

ARTICLE

Thinking What One is Doing: Knowledge-how, Methods, and Reliability

John Turman 

Stanford University, CA, USA

Email: john.michael.turman@gmail.com

(Received 13 August 2020; revised 18 February 2021; accepted 7 April 2021;
first published online 21 June 2021)

Abstract

There has been renewed interest over the last twenty years in Ryle's claims and arguments about knowledge-how. Elzinga (2018) and Löwenstein (2017) have both recently defended independent Ryle-inspired accounts of knowledge-how. In what follows, I will propose and defend an amendment to accounts of knowledge-how like those of Elzinga and Löwenstein. I argue that this amendment provides an additional needed distinction between the performance robustness provided by certain performance *methods* (or *styles*), and the robustness of an agent's ability to perform according to such methods (or styles). Additionally, the proposed amendment, if adopted, will make the amended views even more Rylean. I argue for this, in part, through original exegetical work on an under-discussed theme in Ryle's philosophy of mind: the relation between semi-hypothetical statements, methodological act-description, and knowledge-how.

Keywords: Knowledge-how; Ryle; Reliability; Semi-hypothetical; Intelligence; Methods

Introduction

Elzinga (2018) and Löwenstein (2017) have both recently defended independent Ryle-inspired accounts of knowing how to ϕ . In what follows, I will propose and defend an amendment to these accounts. I will begin by expositing certain essential elements of these accounts, and then propose the amendment. Next, I will argue that the proposed amendment provides for a powerful distinction that allows the amended views to escape certain objections that I will develop. Finally, I will argue that adopting the amendment makes the amended views even more Rylean.

Elzinga has recently defended an account of knowing how to ϕ according to which "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:
 - (a) S is disposed to self-regulate going forward, and
 - (b) 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." (Elzinga 2018: 136)

The normative standards in question vary with different ϕ s, and are the standards of evaluation according to which the ϕ ing at issue is to be judged, variously, as successful, correct, excellent, etc. (Elzinga 2018: 121). The ability to reliably live up to these standards is the ability to meet them in one's ϕ ing across a certain range of performance conditions determined by the context of the knowledge-how attribution (Elzinga 2018: 121–3).¹

Elzinga glosses self-regulation as a kind of self-teaching which involves, at least in part, trying out different approaches in one's ϕ ing and responding to the outcomes of such trials (Elzinga 2018: 132–4). To (a) self-regulate in this way going forward is to be resilient in one's ϕ ing, responding with flexibility to novel performance conditions, opportunities, and obstacles (Elzinga 2018: 123–4, 135). That (b) one's reliable ability has been gained through prior self-regulation helps to ensure that it is because of one's *experience* that one's ϕ ing meets the relevant standards (and in this way (b) is argued to provide a practical analogue of epistemic warrant) (Elzinga 2018: 136–8).

Elzinga's account of such self-regulation is at the same time an interpretation of the *responsibility* that Ryle claims is characteristic of knowledge-how manifesting performances (Elzinga 2018: 130–2). Ryle claims that a person who knows how to ϕ , unlike an accurate clock or a well-trained (non-human) animal, is said to be responsible for the fact that their ϕ ing meets certain standards (Ryle 2009 [1949]: 17). Elzinga's account of self-regulation draws much from Ryle's various claims and arguments about our concepts of knowledge-how manifesting performances, teaching, and learning.

Löwenstein has recently defended an account of knowing how to ϕ according to which knowing how to ϕ is having an ability to reliably meet normative standards in ϕ ing, where this reliability is explained by *normative guidance* – i.e. by the agent's exercising a capacity (though not necessarily during each instance of ϕ ing) to assess, in a certain way, their performances and options, and acting in light of these assessments (Löwenstein 2017: 122–3).

For Löwenstein, normative standards for ϕ ing are not just standards by which we evaluate token ϕ ings – they are criteria by which we *individuate* different activities in the first place (Löwenstein 2017: 19). A reliable ability to meet such standards is an ability to meet them under 'normal' performance conditions, which Löwenstein characterizes as conditions under which ϕ ing is not substantially impeded (Löwenstein 2017: 25–8). The reliability of the ability, the normality of the performance conditions, and knowing how to ϕ are all understood as gradable (Löwenstein 2017: 25–8).

For Löwenstein, assessing one's own performances and options in this way involves the deployment of concepts for act types, performance condition types, and degrees of performance quality, as well as indexical identifications of acts and performance conditions (Löwenstein 2017: 116–20). On Löwenstein's view, such correct assessments about one's own performances and options constitute propositional knowledge (Löwenstein 2017: 120–2).

1. A Hawlean strategy

Both Elzinga (2018) and Löwenstein (2017) incorporate elements of Hawley's (2003) account of knowing how to ϕ , and both accounts (Elzinga's explicitly, Löwenstein's less explicitly) follow a Hawlean strategy for ensuring that their respective diagnoses

¹I read Elzinga this way because, on p. 123, it is explained that the subject of the knowledge-how attribution must reliably meet fine-grained constitutive norms determined by the context of the attribution, while earlier on p. 121, reliability itself seems to be characterized in terms of those performance conditions a fine-grained constitutive norm can be met under.

will match our pre-theoretic intuitions about who knows how to φ (i.e. for ensuring the *extensional adequacy* of the views).² The strategy begins by distinguishing agents' abilities to meet standards reliably across a relevant range of performance conditions from the possibility of agents getting lucky and performing in a standard-meeting way as a result. Such reliable abilities form the foundation of these accounts, and provide a response to putative counterexamples involving fluke or lucky success.

The next step in the Hawlean strategy is to put additional conditions on knowing how to φ , conditions that, among other things, provide a response to additional putative counterexamples. For example, Elzinga's second self-regulation condition rules out cases like Hawley's baker, who doesn't know how to bake a cake, yet has the good fortune to be surrounded by all and only the correct ingredients for cake baking and happens to be disposed to throw them together in the correct way (and would thus seem to meet the first condition) (Elzinga 2018: 129–30). Löwenstein's normative guidance conditions provide for the distinction between the exercises of genuine knowledge-how and the performances of clocks and trained animals (who would also seem to meet the first condition) (Löwenstein 2017: 30–1). Hawley's warrant condition rules out cases of several types, including that of a smoker who knows that they annoy someone but is mistaken about why (but who would nevertheless seem to meet the first condition) (Hawley 2003: 28).

Hawlean strategy

- 1 An agent knows how to φ iff and only if they are able to reliably meet normative standards in φ ing across a range of relevant performance conditions, and
- 2 The agent meets some additional condition(s).

These views are most clearly distinguishable from each other by the specific characterization they each provide of (2) the additional conditions, and Elzinga and Löwenstein devote the majority of their respective expositions to articulating them. But right now I'd like to draw attention to the *first* condition, namely the condition that knowing how to φ is, on these accounts, partly a matter of being able to reliably φ (in a standard-meeting way) across a certain range of performance conditions. On Elzinga and Löwenstein's accounts, an agent knows how to φ iff there is a certain range of performance conditions such that, any time the world is such that these performance conditions are met, the agent can reliably φ in a standard-meeting way and certain additional conditions (i.e. self-regulation; normative guidance) are met (Löwenstein 2017: 25–9, 122–3; Elzinga 2018: 120–3, 136).

2. Conflating two types of conditions

But characterizing knowing how to φ in this way conflates the range of conditions a particular performance *method* (or *style*) equips an agent to cope with, and the range of conditions under which a particular agent can reliably perform according to that method (or style). Consider a baker who knows how to bake a cake. According to Hawley, Elzinga, and Löwenstein, the baker's knowing how to bake a cake is partly grounded in their ability to reliably bake a proper cake across a certain range of per-

²Elzinga (2018: esp. 120–30) explicitly follows a Hawlean strategy. See Löwenstein (2017: 25–9) for a discussion of performance robustness, and (2017: 158–66) for a rich discussion of puzzle cases from the literature in which abilities that are robust nevertheless do not amount to knowledge-how.

formance conditions.³ For Elzinga and Löwenstein, the ability must be robust in the right way – the relevant range of performance conditions must not be restricted to only those conditions under which the agent gets lucky somehow (Löwenstein 2017: 25–9; Elzinga 2018, 120–3). It is plausible that, for many agents who know how to bake a cake, this range of tolerable performance conditions will include some variation in things like the availability of different kinds of cake pans (e.g. shallow pans vs deeper pans), and local altitude and humidity (e.g. Kansas City vs Tucson). It is also plausible that for *some* agents this range will also include variation in things like the amount of coffee consumed before cake baking, and the presence or absence of certain psychological stressors.

But the conditions just mentioned are of *two different types*. Most competent bakers' ability to succeed with either a shallow pan or a somewhat deeper pan, in Kansas City or in Tucson, is grounded in the fact that they *check* their cakes with a skewer (or fork; toothpick; fingertip; etc.) to determine how far along they are in the baking process. (Deeper pans tend to require longer baking times; higher altitudes can require shorter baking times.) A cake is done when it is done, and if a baker checks their cake, they can cope with a certain amount of variation in baking time that can occur for a variety of reasons. Checking one's cake is part of a *method* of baking cakes.

On the other hand, if a competent baker can make it work after having four cups of coffee, or after having no coffee at all, this usually isn't grounded in their method of baking cakes. Rather, it is usually a condition of their ability to execute their cake making method, a condition that is grounded in their not being addicted to, or overly sensitive to, caffeine. There are innumerable conditions that a method cannot control for but must presuppose. To illustrate, a cake *recipe* cannot provide instructions, so to speak, all the way down – it is only against a background of basic competencies that any particular recipe is useful. Recipes don't tell you *how* to measure ingredients, but even if they did, they could only do so in terms of other competencies that are themselves presupposed.

A particular method of baking, once it has been mastered, will *itself* equip the agent using it to cope with certain kinds of variation in performance parameters. Nevertheless, just because someone bakes according to a method that equips them to cope with certain variations in pan shape, altitude, etc., doesn't mean that as long as the pan shape, altitude, etc. are within the tolerable limit they can reliably bake a cake. A particular agent may *also* need to have had at least two cups of coffee, or to have had no coffee whatsoever, or to be free of certain psychological stressors, to reliably execute their cake baking method. (For in times of profound distress, or when suffering caffeine withdrawal, some agents may tend to forget to add an ingredient or to check the cake.)

But to identify the set of conditions under which an agent can reliably ϕ in a standard-meeting way, even where these conditions are ones in which the agent meets certain additional constraints (e.g. self-regulation; normative guidance; warrant), is not yet to distinguish between the variation in these conditions that is provided by the method itself (e.g. shallow vs deeper pans, low altitude vs middle altitude), and the variation that represents the particular agent's ability to reliably execute that method (e.g. plenty of coffee vs no coffee; the presence or absence of certain psychological stressors; etc.).

3. An amendment

I'd like to propose an amendment to accounts like those of Elzinga and Löwenstein. After proposing this amendment, in the next section I'll argue that the distinction the amendment provides for should be a part of any satisfying ability-based account

³Hawley (2003: 19–22), Löwenstein (2017, 25–9) and Elzinga (2018: 120–3, 136).

of knowledge-how. As I will argue in the last section of the paper, this is especially true for accounts that also strive to function as interpretations of Ryle's own view about the nature of knowledge-how.

I'll start with the second condition. I propose that self-regulation, and normative guidance, be understood as characterizing a *method* (or *style*) that an agent is able to φ according to. For Elzinga's account, this would mean that an agent can φ sensitively to the results of previous trials (e.g. by means of a method that includes explicit planning based on previous results, or by means of less conscious behavioral tendencies), *and* in a way that simultaneously allows the current φ ing to itself function as a trial to inform future φ ings (e.g. with sufficient presence of mind for the φ ing and its results to be recallable later).⁴

For Löwenstein's account, this would mean, as per Löwenstein's characterization of intellectual guidance, that an agent is able to φ according to a method that involves conceptualizing the relevant act types or tokens, and the relevant performance condition types or tokens, and making judgements about the degree to which normative standards for φ ing can be met by such acts in such conditions, and acting on the basis of these judgements.⁵

Now for the first condition. I propose that we understand the range of relevant performance conditions as, more precisely, a range of relevant values in relevant performance *parameters* that the methods characterized by the second condition equip an agent to cope with. Knowing how to φ would then require that, so long as the agent is able to execute the method for φ ing described by the second condition, their meeting normative standards in φ ing will not be prevented by certain variation in the values of certain parameters – e.g. whether the baker has glass mixing bowls or metal bowls to work with; whether they are baking in Kansas City or in Tucson; etc.

Understood in this way, the range of relevant performance conditions will be equivalent to a set of possible worlds, but which set this is will depend on the answer to at least one philosophical question that is independent of the issues that I am pressing in this paper. Allow me to explain.

We can characterize the range of relevant performance conditions in a way that is halfway between the informal characterization given above and the precise characterization of a determinate set of possible worlds.⁶ Let a *performance condition for φ ing* be a way that the world can be, a way that is individuated on the basis of the opportunities for φ ing and challenges to φ ing that it presents.⁷ My proposal for the first condition then comes to this: Those performance conditions for φ ing under which an agent can successfully execute their method for φ ing, *and* under which, *if* they managed to meet normative standards for φ ing, they would be meeting these standards *because* they managed to execute their method, are *not* individuated on the basis of certain parameter values that we expect the method itself to equip them to cope with. In other words, for an agent who knows how to bake, having specifically glass (or, alternatively, specifically metal) mixing bowls to work with is *neither* required for their being able to execute their method (as having plenty of coffee might be), *nor* has the power to undermine their successfully meeting normative standards in baking by means of the method (or to bring about a Gettierized meeting of these standards).⁸

⁴cf. Elzinga (2018: 132–4).

⁵cf. Löwenstein (2017: Ch. 4, especially §§4.2–4.4).

⁶I am indebted to David Hills for his suggestions here.

⁷cf. Löwenstein (2017: 28–9) for discussion about characterizing normal performance conditions in terms of (the absence of) performance impediments, or in terms of Gibsonian affordances.

⁸However, the disjunction of the set of allowable values will partly individuate performance conditions for φ ing.

To further refine the set of relevant performance conditions into a determinate set of possible worlds, one must first answer the following independent, philosophically interesting question: To what extent does a (knowledge-how-constituting) competency (or skill, etc.) guarantee the meeting of normative standards in ϕ ing? If one thinks that to know how to ϕ in circumstances C one must have a competency that *guarantees* meeting normative standards for ϕ ing in C , then the set of relevant performance conditions will be identical to the set of worlds in which one would meet standards in ϕ ing if one tried to use one's method for ϕ ing. This draws one closer to Hawley's (2003: 22) view. On the other hand, if, like Elzinga (2018: 121) and Löwenstein (2017: 28) one thinks that one's competency need only *reliably* (though not perfectly) guarantee meeting normative standards, then the equivalent set of worlds becomes a bit more abundant. Alternatively, if one thinks that one's competency need only make meeting standards more likely than some benchmark level (for example, even though the best batters in professional baseball only manage to hit league pitches less than half of the time, their batting average is still much higher than that of an amateur) then the equivalent set of possible worlds might be more abundant still.

If adopted, my proposal for the first condition provides a clarification of the views of Elzinga and Löwenstein, especially as regards distinct dimensions of performance robustness.⁹ My proposal for the second condition, while providing complementary discussion to both the views of Elzinga and Löwenstein, will constitute a revision only to Elzinga's account, since Löwenstein, for the most part, already expounds the second condition of their view this way (Löwenstein 2017: 122–9). Elzinga suggests that agents manage to make trials and respond to feedback in their performances by means of dispositions to do these things, and by 'dispositions' I take Elzinga to mean more-or-less persistent, standing tendencies (Elzinga 2018: 133). Such agents may meet my version of the second condition, but my version of this condition may also be met by agents who, often enough, are unimaginative in their attempts and ignore the lessons of past failures. Such agents may still meet my version of the second condition, provided that they are able, *under some conditions or other*, to perform in a reliably resilient and conscientious way (and provided that they've *already* learned that performance method or style).¹⁰

But is it really enough for knowing how to ϕ that one is *able* (under *some* conditions or other) to execute a method that, when executed, equips the agent to cope with certain variations in the values of certain performance parameters?

It *is* enough, and we don't need to require that the ability to execute such a method be itself robust across any particular conditions. The reason why is one of the most important contributions of this paper. Recall that the preeminent function of the original first condition was to provide a response to putative counterexamples in which someone who doesn't know how to ϕ nevertheless, through sheer luck or with scaffolding from exceptionally accommodating performance conditions, manages to meet normative standards for ϕ ing.

But to depend on luck or exceptional performance conditions when meeting standards is *not* to meet these standards by means of the methods that I have been discussing. (For such methods equip an agent to non-accidentally cope with variation in performance conditions.) Therefore, the agent who can *only* meet standards with the

⁹Löwenstein (2017: 25–9), Elzinga (2018: 121–3).

¹⁰My claim is not that the resilience or conscientiousness of every performance is grounded in that performance's method or style. Rather, my claim is that there are methods and styles such that, to perform according to them is ipso facto to perform resiliently or conscientiously. Consider the methods described by 'if at first you don't succeed, try, try again' and 'stick and stay, and make it pay.'

assistance of luck or exceptional performance conditions *cannot execute such methods at all*. So it's not that the novice can do what the expert can do, but only in a smaller range of performance conditions. Rather, the novice cannot do what the expert does under *any* performance conditions, at least not until the novice masters the methods, styles, techniques, etc. by which the expert non-accidentally copes with variation in performance parameters. Once this is recognized, we no longer need to appeal to the *general robustness* of an ability in excluding cases of lucky success as cases of knowing how to φ , for the relevant ability (i.e. the ability to execute certain *methods*) is not possessed *at all* by the agent that must get lucky. (The argument for this will be elaborated on in the last section of the paper.)

4. Motivating the amendment

Above, I've argued that there is an important distinction between the range of performance parameter values that a particular performance method (or style) equips an agent to cope with, and the range of conditions under which a particular agent can reliably perform according to that method (or style). I've also argued that, once we identify knowing how to φ with an ability to φ according to a certain kind of method, we don't need to appeal to the conditions under which an agent can reliably execute such a method to secure the extensional adequacy of the account.

I'll now argue that these theses can do important work in an account of knowing how to φ . I will do this by first considering a kind of objection to the accounts of Hawley and Elzinga involving two agents who, plausibly, possess the same piece of knowledge-how. I will then show how my proposed amendment to these accounts provides a response to this kind of objection. Following this, I will consider a case in which two agents each possess what is plausibly a different piece of knowledge-how. After discussing the way in which Löwenstein's view can accommodate such a case, I will then show how Elzinga's account can, if amended in the way that I propose, accommodate it as well.

5. Sleeping on the job

Hawley's (2003: 21–2) view relativizes the content of what is known when one knows how to φ to those circumstances under which the agent would succeed in φ ing if they tried. That is, according to Hawley, the content of what one knows when one knows how to φ is always *how to φ under circumstances C*. Waights Hickman (2019: 316–17) objects to relativizing the content of knowledge-how in this way, partly because doing so seems to entail the implausible claim that Mozart and Bach possessed different pieces of knowledge-how if Mozart could compose drunk but Bach could not.

Elzinga's view escapes this version of the objection for, according to my reading of the view, one who knows how to φ need only be able to φ (in a standard-meeting way) reliably across a range of performance conditions determined by the context of the knowledge attribution (Elzinga 2018: 121, 123). In many contexts, the performance conditions relevant to 'so-and-so knew how to compose' include (relative) sobriety, so 'Mozart knew how to compose' and 'Bach knew how to compose' can both be true when said in such contexts.

Even so, a similar objection to Elzinga's view, on my reading of it, can be developed. Consider Ben and Jerry, two partners in an ice cream business. Both Ben and Jerry know how to catch their partner sleeping on the job. Their methods for doing so may both be identical, for example, *check the storeroom, and the area behind the freezer, between 2:30 and 3:30 in the afternoon*. Such a method equips both Ben and Jerry to

cope with certain variations in performance conditions, for example, they can each succeed regardless of whether their partner has chosen to nap in the storeroom or behind the freezer. Even so, there is no possible world in which both Ben and Jerry succeed (at least within a certain interval of time), because for Ben to catch Jerry sleeping, Ben cannot himself be sleeping, and thus cannot himself be *caught* sleeping. For this reason, if the contextually relevant performance conditions are understood as a single set of worlds in which both Ben and Jerry can succeed, it appears that, according to Elzinga's account, there is no single context in which 'Ben knows how to catch a partner sleeping' and 'Jerry knows how to catch a partner sleeping' can both be true.

Of course, one might try to understand these contextually relevant performance conditions as something *other* than a single set of possible worlds in which both Ben and Jerry can succeed. And this is partly what my amendment provides by recasting these conditions as performance *parameters*, or to put it another way, as a certain required *variability* within the set of performance conditions under which an agent can execute their method and meet normative standards as a result (if standards are indeed met, and regardless of whether the equivalent set of possible worlds is, so to speak, abundant or sparse). Since both Ben and Jerry can cope with the same variations in performance parameters as a result of the same methods, they both know the same thing, despite not both being able to exercise this knowledge at the same time.

Löwenstein's view is already immune to this kind of objection, for according to that account, the content of the knowledge-how attribution is relativized, not to a single set of contextually relevant performance conditions, but to those performance conditions under which the subject of the attribution can *themselves* reliably meet normative standards in ϕ ing (Löwenstein 2017: §1.4). This is how Löwenstein's version of the first condition works: If an agent knows how to ϕ , then they are able to reliably meet normative standards in ϕ ing in those performance conditions in which *they* are not substantially impeded from doing so (Löwenstein 2017).

But what counts as a substantial impediment can vary from agent to agent. Ben cannot catch Jerry sleeping if Jerry is awake, and vice versa. Nevertheless, Löwenstein can account for Ben's possessing the same knowledge-how as Jerry by appealing to the identical stock of concepts through which Ben and Jerry each assess their performances and opportunities, and each non-accidentally succeed at catching their partner sleeping as a result (within what, for each of them respectively, are the normal conditions) (Löwenstein 2017: 115–22). Part of what my proposed amendment does is generalize this aspect of Löwenstein's view by recasting any performance robustness that is to be partly constitutive of knowing how to ϕ in terms of performance parameters.

6. Two ways of solving a Rubik's cube

In the previous section I discussed an objection to Elzinga's account (adapted from an objection to Hawley's account made by Waights Hickman (2019: 316–17)). This objection involved two agents who, plausibly, possess the *same* piece of knowledge-how. I discussed how Löwenstein's view escapes this objection, and how my proposed amendment generalizes this feature of Löwenstein's account such that accounts like Elzinga's can be amended to include it as well.

In this section I'll discuss how much the same is true of certain cases involving agents who each possess what is plausibly a *different* piece of knowledge-how. Consider two agents who can both speed-solve the Rubik's cube within the same amount of time, and with the same degree of reliability. (If, like Hawley, we wished to additionally relativize the content of knowledge-how to the conditions under which an agent would, or could, succeed, we can stipulate here that both of these agents

can similarly speed-solve the Rubik's cube under exactly the same conditions, down to the last detail.)

But suppose additionally that each of these agents has learned only one method for speed-solving the Rubik's cube, and that they have each learned different methods. The first agent uses only the Advanced CFOP method, whereas the second agent uses only the Roux method.

Both of these agents know how to speed-solve the Rubik's cube, but there is a difference in their knowledge. But if knowing how to speed-solve the Rubik's cube is a matter of being able to reliably meet normative standards in doing so (e.g. speed; no cheating; etc.), and this across a certain range of performance conditions, then how are we to distinguish the difference in what they each know?

The only available response seems to be this: One agent knows how to *speed-solve the Rubik's cube using the Advanced CFOP method*, while the other agent knows how to *speed-solve the Rubik's cube using the Roux method*. But this is, in a sense, just what my proposed amendment is all about. For once we describe the relevant ϕ ing methodologically, we can exclude, as cases of knowledge-how, cases of lucky success without needing to appeal to a generally reliable ability.

This is because, as discussed above, someone who, through sheer luck, manages to speed-solve the Rubik's cube isn't solving it using the Roux method. Even if, by sheer luck, this person manipulates the Rubik's cube just as someone who is using the Roux method would, still this person isn't using the Roux method. (An argument for this would fall out of my reading of Ryle's various discussions of intelligence, semi-hypothetical statements, and 'heed' concepts. This reading will be articulated in the next section.)

What's more, while we expect that an agent who knows how to speed-solve the Rubik's cube can cope with certain kinds of variation (e.g. the agent can speed-solve a Rubik's cube that has been scrambled into a variety of starting states), it is plausible that the ability to cope with these kinds of variation is provided by the method of puzzle solving itself.

Löwenstein's view can explain both the similarity, and the differences, in what these agents know. According to this view, their knowledge-how is similar in that they both meet the same normative standards in speed-cubing under the same performance conditions, and both as a result of normative guidance (Löwenstein 2017: 122–3). But this view can also diagnose the difference in their knowledge-how by appealing to the different concepts that structure this normative guidance: Whereas one of the agents deploys the concept of the *First Two Layers*, the other agent deploys the concept of the *First Two Blocks*, etc.¹¹ My amendment makes a similar resource available to any ability-based view that follows the Hawlean strategy, in particular Elzinga's view, and does this by recasting Elzinga's self-regulation condition in methodological terms.

7. Ryle on 'thinking what one is doing'

In this section I'll argue that amending Elzinga's and Löwenstein's Ryle-inspired accounts in the way that I have proposed will make these accounts even more Rylean.

Elzinga's discussion of self-regulation, and Löwenstein's discussion of normative guidance, are both, in part, ways of trying to spell out what Ryle has in mind when he discusses the 'responsib[ility]' that is partly characteristic of the exercises of knowledge-how:

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

¹¹Cf. Löwenstein (2017: 115–22).

of what is meant is that, when they perform these operations, they tend to perform them well, i.e. correctly or efficiently or successfully. Their performances come up to certain standards, or satisfy certain criteria. But this is not enough. The well-regulated clock keeps good time and the well-drilled circus seal performs its tricks flawlessly, yet we do not call them ‘intelligent’. We reserve this title for the persons responsible for their performances. (Ryle 2009 [1949]: 17)

Appropriately, the above passage plays an important role in Elzinga’s and Löwenstein’s respective presentations of Ryle’s view about knowledge-how (Löwenstein 2017: 32; Elzinga 2018: 121). Ryle immediately elaborates:

To be intelligent is not merely to satisfy criteria, but to apply them; to regulate one’s actions and not merely to be well-regulated. A person’s performance is described as careful or skilful, 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, that is, in trying to get things right. This point is commonly expressed in the vernacular by saying that an action exhibits intelligence, if, and only if, the agent is thinking what he is doing while he is doing it, and thinking what he is doing in such a manner that he would not do the action so well if he were not thinking what he is doing. (Ryle 2009 [1949]: 17–18)

As these remarks show, Elzinga is right to draw attention to self-regulation, and to learning from experience, in explicating the nature of such ‘intelligent’ or ‘responsible’ agency. Likewise, Löwenstein is right to draw our attention to the application of criteria in performance, and to the ways in which doing so makes agents responsible and their performances intelligent.

But also important to understanding the nature of intelligent or responsible agency is Ryle’s paraphrase of self-regulation, and of the application of criteria, etc. as ‘thinking what [one] is doing while one is doing it, and thinking what [one] is doing in such a manner that [one] would not do the action so well if [one] were not thinking what [one] is doing’ (Ryle 2009 [1949]: 17–18).

Sax (2010: 519–20) observes that, in order to escape his own regress arguments, Ryle characterizes such application of criteria as the exercise of *skill*, or as the mindful exercise of certain abilities. But by going deeper into Ryle’s characterization of the *concept* of mindful performance itself we can see how the specific amendment argued for in this paper is Rylean.¹²

In a subsection of *The Concept of Mind* titled ‘Mental Occurrences,’ Ryle classifies the concept of *thinking what one is doing* as a ‘heed concept,’ or as a form of ‘minding’ (Ryle 2009 [1949]: 119). Inspecting and monitoring are special exercises of heed (Ryle 2009 [1949]: 119). Actions done with heed are actions that possess a ‘special character’:

Doing something with heed does not [essentially] consist in coupling an executive performance with a piece of theorising, investigating, scrutinising or ‘cognising’ ... [Moreover,] ... it is quite idiomatic to replace the heed verb by a heed adverb. We commonly speak of reading attentively, driving carefully and conning studiously, and this usage has the merit of suggesting that what is being described

¹²One might also adapt Sax’s (2010: 521) thesis – that the performance criteria met supply propositional content implicit in intelligent performance – and argue that the inferences licensed by the intelligent manner of a performance (as will be discussed below) also help to supply this implicit content.

is one operation with a special character and not two operations executed in different 'places', with a peculiar cable between them. (Ryle 2009 [1949]: 120–1)

That intelligent actions are actions done 'while thinking what [one] is doing' and that thinking what one is doing is acting with heed, and that acting with heed is undertaking an act that has a special character, fits nicely with Ryle's earlier claim that 'adverbs expressing intelligence-concepts ... signalize ... that the performance itself possesses a certain style, method, or *modus operandi*' (Ryle 1945: 3).

But what kind of 'style, method, or *modus operandi*' distinguishes intelligent actions from other kinds of actions? As it happens, Ryle doesn't say much about these styles and methods specifically (though his failure to do so is likely part and parcel of what Stout (2003: paras 15–18) calls Ryle's 'nonreductive behaviorism').¹³ Even so, Ryle makes certain arguments about the *logic* of heed concepts. Specifically, in *Mental Occurrences*, Ryle argues that the logical profile of heed concepts makes the act-descriptions that employ them 'semi-hypothetical,' or 'mongrel-categorical,' statements (Ryle (2009) [1949]: 124).

According to Ryle, semi-hypothetical statements are 'at once narrative, explanatory and conditionally predictive, without being a conjunctive assemblage of detachable sub-statements' ((2009) [1949]: 123). Allow me to break that down a bit. To describe a driver as *driving carefully* is to state *categorically* what the driver is doing, but in a way that also functions as an *explanation* of the very same act under a different description (e.g. that a particular driver is driving carefully explains this driver's looking over their right shoulder to check a blind spot, though these are not two different acts in the sense that we could ask which one was done first), and also licenses *predictions* about what the driver will or would do (e.g. that the driver is driving carefully licenses the prediction that, should the driver need to merge into the left lane, they will first look over their left shoulder to check the left-side blind spot) (cf. Ryle (2009) [1949]: 125).^{14,15}

But to describe the driver as driving carefully is *not* to say, elliptically, several different things at once – a categorical report; an explanation of what is reported categorically; and a prediction of further occurrences. Rather, a semi-hypothetical statement is supposed to function in all of these ways at once (Ryle (2009) [1949]: 123).¹⁶

¹³Kremer (2017a: 37) argues that Ryle relies on the evolving relationship between teacher and student as a model, not just of the development of competent performance generally, but also of the development of the heeding relied on in such performance: a learner comes to 'double the roles of instructor and pupil ... to coach himself and heed his own coaching' (Ryle (2009) [1949]: 130). (Elzinga discusses closely related themes.) There is a suggestive overlap between what a coach would instruct a pupil to do, and the inferences about competent behavior licensed by certain semi-hypothetical characterizations of such behavior, though discussion of this is beyond the scope of this paper.

¹⁴Cf. White (1964, 1967) for a narrower construal of the explanatory function of semi-hypotheticals.

¹⁵We can expect that the specific inferences licensed will vary depending on the context or purpose in/for which the semi-hypothetical description is used given that, as Tanney (2017: 7) argues, Ryle is committed to meaning being, in the first instance, a property of whole sentences (or of even larger units of discourse).

¹⁶Small (2017: 69–70) catalogs four features of the exercises of skill that Ryle uses to distinguish them from the exercises of mere habit. Among these are that (1) exercises of skill involve what Ryle calls 'paying heed' to what one is doing, and that (2) when exercising a skill an agent is poised to cope with variation in the particularities of the performance situation. Later, in his discussion of what, for Ryle, 'paying heed' is supposed to amount to, Small focuses on the idea that the dispositional excellences subserving a skill are actualized in a special (practical knowledge-entailing) way when the skill is exercised. Part of my aim in this paper is to examine the ways in which the special logic of (heed-type) semi-hypothetical statements explains how features (1 and 2) are related, and in this way to contribute to our understanding of what, for Ryle, 'paying heed' is supposed to amount to.

Part of Ryle's argument for the semi-hypothetical interpretation of heed concepts involves demystifying semi-hypotheticals, and this he does by arguing that they are ubiquitous in our everyday talk about non-mental phenomena (Ryle (2009) [1949]: 124). For example, Ryle argues that 'the sugar lump is dissolving' and 'the bird is migrating' are both semi-hypothetical statements, and claims that most adduced examples of categorical statements are in fact semi-hypothetical statements (Ryle (2009) [1949]: 124).

But here is where a number of Ryle's expositors seem to have missed an important part of Ryle's thesis.¹⁷ As these expositors observe, Ryle's discussion often focuses on semi-hypotheticals like 'the sugar lump is dissolving' or 'the bird is migrating' or 'you *would* miss the last train' (Ryle (2009) [1949]: 123–4). In these examples, the explanatory and predictive power of what is described semi-hypothetically is grounded in a standing *tendency* of the subject of the statement (or, perhaps, in the case of the migrating bird, in a *telos*). For example, the explanatory and predictive power of the fact that the sugar lump is dissolving (e.g. that the sugar lump is dissolving can explain why its sugar molecules are separating) is grounded in a stable, standing tendency of sugar to interact in a certain way with certain liquids.

Given this, one might be tempted to think that semi-hypothetical description is just a way of characterizing an event as the *manifestation of a tendency* for that kind of event to happen under certain conditions. This is what seems to be the case with 'the sugar lump is dissolving,' 'you *would* miss the last train,' and, on a certain (non-teleological) reading, 'the bird is migrating.'

But this would be a mistake, for 'Gilbert is driving carefully' does not entail that Gilbert *tends* to drive carefully, or otherwise *tends* to do what careful drivers do.¹⁸ Gilbert may *now* be driving carefully, despite his usual tendency to drive carelessly. Furthermore, 'it is proper to order or request someone to apply his mind, as it is not proper to order him to be able or likely to do things' (Ryle (2009) [1949]: 122). Still, the fact that (a usually careless) Gilbert is right now driving carefully has explanatory and predictive power (it may explain his now checking his blind spots, and license the prediction that he will modulate his speed in the near future if posted speed limits require it, etc.).

Since Gilbert can be driving carefully, despite his usual tendency to do otherwise, the explanatory and predictive power of the fact that he *is* driving carefully cannot, in general, be grounded in Gilbert's standing tendencies, but must instead be grounded (at least proximally) in a different 'special character' of his token driving performance, a special character that we indicate with an adverb of *manner* (viz 'carefully') (Ryle 2009 [1949]: 120–1).^{19,20}

¹⁷Expositors that seem to have missed what I will discuss presently include Weldon (1950), Hofstadter (1951), King (1951), Lyons (1979), Gram (1974), and Parry (1980).

¹⁸For some discussion of this, see White (1964, 1967).

¹⁹In this way (heed-type) semi-hypotheticals can be seen as restoring the significance of token behavior events in an otherwise dispositional treatment of the concepts of knowledge. Cf. Hornsby (2011: 91) who argues that 'the difference between the practiced typist and the learner is a difference in what they know how to do but is not apparent in the "token events" they participate in when they exercise their knowledge.' Ryle would deny this, and he introduces (heed-type) semi-hypotheticals, in part, as a way to do so consistently with his other commitments.

²⁰Bäckström and Gustafsson (2017: 50) argue convincingly that Ryle is not a traditional behaviorist in part because he characterizes intelligent performances in terms of dispositions that are categorically different (in Ryle's sense of 'category') from the kind of dispositions adverted to by the traditional behaviorist. Plausibly, the distinction (made in this paper) between heed-type, and tendency-type, semi-hypothetical descriptions can add to our understanding of what, for Ryle, this categorical distinction is supposed to amount to.

To recapitulate, Ryle's semi-hypothetical statements are statements that are supposed to indicate, categorically, that a certain kind of event is (or was, or will be) happening which in turn licenses certain explanations and predictions (Ryle 2009 [1949]: 123–5).²¹ Some semi-hypothetical statements do this by characterizing an event as the manifestation of a tendency for that kind of event to happen (e.g. 'the sugar lump is dissolving' and 'you *would* miss the last train'). But, crucially, other semi-hypothetical statements do this by instead adverting to a different 'special character' of the event itself, such as that adverted to by certain adverbs of *manner* (e.g. 'Gilbert was driving *carefully*'). (This is the distinction that has often been overlooked by Ryle's expositors.) According to Ryle, though attributions of heed are, in general, semi-hypothetical statements of the second type, their logical profile is no more philosophically 'scandalous' than those of the first type (Ryle 2009 [1949]: 123).

As a heed concept, *thinking what one is doing* is evidently like *ϕing carefully* in that one can think what one is doing on a particular occasion despite a general tendency to perform absentmindedly. And even if one generally tends to perform absentmindedly, still, the fact that one *is* thinking what one is doing on a particular occasion will have explanatory and predictive power that is, according to Ryle, grounded (at least proximally) in a special character of the performance itself. So 'Gilbert is thinking what he is doing while he is doing it' is also a semi-hypothetical statement that adverts, not necessarily to any particular standing tendency of Gilbert's, but to some other non-etiological 'special character' of his performance.²²

But what kind of 'special character' of Gilbert's performance is this? To find the answer, let us recall Ryle's earlier claim that 'adverbs expressing intelligence-concepts ... signalize ... that the performance itself possesses a certain style, method, or *modus operandi*' (Ryle 1945: 3) Though Ryle doesn't mention performance methods or styles explicitly in *Mental Occurrences*, methodological act-description (e.g. 'solving the Rubik's cube using the Roux method') would certainly qualify as semi-hypothetical: The fact that someone is using the Roux method could explain what they are doing under a different description (e.g. that they are now making a certain *M* move), and could license predictions about what they will or would do (e.g. that the last thing they will do is solve the Last Six Edges).

Consider again the second passage quoted in this section of the paper. Here, Ryle claims that the criteria for ascribing intelligence to an agent include the agent's satisfying and applying criteria, being ready to detect and correct lapses, being ready to repeat and improve upon successes, and profiting from examples when they are available (Ryle 2009

²¹One might object that all categorically stated matters of fact are such that certain explanations or predictions follow from them. But the idea seems to be that, for Ryle, there are certain concepts whose distinguishing function is to license a certain kind of inference. Such concepts include those expressed in dispositional (e.g. 'tends to' or 'can'), hypothetical (e.g. 'if ... then'), and semi-hypothetical (e.g. 'carefully') statements.

²²Kremer (2017b: 27) argues that Ryle was positioning himself outside of a period dichotomy between self-styled intellectualists and anti-intellectualists of his day. Despite their disagreement about the (potential) role of explicit reasoning in human behavior, both of these camps seemed to agree that all human behavior was either intelligent in virtue of explicit intellectual processes; or was the product of mere instinct; or was the product of mere conditioning. However, as Kremer observes, Ryle was at pains to argue that this scheme was not complete in that many performances are intelligent in virtue of a learned skill exercised therein. Ryle argues that such intelligent performances involve a kind of thinking, and for this reason are more than the products of mere conditioning, but neither do they involve the explicit reasoning of interest to the intellectualists of his day. This paper's exegetical work on (heed-type) semi-hypotheticals adds to our detailed understanding of how Ryle attempted to characterize the kind of thinking that he appealed to in escaping the dichotomy, a thinking that does not involve explicit reasoning.

[1949]: 17–18). What's more, Ryle paraphrases an agent's meeting these criteria as the agent's *thinking what [they are] doing in such a manner that [they] would not do the action so well if [they] were not thinking what [they are] doing* (Ryle 2009 [1949]: 17–18).

A plausible interpretation now presents itself: The 'special character' of Gilbert's performance adverted to by 'Gilbert is thinking what he is doing in such a manner that' is 'a certain style, method, or *modus operandi*' (Ryle 1945: 3) characterizing Gilbert's performance that both *explains* and *predicts* (cf. Ryle 2009 [1949]: 125) his doing a number of finer-grained things in the course of performing in this way – satisfying and applying criteria, detecting and correcting lapses if need be, repeating and improving upon successes where possible, and profiting from the examples of others when they are available, etc. (cf. Ryle 2009 [1949]: 17–18).

This should remind the reader of the amendment to the views of Elzinga and Löwenstein that was proposed above. I proposed that Elzinga's self-regulation condition be understood as an ability to ϕ in a certain *way*, namely ϕ ing sensitively to the results of previous trials, etc.²³ And, according to Ryle, 'Gilbert is now ϕ ing sensitively to the results of previous trials' is a heed-type semi-hypothetical statement which describes (in methodological or stylistic terms) a fact that can explain, for example, Gilbert's now repeating something that has worked for him in the past.

Similarly, I proposed that Löwenstein's normative guidance condition be understood as an ability to ϕ according to a certain *method* that includes, as a step, exercising a certain capacity for judgement. (Indeed, this is how Löwenstein, for the most part, already exposit their view.)²⁴ Understood in this way, 'Gilbert is ϕ ing in way *w* on the basis of his judgement that doing so is his best bet for meeting normative standards for ϕ ing under present conditions' is also a heed-type semi-hypothetical statement which describes a fact that can explain, for example, Gilbert's ϕ ing in way *w*. Plausibly, both 'Gilbert is ϕ ing sensitively the results of previous trials' and 'Gilbert is ϕ ing in way *w* on the basis of his judgement that ...' capture (at least part of) the heed-type semi-hypothetical content that Ryle would ascribe to 'Gilbert is thinking what he is doing in such a way that ...'

Gilbert's readiness to do these things equips him to cope with variation in the values of certain performance parameters. Consider Ryle's example of the intelligently performing mountaineer who, 'walking over ice-covered rocks in a high wind in the dark ... *thinks what he is doing* ... [and as such] is ready for emergencies ... economises in effort ... makes tests and experiments' etc. [*italics mine*] (Ryle 2009 [1949]: 30). Such a mountaineer knows how to mountaineer, and as such performs well under a variety of conditions in virtue of his ability to think what he is doing. But this variety of conditions under which the mountaineer can perform well do not so much define a set of possible worlds, since we don't know what the mountaineer needs to be ready to think what he is doing (e.g. coffee or no coffee). Rather, I propose that it is in virtue of his ability to *think what he is doing in such a manner that ...* that the mountaineer is able to perform well given a variety of values for certain performance *parameters*, such as the presence or absence of hidden crevasses, more and less slippery surfaces, changes in wind strength, etc.

But even without knowing how robust the mountaineer's ability to think what he is doing is, still we can distinguish this mountaineer from the novice who must get lucky to perform well. For when someone is *thinking what [they] are doing in such a manner that ...*, this fact has an explanatory and predictive power that the fact of a novice getting lucky does not have.

²³Cf. Elzinga (2018: 132–4).

²⁴Cf. Löwenstein (2017: Ch. 4, especially §§4.2–4.4).

For example, the fact that an experienced bread baker is *thinking what [they] are doing in such a manner that ...* could explain why they cover their dough before they let it rise, whereas the fact that a novice is getting lucky would not explain their covering the dough in the same way. That an experienced bread baker is *thinking what [they] are doing in such a manner that ...* may license the prediction that they will add steam to the oven as necessary, whereas the fact that a novice is getting lucky does not license genuine predictions about what they will or would do under particular conditions.

If Ryle is right about this, then these differences show that our concept of what the experienced baker *does* (when they are *thinking what they are doing in such a manner that ...* – i.e. when they are *exercising* their knowledge-how) is different from our concept of what the lucky novice *does*, and this even on the unlikely occasion when their performances are difficult to ‘photographically [or] gramophonically’ distinguish. For this reason, the novice *cannot* do what the experienced baker can do, even when the novice gets lucky. So an agent’s ability to *think what they are doing in such a manner that ...*, regardless of the enabling conditions that must be met if this agent is to exercise this ability (e.g. several cups of coffee; no coffee at all; the appeasement of personal superstitions; etc.), is enough to distinguish this agent from other agents who must get lucky if they are to φ in a standard-meeting way.

Nevertheless, doesn’t Ryle commit himself to excluding cases of lucky success on the basis of the *unreliability* of luck? I don’t believe that Ryle is committed to this. Consider again the first passage quoted in this section of the paper (*viz.* Ryle 2009 [1949]: 17). There, Ryle claims that when describing someone as knowing how to make jokes, fish, etc., we mean in part that ‘when they perform these operations, they tend to perform them well’ (Ryle 2009 [1949]: 17). This might seem to support reading Ryle as holding the view that knowing how to φ is, in part, constituted by a generally reliable ability to φ well. But I suggest that we instead read Ryle as claiming that when someone who knows how to φ *uses this knowledge* in the form of thinking what they are doing, they tend to perform well.²⁵ This reading makes Ryle’s view consistent with the plausible thesis that some people can only exercise certain knowledge-how under favorable conditions, including prior cups of coffee or the absence of certain kinds of psychological distress.

Further support for this reading can be found in *Mental Occurrences*, where Ryle considers an acquired capacity for meeting performance standards that does not issue in genuinely intelligent performances (Ryle 2009 [1949]: 129). Ryle claims that part of the difference between such a capacity and genuinely intelligent capacities lies in the fact that the latter are capacities for coping with certain kinds of *variation*. For example, a child:

has not learned [the multiplication tables] properly unless he can promptly give the right answer to any snap multiplication problem (lower than 12×13), and unless he can apply his tables by telling us, e.g. how many toes there are in a room in which there are six people. Nor is a man a trained rock-climber who can cope only with the same nursery-climbs over which he was taught, in conditions just like those in which he was taught, and then only by going through the very motions which he had been then made to perform. Learning is becoming capable of doing some correct or suitable thing in any situations of certain general

²⁵A complete defense of my reading of Ryle would include providing a similar interpretation of other passages from *The Concept of Mind*. Considerations of space prevent me from doing this here.

sorts. It is becoming prepared for *variable* calls within certain ranges. [original italics] (Ryle 2009 [1949]: 129)

In both of the above examples, the kind of variation Ryle mentions is what I have discussed earlier as variation in the values of certain performance *parameters*. The rock climber must be able to cope with variation in climbing challenges, but this isn't to say anything about what the rock climber requires to be ready to climb – two cups of coffee; or no coffee at all; etc.

8. How my amendment is Rylean

Above, I proposed an amendment to the views of Elzinga and Löwenstein respectively, according to which the second condition of their accounts (self-regulation in the case of Elzinga (2018: 132–8); normative guidance in the case of Löwenstein (2017: §§1.4, 4.3, 4.4)) be understood as an ability to perform according to a certain *method* (or *style*). I proposed also that the first condition of these accounts be understood as an ability to reliably cope with certain variation in the values of certain performance *parameters*, and this in virtue of the method described in the second condition.

Amending these accounts in this way would make them more Rylean. This is because the Rylean self-regulation to which Elzinga appeals, and the Rylean normative guidance to which Löwenstein appeals, are both understood by Ryle as concepts for precisifying the 'special character,' or the 'style, method, or *modus operandi*,' of performances in which someone who knows how to φ uses this knowledge by *thinking what they are doing in such a manner that they would not do the action so well if they were not thinking what they are doing*.²⁶

References

- Bäckström S. and Gustafsson M. (2017). 'Skill, Drill, and Intelligent Performance: Ryle and Intellectualism.' *Journal for the History of Analytical Philosophy* 5(5), 40–55.
- Elzinga B. (2018). 'Self-regulation and Knowledge How.' *Episteme* 15(1), 119–40.
- Gram M.S. (1974). 'Mongrel Categoricals and Category Mistakes.' *Southern Journal of Philosophy* 12(1), 35–47.
- Hawley K. (2003). 'Success and Knowledge-how.' *American Philosophical Quarterly* 40(1), 19–31.
- Hofstadter A. (1951). 'Professor Ryle's Category-mistake.' *Journal of Philosophy* 48(9), 257–70.
- Hornsby J. (2011). 'Ryle's Knowing How and Knowing How to Act.' In J. Bengson and M.A. Moffett (eds), *Knowing How: Essays on Knowledge, Mind, and Action*, pp. 80–100. New York, NY: Oxford University Press.
- King H.R. (1951). 'Professor Ryle and the Concept of Mind.' *Journal of Philosophy* 48(9), 280–96.
- Kremer M. (2017a). 'A Capacity to Get Things Right: Gilbert Ryle on Knowledge.' *European Journal of Philosophy* 25(1), 25–46.
- Kremer M. (2017b). 'Ryle's "Intellectualist Legend" in Historical Context.' *Journal for the History of Analytical Philosophy* 5(5), 16–39.
- Löwenstein D. (2017). *Know-how as Competence. A Rylean Responsibilist Account*. Frankfurt am Main: Vittorio Klostermann. <https://philarchive.org/archive/LWEKACv2>.
- Lyons W. (1979). 'Ryle's Three Accounts of Thinking.' *International Philosophical Quarterly* 19(4), 443–50.
- Parry R.D. (1980). 'Ryle's Theory of Action in the Concept of Mind.' *Philosophy and Phenomenological Research* 40(3), 379–92.

²⁶I would like to thank David Hills, Krista Lawlor, Rosa Cao, and Thomas Icard III for their discussion, comments, and encouragement on earlier drafts of this project. I would also like to thank three anonymous referees, and one anonymous referee for *Episteme*, for very useful comments on earlier drafts of this project.

- Ryle G.** (1945). 'Knowing How and Knowing That: The Presidential Address.' *Proceedings of the Aristotelian Society* **46**, 1–16.
- Ryle G.** (2009 [1949]). *The Concept of Mind*. London: Routledge.
- Sax G.** (2010). 'Having Know-how: Intellect, Action, and Recent Work on Ryle's Distinction Between Knowledge-how and Knowledge-that.' *Pacific Philosophical Quarterly* **91**(4), 507–30.
- Small W.** (2017). 'Ryle on the Explanatory Role of Knowledge How.' *Journal for the History of Analytical Philosophy* **5**(5), 56–76.
- Stout R.** (2003). 'Ryle's Behaviourism.' *Revue internationale de philosophie* **1**, 37–49.
- Tanney J.** (2017). 'Volume Introduction: Gilbert Ryle on Propositions, Propositional Attitudes, and Theoretical Knowledge.' *Journal for the History of Analytical Philosophy* **5**(5), 1–15.
- Waights Hickman N.** (2019). 'Knowing in the "Executive Way": Knowing How, Rules, Methods, Principles and Criteria.' *Philosophy and Phenomenological Research* **99**(2), 311–335.
- Weldon T.** (1950). 'The Concept of Mind.' *Philosophy* **25**(94), 266–270.
- White A.** (1964). *Attention*. Oxford: Basil Blackwell.
- White A.** (1967). *The Philosophy of Mind*. New York, NY: Random House.

John Turman is a PhD candidate in philosophy at Stanford University and is currently writing a dissertation on competence and philosophical neo-behaviorism. His research interests include philosophy of mind, epistemology, philosophy of action, and Ryle.