My goal, however, is not primarily exegetical. My goal is to respond to recent worries about anti-intellectualism using Rylean resources. In the second half of the paper, I argue that the particular promise of Ryle's view is that it allows us to clarify the relationship between know how and warrant. To have know how for Ryle is not merely to have an ability to perform well but also to perform responsibly. If we take Katherine Hawley's (2003) challenge to develop a sense of warrant related to know how seriously, Ryle's second condition suggests a solution. Roughly put, warranted performance is responsible performance. To make good on the promise of this suggestion, we need to understand the nature of responsible performance and its relationship to warrant. Ryle's account of responsibility as *self-regulation* provides this understanding. Self-regulation is equivalent for Ryle to self-teaching, and so to understand Ryle's account of warrant we will have to wade into his theory of education.

1. RYLE'S POSITIVE ACCOUNT OF KNOW HOW

Ryle lays out the two conditions in his initial characterization of know how in *The Concept of Mind*. Following a preliminary comparison of knowledge how and knowledge that ascriptions, he writes:

> 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 (Ryle 2009a: 17).

In saying that they "tend to perform them well", Ryle claims that the knower's performances reliably live up to normative standards. This is the first condition on knowing how. Ryle continues:

> 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' (Ryle 2009a: 17).

One can satisfy the first condition and be merely well-regulated. This provides reason for taking Ryle's first condition to make important use of the concept of reliability. A

1 See especially Stanley and Williamson (2011) and Stanley (2011). Fantl (2008) suggests Ryle had a sophisticated view, but doesn't develop it.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

well-regulated seal doesn't only sometimes perform well, and a well-regulated clock doesn't only occasionally keep good time. If they are well-regulated, they reliably perform well. Ryle also flags that the first condition is insufficient for knowing how by itself. Although they perform reliably well, the clock and seal lack know how. Ryle appears to treat "know how" as an honorific term:

> 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* (Ryle 2009a: 17, emphasis added).

Someone who knows how to perform a task not only reliably performs well, they are also responsible for their performances. Therefore, S knows how to Φ iff

(1) S is able to reliably live up to the normative standards governing Φ-ing, and
(2) S Φ's responsibly.

Ryle claims that the responsible performer is someone who regulates their performances. So we can say that to have know how is to be both well-regulated, given (1), and self-regulated, given (2). Even before unpacking the two conditions, it should be obvious that Ryle maintained a sophisticated ability view of know how. It must be said, however, that Ryle's comments on (1) and (2) are sparse and indirect, and they will have to be creatively pieced together.

2. THE RELIABILITY CONDITION

Ryle doesn't have much to say about the reliability condition except that it is invoked when we describe someone as knowing how to play chess, argue or make jokes. To make the condition intuitive, consider the conditions under which we'd say that someone knows how to move a knight in a game of chess. We don't say this of someone who, for example, consistently fails to move a knight in the correct L-shaped pattern or moves the piece into a space occupied by one of their other pieces. We do, however, ascribe such know how to someone who moves the piece correctly or in accordance with the norms – at least so long as they do so reliably and responsibly (but I'll ignore the latter for now).

The reliability condition rules out fluke performances as exercises of know how. A squirrel may stumble across a chess board in the park and move a knight to an unoccupied square two squares forward and one to the right, but we wouldn't on that basis say it knows how to move the knight. Consider instead the child who proudly demonstrates her burgeoning chess prowess by moving the knight consistently in accordance with the norms throughout the game. Unlike the squirrel, she moves the piece in the right way reliably, and partly on this basis we attribute know how to her. Reliability, however, should not be confused with perfect conformity. A seasoned chess player may know how to move a knight, but nevertheless, due to inattention or clumsiness, move the piece incorrectly on some occasion. Relatedly, the reliability condition only requires reliable performance across some relevant range of circumstances. It doesn't require counterfactually robust success. I can know how to drive without knowing how to drive a stick or without being able to handle particularly adverse weather conditions. This is important because

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

any account that requires counterfactually robust success would risk setting the standards too high (Hawley 2003: 20–2).

The reliability condition automatically rules out various counter-examples to the simple ability view discussed in the recent literature. Take David Carr's novitiate trampolinist, who miraculously performs a difficult somersault on his first try (Carr 1981: 53). Intuitively, the trampolinist doesn't know how to perform the somersault, but we can hardly deny that he was able to do it. After all, he did it. This appears to be a problem for a view that says merely having an ability is sufficient for knowing how, but for Ryle knowing how requires a reliable ability to live up to the norms of a task. Ask the trampolinist to perform the stunt again and he will likely fail. After all, his first success was only a matter of luck. Since he lacks a reliable ability, he lacks know how on Ryle's view.[2]

The first condition also helps with another case that is supposed to trouble ability accounts:

> Consider the difference between someone who can bench press a maximum of 100 pounds and someone who can bench-press 150 pounds. We may suppose that both employ the same technique; only brute strength makes the difference between them. Both are equally skilled, but clearly have different abilities. (Stanley and Williamson 2016: 15)

This would be a problem if reliably living up to the norms of pumping iron were simply a matter of reliably lifting the weights. Correct performance would consist in lifting the weight, and incorrect performance would consist in failing to lift the weight. But this doesn't fully capture what it means to say the task is norm-governed. First of all, to say that performances are norm-governed is not necessarily to say that there are well-defined rules written down somewhere or that the practitioners or active spectators can recite them on command or even put them in words. At the very least, it is to say that performances are subject to normative assessment (see Haugeland 1982). There is a correct and incorrect way to use your soup spoon, brush your teeth or inbound the ball in a basketball game. When you don't follow the norms governing these tasks, you open yourself up to criticism.

This suggests that when we think of tasks as norm-governed, the accompanying image should be that of a skilled critic evaluating someone's performance. Think of the kind of criticism a burly but incompetent weight-lifter opens himself up to when he muscles the barbell up back-arched, veins popping. A skilled observer will presumably be able to detect many failures in his performance related to his breathing, posture, attack and so on. On the view defended here, strength, speed and stamina are not themselves skills and to deploy them is not necessarily to exhibit know how. To exercise know how or skill is to reliably live up to the norms of the practice or, in other words, to reliably avoid criticism. To say that the two lifters are equally skilled and employ the same technique is therefore to say that they both reliably live up to the norms governing the task in the same way. What the one does reliably right, the other does reliably right as well. Since they do not differ in their ability to live up to the norms of the task, we can still say they are equally skilled on Ryle's view.

2 Some have responded by distinguishing stronger and weaker senses of ability. See Fantl (2008: 456). The reliability condition achieves the same effect while retaining a simple notion of ability.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

The claim that someone who knows how to Φ reliably avoids criticism is not meant to imply that they reliably perform optimally or beyond criticism. Notice that this is slightly different than saying that they don't have to perform with perfect reliability. The point is that someone can reliably perform more or less competently and still count as knowing how. Katherine Hawley argues that the standards for knowing how vary with the context of the knowledge attribution:

> a child might be said to know how to cook if she knows how to use the stove safely, whilst we would set standards higher (have a different task in mind) for an ordinary adult's "knowing how to cook," and set them higher still when discussing a chef in training. (Hawley 2003: 21)

The child who knows how to move the chess pieces but perhaps doesn't know about certain rules, like capturing en passant, still knows how to play chess. To say that the child knows how to play chess is to say, in part, that he reliably lives up to certain norms that we have in mind in the context of the attribution. The same goes for the professional, but in her case we invoke higher standards.[3]

3. RESPONSIBILITY AS RESILIENCE

The clock and seal don't count as knowing how to perform their respective tasks because they are not responsible for their performances. To perform responsibly is to regulate one's performances, but it's not immediately obvious what self-regulation involves. A full treatment of the concept will have to wait until section 6, but a preliminary sketch will help us rule out some of the simpler counter-examples to the ability view upfront.

In a later section of Chapter 2 called "Intelligent Capacities Versus Habits", Ryle distinguishes habitual behavior or rote performance from responsible performance. The abilities one exhibits in undertaking performances of either kind are "second natures or acquired dispositions", but habits are "mindless" or "automatic" and responsible behavior involves "care, vigilance and criticism" (Ryle 2009a: 30). Ryle writes:

> It is of the essence of merely habitual practices that one performance is a replica of its predecessors. It is of the essence of intelligent practices that one performance is modified by its predecessors. The agent is still learning. (Ryle 2009a: 30)

The mountaineer, for example, doesn't navigate a tricky route on autopilot. Colloquially, we would say that he "thinks about what he is doing", he exercises judgment, he attends to and learns from his mistakes. In Ryle's words he is "concomitantly walking and teaching himself how to walk in conditions of this sort" (Ryle 2009a: 30). Later we will see that according to Ryle the responsible performer doesn't merely perform reliably well, he also "applies criteria in preforming critically, this is, in trying to get things right" (Ryle 2009a: 17). In other words, he not only reliably lives up to the norms of the task, he also tries to live up to them.

3 As the example suggests, when I speak of living up to norms I have *constitutive norms* (norms that one must follow to count as playing a game or participating in a practice at all) rather than *regulative norms* (norms that specify, for example, how to perform well) in mind (see Haugeland 1998: 318–22).

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

At least part of what Ryle has in mind is here captured by the concept of *resilience*. John Haugeland's (1998) comments on resilience, a central component of his own deeply Rylean account of know how, are instructive. Haugeland describes resilience as:

> a kind of perseverance born simultaneously of adaptability and self-assurance. I have in mind, as a paradigm of resilience, an expert who "knows full well" that he or she can do something – and so is not turned aside or discouraged at the first, or even the second, sign of recalcitrance. Adjust a bit here, try that a little longer, don't fall for every semblance of trouble: these are the stuff of resilience as we admire it in physicians and mechanics, scientists and school teachers. (Haugeland 1998: 322)

Haugeland does not explain why he adds resilience as a condition on knowing how, but the reason is not hard to discern. The variation in the circumstances of application for any skill always outstrip the conditions under which it was learned and subsequently exercised up to the point of novel performance. To know how to perform a task is to be able to perform it not only in familiar circumstances but to be prepared to adapt to novel ones as well.

Haugeland points out that we are all skillful experts at lots of things (Haugeland 1998: 322). If I know how to describe my friend, I persist and adapt when you don't seem to get the right impression of her at first. If I know how to tie my shoes, then I'll be able to do it in the cold when my hands are shaking. To return to Ryle's example: if the mountaineer knows how to climb mountains, he is prepared to try to scale cliff faces he has never encountered before, make his way through novel trails, and, in short, adapt to new conditions, new perils and so on. This is not to say that he can conquer every terrain. Just as reliability should not be conflated with perfect conformity, resilience should not be conflated with unfailing adaptability. I may not be able to successfully describe my friend to you when you are being particularly thick-headed or tie my shoes while being attacked by a flock of speckled pigeons, but I still may have (more or less) resilient know how with respect to those tasks. The resilience condition does not rule out routine performances as expressions of know how either. The resilience condition specifies what someone does when they encounter resistance, problems or breakdown situations, not how they respond when things are going well.

It's worth stressing that the resilience condition does not imply that knowing how requires counterfactually robust success. If it did, it would risk setting the standards for know how too high. I take it for granted, following Hawley (2003), that one can know how to Φ only in a limited range of circumstances, and that when we say that S knows how to Φ features of the context determine what the relevant circumstances are. To be resilient is to try to live up to the norms of the task, and this implies trying to adapt in novel circumstances (though I don't take it to require unrelenting perseverance). Often, if one is resilient, one will adapt to new circumstances, but often one will try and fail. The resilience condition assures us that rote performance or mere habitual behavior doesn't count as an expression of know how, but unfailing adaptability or counterfactually robust success is not required to rule out rote performance. The idea of trying to perform and trying to adapt is enough.

The resilience condition helps with another case due to David Carr, as well as a variety of similar cases that have been discussed in the literature.[4] A dancer performs a choreographed routine that just happens to resemble a semaphore version of Gray's "Elegy".

4 See, e.g., Hawley (2003: 27) and Fantl (2008: 455–9).

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

According to Carr, the dancer is able to perform a semaphore version of the poem, but doesn't know how to perform it.[5] Notice that Ryle's first condition on knowing how doesn't obviously rule this case out. The dancer lives up to the norms of the semaphore performance whenever she performs her intended dance, so there is a sense in which she is able to reliably live up to the norms of the performance. The reliability condition rules out the trampolinist case, and what the trampolinist lacks in ability is clearly different from what the dancer lacks. On the other hand, the resilience condition implies that when someone exhibits know how in undertaking a task, they try to perform the task according to the norms governing it. In Ryle's words, the responsible performer "tr[ies] to get things right" (Ryle 2009a: 17). The dancer may be able to perform the semaphore poem, but they don't try to live up to the norms governing that particular performance. They merely happen to live up to them while trying to live up to the norms of the routine they intended to perform. Cases like this one illustrate how know how attributions create opaque referential contexts. One can know how to Φ and Φ-ing may be coextensional with Ψ-ing, but it doesn't follow one knows how to Ψ (Hawley 2003: 26). The resilience condition ensures that this feature of know how attributions is preserved.

For expository purposes, I will assume for now that responsible performance is performance that exhibits resilience and treat the responsibility condition as a resilience condition. This assumption will be challenged by the problem of warrant in section 5, and Ryle's full view of self-regulation will have to be developed in response. Later it will be shown that resilience is nevertheless a part of self-regulation and that the revised view is therefore immune to cases like the trampolinist case as well.

4. FIRST AND SECOND-ORDER KNOW HOW

The reliability and resilience conditions handle cases that show that having an ability is not sufficient for knowing how. What about cases that show that having an ability isn't necessary? Consider the following pair of cases:

> 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. But she has lost her ability to do so. (Stanley and Williamson 2011: 416)

If the ski instructor and pianist both know how to perform their respective tasks, then having an ability, even a reliable and resilient ability, is apparently not necessary for knowing how. After all, they lack the relevant abilities.

Several recent commenters point out that Ryle's view is not obviously subject to counter-examples of this sort (e.g., Fantl 2008: 468). Recall that Ryle thinks that "[t]o be intelligent is not merely to satisfy criteria, but to apply them". The clock satisfies criteria; the responsible agent applies criteria. Someone who applies criteria not only performs well but also discriminates correct from incorrect performance. Ryle stresses the

5 Some anti-intellectualists might be inclined to say that the dancer knows how to perform the semaphore poem, but they don't *know that* they know how to do it. I would argue that this response fails to take seriously Hawley's injunction to develop an account of know how that rules out epistemic luck. See section 5.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

importance the latter capacity throughout the second chapter of *Concept*. Consider the following passage:

> If I am competent to judge your performance, then in witnessing it I am on the alert to detect mistakes and muddles in it, but so are you in executing it; I am ready to notice the advantages you might take of pieces of luck, but so are you … The intelligent performer operates critically, the intelligent spectator follows critically. Roughly, execution and understanding are merely different exercises of knowledge of the tricks of the same trade. (Ryle 2009a: 42)

For Ryle, Φ-ing is not the only way to exhibit one's knowledge how to Φ. It is also manifest in the skillful assessment of performances. There are a number of ways one could take this suggestion, but again I find the most promising approach in Haugeland's positive account of know how. In what follows, I develop a distinction between *first-order know how* and *second-order know how*, based on a similar distinction made by Haugeland (1998: 322–5). Once the distinction is on the table, I will use it to address the counter-examples.

First-order know how or skill is a reliable and resilient ability to perform a task, participate in a practice or activity. Specifically, first-order skills are the reliable and resilient abilities to recognize and cope with phenomena in accordance with governing norms. In solo engagements these include, for example, the abilities to recognize aspects of the material world and manage, predict or otherwise react to their behavior. These can include the skills required to hypothesize, imagine alternatives, plan, strategize as well as the abilities required to test a hypothesis or execute plan, the ability to improvise and so on. In cooperative engagements, these include also the abilities to recognize, predict and interpret others' performances and coordinate one's performances with theirs. In short, the category includes any and all (reliable and resilient) abilities implicated in performing a task as opposed to abilities implicated in assessing the performance of a task.

Second-order know how or skill is a reliable and resilient ability to tell whether phenomena are in accord with the normative standards governing them. They are evaluative skills or skills of assessment. "Phenomena" includes agent's performances, material features of the world and the paraphernalia related to the task. For example, someone with the second-order know how required to play a passable game of chess can distinguish legal from illegal moves (made by her hand or another's) and can also tell when the game paraphernalia themselves are well behaved. That is, she can tell whether the objects are suitable to their roles within the task. A sparrow makes a poor king because the moment you put it on a square it wanders off. For a piece to play the role of king, it must stay in its square unless it is moved by a player and it must be amendable to the player's legal moves. Someone with the second-order skills required to play chess can tell when both players and pieces are well-behaved. The reliable and resilient ability to recognize *performances* as in accord with governing norms (as opposed to other phenomena) is most important for the following discussion. When I speak of second-order skills, this is primarily what I'll have in mind.

Few epistemologists explicitly recognize the importance of evaluative skills in the recent debate, so a few additional comments are necessary. First, second-order skills are skills of appraising norm-governed phenomena, but to exercise second-order skill is to engage in norm-governed activity as well. It is the norm-governed activity of assessing norm-governed activity, and to have this skill is to know how to assess know how. Relatedly,

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

the norms governing the deployment of second-order skills are different than the norms governing the activity being appraised. Recognizing a chess move as permissible is different from making a permissible move, and the standards governing the appraisals are different than the standards governing the moves. Second, there can be third-order skills, skills of telling whether second-order skills are exercised according to the norms governing them, just as there can be fourth-order and fifth-order skills and so on. Plausibly, the possession of skills of a higher order is parasitic on the possession of second-order skills, and they are presumably difficult to deploy given constraints on working memory. Third, notice that as defined, the possession of second-order skills doesn't require the ability to describe or codify the norms of the relevant practice. One can have the second-order skills required to play a game of chess (or spectate in the sense of the term that requires knowing what's going on) and indeed be able to sanction others' performances without being able to state the norms of the game.

Finally, second-order skills are often deployed as *recognative skills*. A recognative skill is a reliable and resilient ability to give uptake to aspects of the world. It seems natural to describe receptive encounters with the world in terms of skill or know how (Kukla 2006). The surfer possesses a *skill* of recognizing good waves that I don't have. The socially attentive agent *knows how* to recognize sexist behavior. Cognitive scientists routinely couch their research into specialized recognative abilities in the language of "perceptual skill". Think of a radiologist's ability to read an MRI, the forensic scientist's ability to detect signs of a struggle, the bird watcher's ability to recognize various species, or the chick sexer's ability to distinguish cockerel from pullet. In each case, the expert is reliably and resiliently responsive to features of the world. Radiologists know how to detect features of the image – distinguish noise from signal, normality from abnormality – and make perceptual judgments about their significance. This suggests that what is common to recognative skills is that they involve detecting features of the world and arriving at non-inferential, perceptual judgments (see McDowell 2008). When second-order skills are deployed as recognative skills, the sort of judgment arrived at is a normative judgment.[6] In response to features of a performance, you non-inferentially judge whether it was executed in accordance with the norms. Someone who knows how to do this does it reliably and resiliently well.

The ski instructor knows how to perform the complex stunt without being able to do it. As the tired saying goes, "those who can't do, teach", to which we might add "but those who teach, know how to do." But the saying is misleading when applied to this case. Describing him as a *ski instructor* suggests that he can articulate his knowledge, but merely saying that he knows how to perform the stunt does not. To teach someone how to do something is to exercise a bit of know how itself, and not all experts are skilled teachers. Imagine instead an inarticulate skier who is no longer able to perform. He may make a competent judge without making a competent teacher or color commentator for televised ski jump competitions. Although he cannot articulate his knowledge or pass it

6 I don't identify second-order skills with normative recognative skills because I take second-order skills to have two modes of expression. Either one non-inferentially judges a performance to be correct or incorrect or one's normative judgment has an inferential basis (see McDowell (2010: 141) on the difference between inferential and non-inferential judgments). To deploy second-order skills is to make evaluative judgments in either case, but it doesn't matter how you make them so long as you make them reliably and resiliently.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

on, it still sounds right to say that he knows how to perform the trick. This suggests that the minimal requirement for knowing how to perform without in fact being able to perform is being able to recognize when someone lives up to the norms of a given task and when they fail to live up to them. In my view, when we say that the ski instructor and the pianist know how to perform their tasks, we attribute second-order skill to them.[7] Since second-order skills are compatible with anti-intellectualism – indeed, to have one is to have a reliable and resilient ability – these cases pose no problem to the anti-intellectualist.

One might object along the following lines. Suppose someone acquired first-order piano playing skills without ever gaining the complementary second-order skill. When they suffer some tragedy and have their arms amputated, don't they still in some sense know how to play the piano? Isn't what matters that they still have some residual connection to first order-skills they are now unable to deploy? First of all, from the armchair I find this scenario psychologically implausible. In section 8, I argue that human beings simply don't tend to acquire first-order skill without acquiring some second-order skill. Second, whether or not the scenario is *conceptually possible* is an open question. Third, I am simply inclined to say that the individual in question does not know how to play the piano. Ask them to play, and (unfortunately) they cannot. Ask them for advice, they have none. Ask them to judge a recital, and they are no more help than the causal listener. Finally, if there is any residual intuition that they know how to play the piano, Alva Noë's distinction between having an ability to Φ and being able to Φ should help. The amputee pianist would have an ability to Φ, but tragically would be unable to Φ (Noë 2005: 283). Though I lack the space to elaborate, this move is perfectly compatible with my view.

The reliability condition and responsibility condition, the latter provisionally understood as a resilience condition, along with the distinction between first and second-order know how avoid the counter-examples that plague the simple ability view. Although the sophisticated ability account described so far avoids the counter-examples to the simple ability view, the problem of warrant is still outstanding, and it will occupy me for the remainder of this paper.

5. THE PROBLEM OF WARRANT

The problem of warrant is one of the deepest of the recent challenges to anti-intellectualism because it applies to sophisticated ability accounts. Katherine Hawley (2003) introduced the need for a concept of warrant related to know how into the recent discussion. According to Hawley, just as mere true belief is not enough for propositional knowledge, mere successful action is not enough for knowing how (2003: 19). "Warrant" is understood in Plantinga's neutral sense as picking out whatever turns true belief into

7 To be clear, my view is that know how ascriptions ordinarily attribute both first and second-order skills. When someone tells you that Sue knows how to play the piano, it's reasonable to expect that she has the first order-skills and the second-order skills. It's only when we know the pianist is not able to play the piano that we don't take the attribution to imply the possession of first-order skills. In these cases, I would argue that the unqualified attribution of knowledge ("she knows how to play piano") is shorthand for a qualified attribution ("she has the relevant second-order skills"). Our intuition that the pianist has know how is thereby explained by the fact that they have second-order know how.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

knowledge or, in the case of know how, turns successful action into knowing how (Hawley 2003: 28). Another way of putting the idea is that warrant is whatever rules out cases involving epistemic luck from counting as cases of knowledge. Although Hawley does not claim that her challenge poses a special problem for anti-intellectualists, the burden appears to fall more heavily on anti-intellectualists' shoulders. One might suspect, and I would suggest that many new intellectualists do, that any ability account will fail to rule out counter-examples involving epistemic luck. Notice that most of the objections to the simple ability view had to do with the problem of lucky success, and so will the further objections that motivate a concept of warrant to be discussed momentarily. The intellectualist might argue that if ability accounts can't rule out epistemic luck, then they can't do justice to know how as a species of knowledge and so must be rejected. In what follows, I argue that the anti-intellectualist is on just as firm a footing with regard to the problem of epistemic luck as the intellectualist.

First we need some familiarity with the kind of cases that motivate the need for a concept of warrant and how they apply to the current account. Hawley offers the following case:

> Consider Shelley, who attempts to make a cake by throwing together the ingredients she discovers on opening the kitchen cabinet. Luckily for her, the cupboard contains flour, sugar, butter, and eggs, and she makes a passable cake. Yet it seems that Shelley does not know how to make a passable cake … If she notices her success, and recalls her method, she may come to know how to make a cake. But the mere fact of her lucky success does not initially qualify her as knowing how to make a cake. (Hawley 2003: 27)

Shelley has a reliable ability. Presumably, she won't successfully bake a cake in many other circumstances, but recall that the reliability condition does not entail a counterfactually robust ability to live up to the norms of a task. It just describes the knower as reliable across some relevant range of circumstances. Shelley is able to successfully bake a cake when, for example, she happens upon the right ingredients or when a more experienced baker scaffolds her performance. That is enough for having a reliable ability.

Shelley also has a resilient ability. To say that someone is resilient is to say that they try to live up the norms of the practice, and this is taken to imply that they will vary their approach in response to recalcitrance. Often this means they will be able to perform across a variety of circumstances, but it's no guarantee. It's perfectly coherent to imagine that Shelley tries to get the norms of cake making right and that she is disposed to vary her approach in response to negative feedback, but that she nevertheless lacks a counterfactually robust ability. It appears then that Shelley has a reliable and resilient ability to make cakes. Intuitively, however, she lacks know how. Therefore, Ryle's two conditions, on their present interpretation, are insufficient for knowing how.[8]

Notice that the claim is not just that Shelley doesn't know how to bake a cake, but that Shelley doesn't know how to bake a cake in these specific circumstances. I am inclined to agree with Hawley on this point. Shelley lacks something that an experienced cake baker has, and describing Shelley as knowing how to bake a cake even in these circumstances

8 Why not simply strengthen the two conditions so that they imply counterfactually robust success? Perhaps this move could work, but I suspect it would set the standards for knowing how too high. Moreover, the solution I offer does a better job of addressing the intuitive problem with the Shelley case.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

doesn't do justice to the difference between novice and seasoned expert. The intuitive difference is that while Shelley is only able to make a cake because she happens to finds herself with the requisite ingredients, the skilled cake maker is able to make a cake because she has a history of cake making to draw on. In other words, she has *experience* in the thick classical pragmatist sense of the word exemplified by the job advisement that specifies "two years' experience required". Even if she finds herself in Shelley's circumstances, I am still inclined to say that her successful performance would not be merely a result of luck. It is a result of her experience.[9]

Ryle's theory of self-regulation can help us avoid the Shelley problem because it is at its core a theory of skill acquisition. At the same time, exploring Ryle's theory of self-regulation gives me an opportunity to make good on the claim that resilience is a form of self-regulation.

6. RYLE'S POSITIVE ACCOUNT CONTINUED

Immediately after equating responsibility with self-regulation in his initial characterization of know how, Ryle offers a gloss on self-regulation:

> A person's performance is described as careful or skillful, if in his operations he is ready to detect 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, this is, in trying to get things right. (Ryle 2009a: 17)

Think of how clocks and seals come to be well-regulated. They are well-regulated because they are each in the care of an agent, who is able detect and correct poor performance. Perhaps Ryle's idea is that in the case of the self-regulation, the agents who apply criteria and the subjects to whom the criteria are applied are one and the same.[10] The self-regulator would be someone who is aware of the standards that govern the task and monitors their performances for errors, adjusting as needed. This suggests that self-regulation is a form of what Peter Railton calls "normative guidance" (2006).

This interpretation captures the idea that the self-regulator is ready to detect and correct lapses and repeat successes, but it doesn't seem to capture the idea that they are prepared to "improve upon successes" or "profit from the examples of others". For Ryle, self-regulation isn't merely a matter of normative self-maintenance, it is a matter of learning what the norms of performance are. This makes sense when we take seriously the dual role of the self-regulator. The seal and clock don't need to apply criteria, to detect and

9 A sympathetic reader may still insist that the resilience condition rules out the Shelley case. If Shelley is resilient and she does encounter recalcitrance, won't she then make adjustments and won't her success be the result of her resilience rather than luck? I agree that if she were to encounter resistance and make adjustments, she would count as knowing how. However, in this case her success would be still the result of her learning experience, and so that wouldn't undermine my main point. One might also argue that Shelley lacks know how because she lacks second-order baking skills. While she does lack them, she still meets the conditions for possessing the first-order skills and intuitively lacks first-order know how. This reply would only work if having first-order skill required having the complementary second-order skill. I do not argue for this claim here.

10 This seems to be confirmed when Ryle describes the learner as someone who comes to "double the roles of instructor and pupil … to coach himself and heed his own coaching" (Ryle 2009a: 130).

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

correct lapses, that is the job of their caretakers. But if the self-regulator is both performer and self-trainer, they need to learn how to perform well and how to assess their own performances. In other words, they need to acquire both first and second-order skills. This means that the self-regulator is an autonomous performer: unlike subjects who are merely well-regulated, they are engaged in the practice of figuring out the norms for themselves.

7. SELF-REGULATION AS SELF-TEACHING

In *Concept*, Ryle makes these commitments clear with his comments on the difference between know how and mere habit. Know how is acquired through *training*, and habits are acquired through *drilling* (Ryle 2009a: 30–1). In a brief but illuminating essay called "Teaching and Training" Ryle expands on the concept of training in an explicitly pedagogical context (Ryle 2009b). In the article, he claims that the goal of a teacher is to teach their pupil to be self-teaching. Although he doesn't frame the article as an extension of his thoughts about responsibility in *Concept*, it's clear that he is picking up on the same themes and in fact returns to some of the examples. For this reason, I take "self-teaching" to be equivalent to "self-regulating" for Ryle.

If the goal of the teacher is to teach their students to be self-teaching, then to train someone is to guide their practices of self-regulation. Ryle says that it "ought to be one main business of a teacher precisely to get his pupils to advance beyond their instructions and to *discover new things for themselves*, that is, to get them to think for themselves" (Ryle 2009b: 466, emphasis added). For Ryle, education as a process of discovery contrasts with the view of education as a direct transfer of knowledge from teacher to student through drilling. We are often misled into this image of the teacher/student relationship because we often conflate education with developing habitual responses, inculcating rote or automatic behaviors, or getting our students to parrot facts back to us. The role of the teacher is not to drill in behavior that meets predesigned, fixed standards but to guide the student's exploration of a given subject matter: to facilitate or create an environment where a student can approach and solve problems suitable to their present knowledge and capabilities for themselves. The sort of teaching that interests Ryle is teaching that results in a pupil who makes his own moves, comes up with novel solutions to problem or develops new methods or strategies for investigating the world or refining some practice or skill. Education as a process of discovery is therefore indirect in the sense that passing on knowledge is environment mediated.

The distinction between direct and indirect knowledge transmission relates to a difference in the transmission of propositional knowledge and know how. Propositional knowledge can usually be acquired directly via testimony, but often know how can only be acquired indirectly through training.[11] This is not to deny that it is possible to pass on know how directly in some cases (nor is it to deny that it is possible to indirectly transmit propositional knowledge). Imagine the novice chess player who knows how to play a passable game of chess, but has never heard of capturing en passant. It seems reasonable to suppose that describing the rule to them would be sufficient for them to gain the relevant know how. However, Ryle was more interested in those cases where the learner does not

11 Testimony is direct in that given that the relevant conditions obtain, a speaker saying "P" is sufficient for a hearer to be in a position to know that P. See Hawley (2010) for a general account of testimony and know how.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

possess the relevant background skills and a mere description of a rule is insufficient for skill acquisition (Ryle 2009a: 46). In other words, he was more interested in coming to learn how to play games at all or how to play radically different games – in how to teach someone to play chess, not how to teach someone to capture en passant (compare Small 2014: 86–7). In such cases, know how is passed on indirectly. The teacher guides the student's interactions with the task environment and gets them to self-regulate their own performances.

8. THE STRUCTURE OF SELF-REGULATION

But what is self-regulation and what exactly is the teacher's role if the student is meant to teach themselves? I'll start with the former question. The basic structure of self-regulation or self-teaching is suggested by Ryle in a brief vignette about three young brothers who each attempt to canoe for the first time. Each has no prior experience, and they each try to teach themselves without any aid from a more knowledgeable instructor or fellow novice. The first brother is a self-regulator; the second two are not. Ryle says of the first brother:

> at first he tries a lot of random things, and nearly all of them end in immersion or collision; but he does after a time find out some ways of managing his craft … But it is because he had previously learned by practice, coaching and imitation the 'hows' of lots of other things such as tree-climbing, spelling and skating that he now takes it for granted that canoeing has its 'hows' as well, which similarly can be learned by practice, trial and error, and looking for ways of avoiding the repetition of errors. Here, as elsewhere, he has to study in order to improve; but this time he has nothing to study save his own unsuccesses and successes. (Ryle 2009b: 476)

The other two brothers each lack one of the two necessary components that make the first successful. Ryle says of the second brother:

> His more reckless and impatient brother, though full of go, just makes a dash at it, and then another quite different dash at it, and learns nothing or almost nothing from the failures which generally result, or even from the successes which sometimes just happened to result. He is not a self-trainer. (Ryle 2009b: 476)

The third-brother is not a self-teacher either:

> The third brother is uninterested, slow in the uptake, scared or idle. He never chances his arm. He tries nothing, and so initiates nothing either successfully or unsuccessfully. So he never learns to canoe. (Ryle 2009b: 476)

The example is designed to pinpoint the basic features of self-regulation that make the acquisition of know how possible. What's important about the third brother's failure is not that he is too timid or that he is uninterested per se, but that as a result he makes no trials, performs no experiments in managing his craft. The second brother is too reckless, inattentive and so on, but the essential point is that as a result he doesn't respond to feedback on his performances and adjust his performance accordingly. The first brother, on the other hand, has already learned the "hows" of other practices and so presumes that

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

canoeing has its own "hows". In other words, he has a sense that there are norms governing canoeing practice, he has the goal of living up to the norms, and he has a sense that the correct performance can be distinguished from incorrect performance by:

(a) varying his performances in a variety of circumstances (which the third brother fails to do) and
(b) responding to feedback on those performances (which the second brother fails to do).

We can first clarify this Rylean account of self-regulation by asking what it is to have a "sense" that there is a wrong and right way to perform a task and that these ways can be distinguished through performing operations (a) and (b). Ryle doesn't give us any explicit answers, but a broadly Rylean answer is not hard to imagine. We can say that at a minimum someone has such a sense if they are in fact disposed to vary their performances across a variety of circumstances and respond to feedback when they put themselves to some novel task. They *may* make the experimental procedure explicit to themselves, but it is reasonable to suppose that this is not required to count as having an (implicit) sense of there being a correct way of performing a task and how it is learned. If this is right, then the self-regulator will simply be, at a minimum, someone disposed to do (a) and (b), and these make up the core features of an agent that make acquiring know how possible.[12] Before addressing the role of the teacher, a few comments on (a) and (b) are in order.

(a) Making Trials

The self-regulator's performances can be described as trials or experiments, which test their hypotheses about how to engage successfully in the task. In saying this I do not mean to over-intellectualize the learning process. As I've already intimated, the learner may explicitly think of himself as engaged in the business of hypotheses testing and try to proceed deliberately and systematically, but he needn't proceed this way. He may simply be absorbed in the world in which he performs and make his trials as part of his attempt to achieve his goal. Although I don't wish to over-intellectualize self-regulation, I don't wish to under-intellectualize it either. I think it is at least misleading to describe self-regulation as a process of "random" trial and error. Our canoe boy will probably not initially try to manage the craft by banging his head against the hull of the boat or by stuffing pine cones in his ears and barking like a seal. Instead, the trials he makes will be influenced by his background knowledge and capacities.

(b) Responding to Feedback

To respond to feedback on one's performances is to treat them as trials. In making trial performances, the self-regulator tests his (implicit or explicit) hypotheses about how to perform some task and evaluates his performances in terms of whether or not they contribute to his goal. Those that tend to contribute positively are confirmed, although

12 Compare Daniel Dennett's claim that explanations of human behavior must at some level of analysis invoke some form of the law of effect (1975).

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

perhaps only tentatively, and those that don't are falsified, though perhaps not conclusively. In light of his evaluation, the self-regulator makes the relevant adjustments and thereby develops or refines his method. A strategy, method or "how" of a practice takes on the form of a conditional: in C, perform Φ. For example, the canoe boy will quickly realize that when you want to turn left, you paddle on the right-hand side of the boat. Strategies may be made more or less explicit by the performer during the learning process, but they needn't be, and often one finds that the strategies one deploys are difficult to make fully explicit after the fact. Nevertheless, methods are relatively stable and subject to revision. It is through continually undertaking performances and responding to feedback that the self-regulator develops, tests and refines their strategies.

One's background knowledge and capacities also determine in part how one responds to feedback. That this is so should be obvious given that background knowledge and capacities impact the formation of hypotheses and making an adjustment is nothing other than undertaking a new trial or testing a new hypothesis.[13] Describing performances as trials also helps us think about how first and second-order skills are acquired as a package. Trial performances encode the learner's current sense of how the practice ought to be carried out, and this sense is updated in light of feedback. Future trials then reflect one's updated take on how one ought to perform. The child tries to turn left by rowing on the left side of his craft. This expresses his intuitive sense of how the job is done. When the boat lists to the right, he recognizes his mistake, perhaps performs a new trial, and updates his sense of what one ought to do to turn left. I don't deny that first and second-order skills can come apart to some extent or that it is possible to gain second-order skills without gaining the corresponding first-order skills. A full exploration of how second-order skills are acquired is beyond the scope of the present work, but in the main first and second-order skills are gained as a package through practices of self-regulation.

9. THE TEACHER'S ROLE

Despite the fact that Ryle was interested in developing an account of education in terms of self-teaching, he has little to say about the role of the teacher (or "trainer"). The above elaboration of self-regulation suggests the following. The teacher's goal is to give the student a sense that there is a right and a wrong way to perform a task and that the right way can be distinguished through the procedure outlined (a) and (b). If to have such sense is to actually be disposed to vary one's performances and make adjustments in light of feedback, then the teacher's role is to take the pupil through the process of self-regulation. This is not accomplished by describing the learning process in the abstract, but by teaching them, or rather teaching them to teach themselves, some concrete skill.

To do this, the teacher will give demonstrations, provide explanations, encourage experimentation and reflection, and sanction incorrect performances, but this is just a patchwork of the elements of teaching of this sort. A general account is made available

13 The influence of background knowledge on future experimentation also means that improvisation can be described as either a process of putting methods together in new ways or revising or departing from them on the fly – both being a matter of experimentation in performance (as opposed to experimentation in training), which like any experimentation will be guided by prior knowledge and capacities.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

when we note that teaching is itself a skill. The teacher learns how to teach canoeing in the same way the pupil learns to canoe: by trying different approaches and attending to feedback. Of course she will likely learn how to teach from others, who guide her through the practice just as she will later guide her pupils through their lessons. But if she is taught to discover how to teach for herself, she is encouraged to try different approaches, to attend to what works and what doesn't and thereby develop her own way of teaching, and this is potentially always an ongoing process.

10. SELF-REGULATION AND RESILIENCE

So far I have been elaborating on Ryle's second condition in order to develop an account of warrant. S knows how to Φ only if (2) S Φs responsibly. Responsible performance consists of self-regulation, and self-regulation involves undertaking trial performances and responding to feedback. Before addressing the problem of warrant, two interpretations of (2) must be separated. On the first, S knows how to Φ only if S gained the reliable ability to live up to the normative standards governing Φ-ing through self-regulating their prior performances. On the second, S knows how to Φ only if S is disposed to self-regulate going forward. The former, backward-looking interpretation of (2) is more promising for an account of warrant, and I will return to it in the next section. The current formulation of (2) suggests the latter, forward-looking interpretation, but this apparently leads to an implausible view. To self-regulate is to engage in the learning process. If knowing how requires self-regulation going forward, then the Rylean account would apparently set the standards for knowing how too high.

For example, the case of the mountaineer – walking and concomitantly teaching himself to walk – invokes the image of the dedicated professional continually trying to improve. Contrast that image with that of a fourth canoe boy, who learns how to navigate the local waters and promptly stops the learning process. When he encounters feedback that tells him he is not performing well (like crashing into the riverbank in a swift current), he does little to adjust his future performance. Nevertheless, he is more or less competent, and it seems perfectly natural to say that he knows how to canoe. Speaking for myself, I find that I am often in the same boat as the fourth boy. Apparently, we must attribute know how to him on pain of making us wrong about most of what we take ourselves to know how to do, and a view of know how that rules out the hobbyist sets the standards too high.

Before turning to the backward-looking interpretation, I will try to salvage the forward-looking one. Notice that the backward-looking interpretation doesn't leave room for resilience. Resilience is a putative forward-looking condition on knowing how, and there is a clear sense in which resilience is an aspect of self-regulation. Someone who is resilient (a) varies their performance (b) in response to feedback that indicates they are performing poorly. Since resilience involves both aspects of self-regulation, we can say it is a form of self-regulation. However, resilience does not call for major adjustment in the face of recalcitrance – and rightly so, since that would apparently rule out the fourth canoer as knowing how to canoe.

It's not obvious, however, that self-regulation always requires major adjustment either. Self-regulation primarily shows up in what a performer does in response to poor performance. The self-regulator is someone who "tries to get things right" by making adjustments

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

when they go wrong. After all, successful performance doesn't typically call for any adjustment. So long as the agent can repeat their successes consistently, they can responsibly carry on in the same way. Nevertheless, throughout the course of applying one's skills, small variations in the circumstances of application will still call for improvisation of the sort resilience describes. Since this is a sort of self-regulation, we can say that someone who knows how to Φ regulates their performances going forward, and the forwards-looking interpretation is preserved.

One might worry that this causes a problem for the professional on the other end. We want to say that for the hobbyist suboptimal performance doesn't always call for adjustment, but for the dedicated professional or serious amateur it does. However, this is only because the standards for success are different for the two. The professional must engage more deeply in practices of self-regulation because he has taken on greater responsibility. He strives to meet higher standards. This might lead to another problem. If responsibility is a condition on knowing how, and the conditions for responsible performance vary from person to person, doesn't that mean the conditions for knowing how differ from person to person? Yes, but as Hawley points out, the standards for know how attributions vary with context anyway. Moreover, they vary in the same way as assignments of responsibility: in relation to standards for success.

11. SELF-REGULATION AND WARRANT

With the outlines of a theory of self-regulation in place, we can return to the problem of warrant. The accomplished baker's success is the result of experience. Shelley's success is the result of luck. Luckily successful action is incompatible with knowing how just as luckily true belief is incompatible with knowing that. This suggests that for someone's successful performances to be an expression of know how, their performances must be the result of experience. For Ryle, having know how requires self-regulating one's performances. I have argued that we should take this as placing a forward-looking condition on knowing how in order to capture Haugeland's idea of resilience and avoid the problems discussed in section 3. Here, I suggest that we take the condition to place backward-looking conditions on knowing how as well. All together, S knows how to Φ iff:

(1) S is able to reliably live up to the normative standards governing Φ-ing.
(2) S is a self-regulator along two dimensions:
    i S is disposed to self-regulate going forward, and
    ii S gained the reliable ability to live up to the normative standards governing Φ-ing through self-regulating their prior performances within the relevant practice.

To know how is not only to be able to reliably perform well but also to have warrant, understood in terms of (2.i) and (2.ii). This goes for both first and second-order skills, and so my earlier comments on both should be reinterpreted to reflect the revised interpretation of (2). Shelley does not have know how on this account, because her success is not the result of self-regulating her prior performances. In other words, she fails to meet condition (2.ii). I take it that warrant also requires (2.i) since successful performance in the face of recalcitrance that isn't the result of resilience would be the result of luck as well.

Since it requires regulating prior performances, the backward-looking condition (2.ii) may seem to imply that in order to know how to do something one must have already

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

done it. Surely, however, one can know how to do something without *ever* doing it. I have never (as far as I know) subtracted 178 from 234 and may never do so, but I'm quite sure I know how. However, (2.ii) specifically states that in order to have know how one must have regulated one's prior performances *within the relevant practice in which Φ-ing takes place*, not that one has regulated one's prior attempts at Φ-ing in particular. Remember that Ryle was interested in how we acquire the general competence that makes more specific performances possible – in learning chess rather than learning to capture en passant. I count as knowing how to perform the particular equation because I have self-regulated my prior performances in arithmetical practices generally, not because I have regulated my prior attempts to calculate that particular difference.

One might further worry that (2.ii) makes innate know how or know how acquired through means other than self-regulation impossible. Isn't it perfectly coherent to imagine someone who is born knowing how to do a somersault or solve subtraction problems? Or think of Neo, the protagonist of *The Matrix*, who asserts "I know kung fu!" after having the knowledge downloaded into his brain. Does the present account imply that Neo is wrong about his current epistemic state? There may be several viable responses here (e.g., one might distinguish weaker from stronger senses of know how and argue that Neo knows how to do kung fu only in a weaker sense), but I prefer to insist that innate or downloaded abilities don't count as know how. It seems to me that we already have a word for abilities that you find yourself with as opposed to ones you learn. To have an innate ability, for example, is to be born with a *knack*. To have a knack for something is precisely to possess an ability to perform in accordance with standards which is not an epistemic achievement (Bengson and Moffett 2011: 35–40). Know how is distinguished from knacks precisely because know how is acquired through responsible epistemic practices. So, innate or downloaded abilities are not cases of know how but of knacks or perhaps "mere abilities", and the potential counter-examples miss their mark.[14]

One final comment to ward off potential misunderstanding. The self-regulation account of warrant may appear to be especially useful for thinking about sophisticated skills or expert performance, but it is intended to apply to mundane know how as well. I know how to open a cardboard box, untangle electrical wires, and open a typical door. I can usually do all these quite effortlessly, but I nevertheless implicitly deploy methods for each task, methods that are not, for example, available to a toddler. This comparison is instructive. To gain know how within these domains, the toddler will have to regulate her performances while, for example, opening birthday presents, untangling headphones, and entering buildings. I suspect that the present account will appear to have narrow scope only if one loses sight of the fact that much of what one knows how to do, like much of one's propositional knowledge, was hard won during one's upbringing.

12. CONCLUSION

My goal in this essay was twofold. First, I distinguished Ryle's positive account from the simple ability view. Second, I used Ryle's comments on responsible performance to

14 Nevertheless, just as one can obtain a belief that only later becomes propositional knowledge when you gain warrant for it, an inborn knack or ability can later become know how if you subsequently regulate your performances.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

construct an account of warrant. There is certainly more to say about the concept of self-regulation, but I want to end with some reflections on the problem of epistemic luck.[15] My specific goal in the second half of the paper was to argue that the anti-intellectualist fares just as well as the intellectualist when it comes to the problem of warrant. As stated, however, the above account is subject to potential counter-example. Here is another case from Hawley:

> [C]onsider Susie, who likes to annoy Joe, and believes that she does so by smoking. In fact, Joe is annoyed by Susie's tapping on her cigarette box, which she does whenever she smokes. Susie would succeed in annoying Joe if she tried, but it seems that she doesn't know how to annoy Joe. (Hawley 2003: 27)

It seems perfectly possible that Susie gained her reliable ability to annoy Joe through self-regulating her prior performances, that she tries to live up to the norms for annoying Joe, and that she is prepared to self-regulate going forward. Intuitively, however, she does not know how to annoy Joe. This is a Gettier problem for the present account.

Is it avoidable? Just now I was assuming a causal interpretation of (2.ii). S knows how to Φ only if S's self-regulating their prior performances *caused* S to have a reliable ability to live up to the standards for Φ-ing. This appears to handle the Shelley case, but doesn't handle the Susie case. Although there may be a causal connection between Susie's self-regulation and her reliable ability, she still lacks know how. Perhaps we should say that S knows how to Φ only if she is able to reliably perform well *because* she self-regulates her performances, paralleling a certain move made by Linda Zagzebski to avoid Gettier problems for propositional knowledge (Zagzebski 1999). This appears to help. Susie does not have a reliable ability to annoy Joe *because* she self-regulates her performances. She has a reliable ability to annoy Joe because of luck.

Of course for this to work, we need an analysis of the "because" clause that doesn't reduce to a simple causal view. On this matter Zagzebski is largely silent, claiming the clause resists analysis. One may rightly worry that if there is no way to unpack the clause, the proposed solution amounts to nothing more than a restatement of the problem. Duncan Pritchard (2008) suggests we can analyse the "because" clause in terms of sensitivity conditions and perhaps this is a promising suggestion, but my point is not to wade into this particular debate nor is it to solve the apparent Gettier problem for knowing how. This is not the place to provide a detailed comparison, but what I want to suggest is that the Gettier problem as it applies to knowing how is the same Gettier problem

15 One might worry, for example, that self-regulation sneaks intellectualist commitments through the backdoor. The self-regulator is in the continual process of undertaking new experiments in performance. The trials one undertakes are in part determined by one's background knowledge, and this will arguably often include a lot of propositional knowledge. But according to the anti-intellectualist, know how neither reduces to *nor requires prior propositional knowledge* (Fantl 2008: 452). Doesn't this give ground to the intellectualist? I don't think so. First of all, I am inclined to say that whether know how requires propositional knowledge or not is an empirical question, and I take it as a strength of the view that it leaves this question open. One potential avenue to pursue for the anti-intellectualist is suggested by Seligman *et al.* (2013) and Railton (2014). It might be argued using these resources that the learning process is primarily driven by system-1 processes that implement forms of error-driven learning that apparently don't make use of propositional knowledge. Second, even Ryle accepts that performances exhibiting know how are occasionally guided by propositional knowledge, so admitting this should not undermine a commitment to anti-intellectualism (Ryle 2009a: 18).

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

that applies to knowing that. In both cases, the epistemologist's goal is to find a relationship between warrant and other knowledge conditions that rules out epistemic luck. So this is not a special problem for the anti-intellectualist. In this paper, I have articulated a sense of warrant related to knowing how in order to rule out cases of epistemic luck, but I'll leave solving the Gettier problem to one side.[16]

REFERENCES

Bengson, J. and Moffett, M. (eds) 2011. 'The State of Play.' In *Knowing How: Essays on Knowledge, Mind, and Action*, pp. 3–55. Oxford: Oxford University Press.

Carr, D. 1981. 'Knowledge in Practice.' *American Philosophical Quarterly*, 18: 53–61.

Dennett, D. 1975. 'Why the Law of Effect Will Not Go Away.' *Journal for the Theory of Social Behaviour*, 5: 169–88.

Fantl, J. 2008. 'Knowing-How and Knowing-That.' *Philosophy Compass*, 3: 451–70.

Haugeland, J. 1982. 'Heidegger on Being a Person.' *Noûs*, 16: 15.

—— 1998. *Having Thought: Essays in the Metaphysics of Mind*. Cambridge, MA: Harvard University Press.

Hawley, K. 2003. 'Success and Knowledge-How.' *American Philosophical Quarterly*, 40: 19–31.

—— 2010. 'Testimony and Knowing How.' *Studies in History and Philosophy of Science*, 41: 397–404.

Hornsby, J. 2011. 'Ryle's Knowing How and Knowing How to Act.' In J. Bengson and M. Moffett (eds), *Knowing How: Essays on Knowledge, Mind, and Action*, pp. 80–100. Oxford: Oxford University Press.

Kukla, R. 2006. 'Objectivity and Perspective in Empirical Knowledge.' *Episteme*, 3: 80–95.

McDowell, J. 2008. 'Avoiding the Myth of the Given.' In J. Lindgaard (ed.), *Experience, Norm, and Nature*, pp. 1–14. New York, NY: John Wiley & Sons.

—— 2010. 'Brandom on Observation.' In B. Weiss and J. Wanderer (eds), *Reading Brandom: On Making It Explicit*, pp. 129–44. New York, NY: Routledge.

Noe, A. 2005. 'Against Intellectualism.' *Analysis*, 65: 278–90.

Pritchard, D. 2008. 'Virtue Epistemology and Epistemic Luck, Revisited.' *Metaphilosophy*, 39: 66–88.

Railton, P. 2006. Normative Guidance'. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, pp. 3–34. Oxford: Oxford University Press.

—— 2014. 'The Affective Dog and its Rational Tale: Intuition and Attunement.' *Ethics*, 124: 813–59.

Ryle, G. 2009a. *The Concept of Mind. 60th Anniversary edition*. London: Routledge.

—— 2009b. *Collected Essays, 1929–1968. Vol. 2*. London: Routledge.

Seligman, M., Railton, P., Baumeister, R. and Sripada, C. 2013. 'Navigating into the Future or Driven by the Past.' *Perspectives on Psychological Science*, 8: 119–41.

Small, W. 2014. 'The Transmission of Skill.' *Philosophical Topics*, 42: 85–111.

Stanley, J. 2011. *Know How*. Oxford: Oxford University Press.

—— **and Williamson, T.** 2011. 'Knowing How.' *Journal of Philosophy*, 98: 411–44.

—— **and** —— 2016. 'Skill.' *Noûs*, 50(2). doi: 10.1111/nous.12144.

Zagzebski, L.T. 1999. 'What is Knowledge?' In J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*, pp. 92–116. Malden, MA: Blackwell Publishers.

16 I would like to thank Trip Glazer, Colin Hickey and Bryce Huebner as well as one anonymous reviewer at *Episteme* for their encouragement and comments on earlier drafts of this project.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press

BENJAMIN ELZINGA is a PhD candidate in philosophy at Georgetown University and is currently writing a dissertation on epistemic agency. His research interests include epistemology, philosophy of education, the problems of free will, and the philosophy of mind.

https://doi.org/10.1017/epi.2016.45 Published online by Cambridge University Press