



The Mark of Understanding: In Defense of an Ability Account

Sven Delarivière^{1,2} · Bart Van Kerkhove¹

Received: 19 June 2020 / Accepted: 25 November 2020 / Published online: 18 January 2021
© The Author(s), under exclusive licence to Springer Nature B.V. part of Springer Nature 2021

Abstract

Understanding is a valued trait in any epistemic practice, scientific or not. Yet, when it comes to characterizing its nature, the notion has not received the philosophical attention it deserves. We have set ourselves three tasks in this paper. First, we defend the importance of this endeavor. Second, we consider and criticize a number of proposals to this effect. Third, we defend an alternative account, focusing on abilities as the proper mark of understanding.

Keywords Understanding · Abilities · Counterfactuals · Mathematical practice

1 Introduction

Understanding is a valued trait in any discipline and it is fair to say that we sometimes uncontroversially attribute it to certain subjects. But is there a systematic way to reveal what understanding-attributions should entail? What is it in those situations of uncontroversial attribution or dismissal that guides us—or should guide us—in ascribing (or denying) understanding? Answering this question involves specifying what systematic feature we find so philosophically or epistemically valuable about understanding and thus necessary for (and explanatory about) its attribution, regardless of who it is attributed to or what it is about. We shall call this the “mark of understanding,”¹ because it is what *demarcates* it, and characterising this mark is what the focus of this paper will be. We begin, in Sect. 2, by defending

¹ In the same spirit as the “mark of the cognitive” (Adams and Aizawa 2010, p. 46) in philosophy of mind.

✉ Bart Van Kerkhove
bart.van.kerkhove@vub.be

¹ Centre for Logic and Philosophy of Science, Vrije Universiteit Brussel, Brussels, Belgium

² FWO-Vlaanderen, Brussels, Belgium

the importance of this endeavor. After considering some proposals (sense-, state- and synonym-based accounts) and their flaws in Sect. 3, we shall in Sect. 4 argue that understanding-attributions should and do always boil down to a subject possessing a particular set of relevant abilities, composed of acts (salient to the object for a certain context), and that this is the most coherent and useful conceptualisation of “understanding,” because it sidelines the role of feelings (which are so distrusted) and avoids the pitfalls of mental states (which we can neither discern or value directly).

2 The Value of Understanding and Identifying its Mark

Before considering any possible conceptualisations of understanding, it will be worthwhile to consider what makes both understanding and its mark valuable in the first place. To begin, it is obvious that humans strive for understanding. Indeed, understanding is considered a valued trait or aim in just about any epistemic activity or discipline. Both “scientists and laypeople alike will typically regard understanding as one of the most important and highly valued products of scientific research and teaching” (de Regt 2017, p. 1). The promise of (better) understanding is driving science, art and religion (Baumberger et al. 2016).

Also in mathematics, often taken to be a special case within the sciences, understanding takes a central role. That students should understand mathematics is “[o]ne of the most widely accepted ideas with the mathematics education community” (Hiebert and Carpenter 1992, p. 65), and achieving this is “[t]he main goal of elaborating teaching designs, projects, new software and textbooks” (Sierpinska 1990, p. 24). Yet; understanding is not confined to education alone. To gain understanding is a key motivation even for research mathematicians. This was made all the clearer when computers started contributing to their field. When the Four Color Theorem was proved with the help of a huge amount of automated testing (Swart 1980), one of the most heavily discussed aspects was the value of this contribution. For instance, Frank Bonsall, a leading mathematician, said:

We cannot possibly achieve what I regard as the essential element of a proof—our own personal understanding—if part of the argument is hidden away in a box. (as quoted in MacKenzie 2004, p. 102)

When it comes to mathematics, understanding is not only a motivator to prove (Rav 1999), but also a motivator to re-prove what is already known, but not understood (Dawson 2006; Thurston 1998) and a possible criterion to distinguish proofs that merely demonstrate from proofs that explain (Delarivière et al. 2017; Frans 2020). Understanding may be the driving force of mathematical practice a lot more than formal validity (Delarivière and Van Kerkhove 2017). When someone studies or practices mathematics, their aim is not usually to only acquire or contribute to a list of mathematical results, but also to *understand* them.

The value of understanding has been explicitly discussed in the field of epistemology (see Grimm 2012 for an overview), with some authors (e.g. Zagzebski 2001; Kvanvig 2003; Pritchard 2009) arguing that its value exceeds that of

knowledge, because “we would surely rather understand than merely know” (Pritchard 2009, p. 30). This distinction in value does not, however, enjoy a consensus, with disagreement resting largely on the proposed distinctions between the nature of understanding and knowledge.² But whatever the relative value between knowledge and understanding, or the distinction that motivates it, it is not the value of understanding that is at issue. Some go as far as saying that the value of understanding is an intrinsic one (Griffin 1986). Whether this is true or not, clearly understanding brings with it certain benefits (e.g. being able to predict an outcome, improve powers, avoid dangers, etc.). Where there is understanding, there is an opportunity to exploit that understanding to your one’s gains.

If understanding is a valuable epistemological trait, then we need a fruitful epistemological characterization of it. Understanding is itself yet another target which we may strive to understand. And yet, the notion of understanding has in the past largely resisted such characterization. Philosophers have tended to regard the concept with much suspicion, or not much at all.

In the field of epistemology, the related knowledge concept—which has enjoyed a long tradition of philosophical investigation—has received the characterisation of “justified true belief,”³ along with explications on what that could or should entail, as well as an open discussion about any further conditions or constraints sparked by Gettier (1963). But understanding doesn’t share this longevity in investigation or richness in characterisation. This may, in part, be due to a simple lack of differentiation in terms. Locke’s *An Essay Concerning Human Understanding*, for instance, used “understanding” as a synonym for “knowledge.” Both derive from the word “episteme,” but the focus in philosophy has evolved, or shifted (Grimm 2012), into a characterisation of what we today call “knowledge” and not what we call “understanding” (Baumberger et al. 2016). This is to the lament of recent epistemologists (e.g. Zagzebski 2001; Kvanvig 2003; Elgin 2007; Pritchard 2009). While claims of knowledge can be corroborated or contested with philosophical criteria, when it comes to understanding—which is of no less value—we are largely reduced to intuition and hand waving.

Furthermore, if understanding is a central aim for so many of the sciences, as we saw earlier, would it not be useful to know what is its target, so that they may improve their aim? Yet even in the field of philosophy of science, the concept of understanding was long left out of focus in favour of other concepts. Understanding has close ties with the concept of explanation, which has benefitted from a lot more attention in the literature (e.g. Hempel 1965; Steiner 1978; Kitcher 1989). The tie is not hidden. In fact, “virtually every theory of explanation also places a premium on the power of an explanation to produce understanding” (Trout 2005, p. 198).⁴ Yet

² The proposed difference-maker ranges from understanding’s transparency (Zagzebski 2001) or coherence (Kvanvig 2003), its status as an intellectual achievement or cognitive ability (Pritchard 2009). If understanding is simply a type of knowledge, then the issue is a non-starter based on an unfair comparison (Brogaard 2005).

³ A characterisation not quite as old as it is deemed to be (Dutant 2015).

⁴ This makes the concept of explanation implicitly tied to the concept of understanding. Though not necessarily vice versa. Lipton (2009) argues that one can acquire understanding without an explanation. If this is true, then an account of understanding is less dependent on an implicit account of explanation than the other way around.

this did not curry attention to the concept of understanding in philosophy of science. The justification for this lack of attention was understanding's "pragmatic" nature. If we focus on understanding, so the argument goes, we are limited to investigating the subjective and relative responses of a single individual.

"Very broadly speaking, to explain something to a person is to make it plain and intelligible to [her], to make [her] understand it. Thus construed, the word 'explanation' and its cognates are pragmatic terms: their use requires reference to the persons involved in the process of explaining. In a pragmatic context we might say, for example, that a given account A explains fact X to person P1. We will then have to bear in mind that the same account may well not constitute an explanation of X for another person P2, who might not even regard X as requiring an explanation, or who might find the account A unintelligible or unilluminating, or irrelevant to what puzzles [her] about X. Explanation in this pragmatic sense is thus a relative notion: something can be significantly said to constitute an explanation in this sense only for this or that individual." (Hempel 1965, pp. 425–426)

It is true that concepts like understanding are sterile without a reference to a subject. For instance, it is Marie who understands why the global temperatures are rising, Bill who understands why the square root of two must be irrational, and Wendy who understands why the old theory failed to account for the facts. In this pragmatic context, whether an explanation indeed provides understanding is dependent not only on the explanation, but on the agent receiving the explanation, her beliefs at the time, her intelligence, her critical standards, her personal idiosyncrasies, and so forth. This, according to Hempel (1965), is the wrong focus. The pragmatic concept may be of interest to psychologists or educators, but not to epistemologists or philosophers of science, where the focus is not on whether an argument is subjectively effective for a particular individual, but whether it objectively explains the facts.

There is a credit to this evasion of the pragmatic, in the sense that the concept of explanation is intended to cover more than its effects on a particular individual. But there is also a danger to it, in the sense that the concept of explanation would be vacuous if it is independent of any and all individuals. Remember the premium that accounts of explanation place on its power to impart understanding. If a particular candidate argument provides no understanding to any individual in humanity, but it does fit a philosopher's account of explanation, do we fault humanity or the philosopher? So even within the philosophy of explanation, the concept of understanding is not sterile. Furthermore, the fact that it is pragmatic does not entail a lack of objectivity (unless in a very narrow sense of the word). So much has been argued for by, for instance, Friedman (1974), and further developed by de Regt (2017). The pragmatic nature of understanding is thus no justification for its exclusion, and its exclusion may lead to a more sterile approach to related concepts.⁵

⁵ See (Delarivière et al. 2017) for elaborations on how the exclusion of the role of understanding restricts accounts of explanation.

So far, we have only focused on the value of the mark of understanding as it applies to humans, but the concept of understanding is equally beneficial to android epistemology, where the aim is to have a better grasp of the process and limits of knowledge and understanding in artificial agents (see e.g. Ford et al. 1995, 2006; Delarivière and Van Kerkhove 2017). The use of computers in mathematics research provides yet another interesting example here, because it provoked a fundamental discussion as to their epistemic capacities. The discussion centered on three issues: (a) reliability, (b) surveyability or intelligibility and (c) its providing or being provided by understanding. Based on one or several of these, some people considered computer proofs to be: uninteresting or unsatisfying mathematics, a completely different sort of mathematics, or no mathematics at all (MacKenzie 1999; Vervloesem 2007). The issue sparked a debate about the differences or similarities between computer proof and its traditional human counterpart. However, both computers and humans are subject to reliability and (sometimes) surveyability issues, making it hard to argue for a dichotomy between the two (Delarivière and Van Kerkhove 2017). Nonetheless, humans are considered as more trustworthy due to another quality they possess or supply. For instance, the mathematics community accepts peer-reviewed results without everyone partaking in this process, instead allowing peer reviewers to function as the testimony of trustworthy black boxes (Geist et al. 2010).

So, what computers are currently lacking, and what mathematicians seem to find most unsatisfying about them, is something humans can deliver besides formal rigor (MacKenzie 1999; Avigad 2008). What human surveyors (in the best cases) supply to warrant peer review and what provers supply that empower their proofs is (c) understanding. According to Rav (1999), this focus on understanding means the primary goal of mathematics is the development of mathematical meaning which cannot be derived from formal expressions, but instead requires active interpretation, an “irreducible semantic component” (Rav 1999, p. 11). On the basis of such claims, computers get pushed outside the realm of understanding and thus the locus of mathematics. While this lack of understanding often gets mentioned (MacKenzie 1999), and is assumed to constitute a necessary difference or dichotomy, the critique is vague and little is done to explicate or investigate what this informal understanding might actually or preferably entail, as well as when any of its characterizing criteria are met or left unsatisfied. Avigad (2008) laments this lacuna in philosophy of mathematics in particular:

[T]here is a smaller, but significant, community trying to automate mathematical discovery and concept formation (...) If there is any domain of scientific inquiry for which one might expect the philosophy of mathematics to play a supporting role, this is it. The fact that the philosophy of mathematics provides virtually no practical guidance in the appropriate use of common epistemic terms may lead some to wonder what, exactly, philosophers are doing to earn their keep. (Avigad 2008, pp. 313–314)

Only if we have a better understanding of the concept of understanding, are we truly in a position to evaluate the success and alleviate the failures of understanding (or stipulate its impossibility) in artificial systems. Not only is understanding’s

alleged pragmatic nature no justification for evading its mark (which could even prove detrimental to concepts that are directly or indirectly tied to it), the pragmatic nature of understanding also makes clear that the potential and flaws of the targeted subjects cannot be adequately expressed without a mark of understanding. Therefore, if understanding is a valued trait and aim comparable to that of knowledge, its mark deserves to be described with a philosophical care equaling that of knowledge.

3 The Mark Question I: Dispelling Sense, State and Synonym Accounts

Do we really need a coherent philosophical conception of understanding before we may talk of its applicability? If we ascribed value to understanding, then it should follow that we have at least a *sense* of what the target of that value-ascription is. It would be hard to value something without a sense of what it is that is being valued. Of course, this doesn't entail that we can give a full conceptual description of it, but we may be able to "recognise it when we see it." This may very well be true. Nonetheless, it would be easier to find if we know where to look, and we would see it more clearly if we know where it begins and ends. In other words, both the pursuit and demarcation of understanding would be facilitated by having a clear and coherent philosophical concept that marks it out. So, when we attribute someone with understanding, is there a systematic way to reveal what that should mean?

Where do we begin in constructing an answer? Perhaps by demarcating what should take precedence in the discussion. To do that, allow us to formulate a simple claim about what we are talking about when we talk of understanding, namely that:

Understanding is a trait (T), of a subject (S), concerning an object (X)

We've proposed this fairly uncontroversial sentence such that, if we unpack it, three aspects of understanding present themselves for characterisation, namely (T) the trait itself, (X) the object it concerns and (S) the subject that possesses it. This allows us to branch the discussion in three different sub-topics, each focusing on a different aspect, each of which can be studied and characterized semi-independently of one another. "Semi," because the intention is not to sever the branches, but merely to prevent their conflation. Most epistemological treatments seem to focus on only one or two of these (mostly property or object), while leaving the others ignored or implicit. We hope to show that some of the convoluted problems this creates can be rather easily avoided by taking a step back and seeing the full (coherent) picture. Before tackling some proposed accounts, each of these branches will now be briefly addressed in their own turn.

The first conceptual branch is that of the trait of understanding (T). Before we can begin spelling out who it is that understands and what it is that is understood, as well as why this is so, we need to be able to specify what it is about them that makes us ascribe them with understanding, regardless of who it applies to or what it is about. Obviously, this is not to deny the pragmatic nature of understanding, by treating understanding as if it were some ontological state detached from any individual. It is

to focus our attention on what it is that these ascriptions of understanding (to various subjects regarding various objects) have in common. It is important that we start with this mark, and not its corollaries, so as to weed out the possibility of conceptual entanglement. As a conceptual branch, it has been the primary focus of the epistemology literature concerning understanding, and it will be the main focus of this paper.

Overall, the object of understanding (X) has, as a conceptual branch, received the most attention of all, but specifically in the philosophy of education literature. Most articles that deal with educating students on a particular topic and how to do it well, do focus on how to help students reach understanding of a specific object, X, and how to evaluate what they have learnt. While this literature is interesting, its goal is largely practical and local, and therefore not usually focused on making explicit the notion of understanding as a whole⁶ or how to conceive of the different kinds, ways and qualities (as well as their contextual variations) in which its mark manifests itself. We shall not elaborate further on this here, but for our current purposes, it is important to keep in mind that the mark of understanding needs to allow manifesting itself differently depending on the various objects of understanding, and indicate a way to conceptualise these differences.

This leaves us with the subject enjoying understanding (S). As a conceptual branch, it has been severely under-nurtured in the philosophy of science and epistemology literature. With the exception of Toon (2015), we cannot cite a single article that explicitly covers what it takes for understanding to belong to a subject and what makes out that subject. This is quite strange for a concept which is considered as pragmatic, in the sense of being predicated of a subject. This is not to say that there is no literature at all to draw from. The knowing subject, for instance, is a topic which has received some substantial attention in recent years, by drawing from other lucrative fields, such as philosophy of mind or cognitive science (e.g. Tollefsen 2006; Palermos and Pritchard 2016; Delarivière 2020a). Again, this is a topic we shall not be dealing with in the present paper, but it will be important to keep in mind that whatever marks understanding needs to allow manifesting itself in various subjects.

Let us then consider some proposals of what it is that a subject could (or should) possess when we attribute her with understanding. This involves specifying what we find so philosophically or epistemically valuable about understanding and thus necessary for its attribution,⁷ regardless of who possesses it or what it is about. For instance, something which is invariably present in an account of understanding, be it as a symptom or mark, is the presence of abilities. We will soon argue why treating abilities as a mark rather than a symptom makes more sense. But first, we will consider some accounts that don't put abilities center stage. These can have some intuitive and/or philosophical plausibility, but also lead to troubles on both fronts,

⁶ There are exceptions, such as (Sierpinska 1990, 1994).

⁷ We believe a useful theory of understanding does not come apart from a justified theory of attribution, because it makes no sense to clarify understanding separately from (justifiable) understanding attributions. The proof of the pudding must be in the eating. This does entail that characterisation of understanding needs to be able to address contextual shifts in attribution, which we believe is possible (see Delarivière 2020b).

indicating that they are neither as plausible nor as intuitive as one might first assume.

3.1 Sense

The obviously most salient feature of understanding is the feeling or sense of understanding. The most notable of feelings related to understanding is the “aha-erlebnis,” connected to other feelings like surprise, familiarity or expectation (Lipton 2009). We can also find *confidence* in our newly-found epistemological powers, take *pleasure* in finding transparency and coherence where we thought there was opaqueness and disjoint, and we can take *satisfaction* from fulfilling a drive to find evidence that supports our theories, or forces us to modify them (Gopnik 1998). The sense or senses we associate with understanding can be spread out in time or hit us with a flash, a moment of eureka. This latter sense is so common to us that it sounds familiar to scientists and laypeople alike.

But appearances can be deceptive, and this sense faculty has been rightly criticised as an unsatisfying characterisation for understanding as an epistemological mark. Indeed, there are plenty of examples showing it cannot be a sufficient or necessary condition. It’s not difficult to think of cases of someone having genuine understanding without any accompanying feeling, and even easier to think of cases of someone genuinely experiencing that sense, but being simply mistaken. Trout (2002) has attributed this to a combination of hindsight bias (the observed effect of people systematically overestimating their past powers) and overconfidence bias (the observed effect of people systematically believing they are right even when they aren’t). Indeed, the feeling of understanding may be its most familiar and salient aspect, but it has been widely agreed that it is neither necessary nor sufficient for genuine understanding (Kvanvig 2003; de Regt and Dieks 2005; Wilkenfeld 2013b; Toon 2015; Ylikoski 2009). For this reason, no one defends it as the mark of understanding (unless to discredit understanding as a concept), and detaching understanding from its associated sense has actually been the start of taking the concept of understanding seriously.

This doesn’t entail that feelings are epistemologically irrelevant or mere epiphenomenal by-products. They may very well bear some epistemological relation to understanding other than as a mark. Even when neither necessary nor sufficient, their presence may significantly correlate with acquiring better understanding. In that case, sense could serve as indirect evidence or a strong heuristic in the pursuit of understanding. For instance, under the condition that one’s background beliefs are varied and true, the satisfaction of new information cohering with these background beliefs makes the sense of understanding to be a reliable indicator of understanding (Grimm 2011). Given that this feeling seems generally desirable, this may illuminate at least in part why we value and pursue understanding. Indeed, psychological payoffs are more local and immediate in the pursuit of understanding than some of the epistemological benefits (Lipton 2009). Gopnik (1998) has argued that this sense of fulfillment is due to a useful evolutionary drive for the theory-formation.

But the role of feelings could be more epistemically embedded still. Historically, people have been naive about what sort of qualities or processes underlie human intelligence, so it is conceivable we are underestimating the role of feelings in understanding. It is entirely possible that an agent that lacks such a psychological dimension is condemned to lack in understanding also. This would disprove de Regt's notion that a feeling of understanding "has no epistemic function" (de Regt 2009, p. 587).⁸ But this is mere speculation. Others have pointed out the epistemically misleading or inhibiting role of feelings, namely that feelings can be misleading due to overconfidence and hindsight bias (Trout 2002), or play a predominantly inhibiting role due to overestimating one's detail, coherence, or depth of understanding (Ylikoski 2009). To the extent that a psychological sense of understanding plays an epistemic role, it will merit inquiry in the context of epistemology, but clearly there is no consensus on whether it does. What does enjoy widespread agreement is that it is not the mark of understanding.

3.2 States

Philosophical conceptions of understanding invariably invoke the presence of mental states (be it as explanatory concept or trait) and abilities (be it as trait or symptom). So, two main lines of spelling out the trait of understanding are as either (a) appropriate mental states, potentially with abilities as symptoms (e.g. Zagzebski 2001; Grimm 2011, 2016; Van Camp 2014; Wilkenfeld 2013b) or (b) appropriate abilities, potentially with states as explanatory constructs (e.g. de Regt and Dieks 2005; Ylikoski 2009; Hills 2015; Delarivière and Van Kerkhove 2017). Siding with the latter, we will argue why considering abilities as a mere symptom of mental states would put the cart in front of the horse, creating a variety of needless problems. The argument, in short, is that this is because it is not mental states themselves that are empirically accessible or epistemically valuable to us. We both detect and judge mental states by the abilities, and not vice versa.

To start, it's worth differentiating a mental state from a physical state, such as a brain state. To our knowledge, no one defends physical states as the mark of understanding. It is quite obvious that those physical states themselves are neither what we value about understanding (e.g. we do not say "ultimately, it is the aim in science for our brains to reach state X") nor what we look for in evaluating someone's understanding (e.g. we do not say "Did you understand what I said? Let's get you in the brain scan and find out"). If we characterise understanding through mental constitution or occurrences (conscious or subconscious), then the successes of an understanding subject will need to come from the subject having something appropriate (such as mental representations) in her metaphorical mind's eye or other mind-organ, and the abilities that the subject may display will merely be the fortuitous public effects of those appropriate private occurrences. We can find some explicit or tacit endorsement of this view in characterisations of understanding (e.g. Hiebert and Carpenter 1992; Zagzebski 2001; Barby et al. 2007; Wilkenfeld

⁸ de Regt (2004) concedes that it can be "a source of motivation" (p. 104), but not an aim of science.

2013a; and arguably Grimm 2011, 2014). For instance, Wilkenfeld (2013a) characterises understanding as such:

A statement, attributed in context C, that thinker T understands object o, is true if and only if T possesses a mental representation R of o that T could (in counterfactuals salient in C) modify in small ways to produce R', where R' is a representation of o and possession of R' enables efficacious (according to standards relevant in C) inferences pertaining to, or manipulations, of o. (Wilkenfeld 2013a, pp. 1003–1004)

The mental representation here functions as the mark of understanding. This view fits with some of our everyday language about understanding and its justifications (e.g. seeing how things fit together, pointing to gaps or flaws in a mental image, abilities being treated as evidence rather than proof of understanding). We argue, however, that marking understanding via mental states leads to problems both practical and metaphysical.

Firstly, if understanding is characterised through the appropriate mental occurrences, then the task of gauging another person's understanding is not just difficult, it is impossible. We can't empirically discern what people see "in their mind's eye" or what "state their mind is in," except by inferring these through what they do (e.g. through what they say)—in which case the causal role of their internal imagery or mental state is an instrumental postulate. We cannot look into anyone's mind, private as it is. We would have no idea whether Euclid actually understood geometry, for only Euclid himself knew what went on in the privacy of his own mind. So perhaps our attributions of understanding are not really about these private occurrences. The only mental representations we have access to is our own, and even then, it depends on what we mean with "access" and "representation." Of course, Euclid can tell us about what goes on in his mind, but saying is a kind of doing and if his private occurrences are what mark his understanding (more than his doings or sayings), then we have absolutely no way to check whether he isn't subtly or unequivocally deceiving us. In fact, we don't even have any guarantee that he isn't mistaken himself. We'll return to this possibility soon. What subjects with understanding share (or overlap in) is their abilities, not their presumed states.

Maybe there is a way out of this conundrum by reinvoking physical states. Most philosophers and scientists agree (us included) that there is some relation between what the brain and what the mind are doing. If we'd want to be able to take advantage of that relation to detect or mark understanding, our bets would be safest with the identity theory of mind. According to this, a particular type of mental state (e.g. pain) corresponds with a particular type of physical state (e.g. C-fiber firing), as opposed to, for instance, a functionalist theory of mind (the dominant theory), which says that a particular type of mental state corresponds with a particular type of function (e.g. yelling ouch and withdrawing). If the identity theory of mind *would* turn out to be correct, and mental states *are* the mark of understanding, *then* we could be able to detect the appropriate mental states through their corresponding physical states. Unfortunately, we would still be at a loss at discerning whether someone understands anything at all until we figure out which physical states

correspond with the appropriate mental states. Furthermore, we would still need to figure out which mental states are valuable, which brings us to the second problem.

The second problem is that there's no clear way to decide which occurrences are appropriate by referring only to those occurrences themselves. What is the correct mental state to understand the irrationality of the square root of 2? This is not only hard to characterise, but the avenues in which we look to justify our choices are invariably motivated by the external affordances (i.e. abilities) granted, not by the intrinsic value of the mental occurrence itself. What a group of subjects with understanding share (or overlap in) is the appropriateness of their abilities, not the appropriateness of their presumed states. Calling the occurrence in one's mind's eye "a proof" would be vacuous if what is seen cannot be put to any public use. Conversely, if we found out that someone exhibiting extreme competence had mental states different from what we have hitherto characterised as appropriate, we would broaden the scope of appropriate mental states, not deny her with understanding.

This brings us back to the possibility of being wrong about one's own mind. All of us can have difficulty expressing what kind of mental representations we are operating with. And even when we think we can, our presumptions about the processes of our minds are not always an accurate description about the actual processes that underlie them. Studies in mental rotation, for instance, have shown that while we think we can rotate and compare 3D tetris-figures in our mind's eye, the competences we have do not line up with computer-systems that actually do what we think we do (Dennett 1993, c10). Confronted with such a situation, we are forced to rethink what we know about our own mental representations. This just goes to show that the kind of mental representations we infer is better marked by considering the competences they imbue than vice versa. The nature of the internal model is judged by external displays, not the other way around.

This doesn't imply that we must dismiss all claims made from a phenomenologically "internal" narrative. Our focus is on how people act, but how people act includes their speech acts, which includes speech acts about their own phenomenology. We must take these speech acts seriously (i.e. as genuine evidence of that person's abilities or lack thereof), but we're under no obligation to take their internal seeing-narrative metaphor literally (i.e. as correct claims about the mechanisms of the mind behind those abilities, the mental states).⁹

One way out of this problem would be to say that the appropriate occurrences are those of which the public effects are certain appropriate abilities, a route taken by Wilkenfeld (2013a), for instance:

[O]ne's understanding *x* consists in large part of representing *x* in the right sort of way (...) I will contend that "the right sort of way" is best cashed out as a mental representation the possession of which enables certain abilities (Wilkenfeld 2013a, p. 1002)

But then our understanding-attributions are actually decided by the abilities and not the occurrences which supposedly lie behind them. In short, it cannot get around

⁹ For an exposition on third person phenomenology or heterophenomenology, see (Dennett 1993).

the problem that “the correctness of the internal model [or representation] is judged by external displays of understanding, not the other way around.” (Ylikoski 2009, p. 103). We may as well put the horse in front of the carriage, where it can do its work unencumbered.

Having dispelled the mental state approach does not mean, however, that we cannot use the concept of mental states or representations as instrumental postulates (i.e. a proposed explanation that bind the observable input and output), but then they “are only hypotheses, models designed to explain, to sum up, what you observe” (Wittgenstein 1953, p. 62). Our dispelling is only meant to target mental representations as the mark of understanding, where understanding stands or falls based on what happens inside the mental sphere (i.e. what is in someone’s mind’s eye or other mental organ, conscious or not). In any case, there is still a lot of freedom in how precisely we conceptualise the mark of understanding, but at the very least we have now seen the pitfalls of an approach that puts its premium on mental states directly.

3.3 Synonyms

The literature makes frequent references to understanding involving “grasping” (e.g. Kvanvig 2003; Trout 2005; Khalifa 2013, Grimm 2011, 2014), or “seeing” (Zagzebski 2001; Riggs 2003), or “having” something. However, what exactly is being grasped, seen or had (be it a representation, a proposition, an explanation, a relationship, a belief, a body of information, ...) as well as what that entails (be it a sense, an attitude, an ability, access to, ...), varies widely. In a literal sense, the concept of “grasping” is about manually seizing (and possibly manipulating) something, and the concept of “seeing” is about a visual process that stretches from the eye to the brain. It might be readily conceded that “seeing” here is not primarily about the literal visual process performed between eye and brain, or even primarily about a corresponding mental process performed by the mind’s eye, and that “grasping” is not primarily about a literal seizing performed by hands, or primarily about corresponding mental gestures with the mind’s hands. Their usage could be metaphorical. But that leaves us with the question of the nature and value of the metaphors. Quite often, the work that the metaphor is supposed (and not supposed) to be doing is not made more explicit, which is unfortunate given that it is intended to lead us to a more explicit concept of “understanding” (Gordon n.d.). We’ll now consider some of these metaphors as disguised synonyms and the obscurification or pitfalls they can lead to.

Zagzebski (2001) says understanding involves “seeing how the parts of that body of knowledge fit together, where the fitting together is not itself propositional in form” (p. 244), and that it involves “mental representations” (p. 241) which she thinks “will likely include such things as maps, graphs, diagrams, and three-dimensional models in addition to, or even in place of, the acceptance of a series of propositions.” (ibidem), but it’s quite unclear how literal she takes “seeing” and how phenomenological she takes mental representations to be. Similarly, Kvanvig (2003) uses the term “grasping,” directs the grasping to coherence relations, and

qualifies it as “internal” (p. 192), but never spells out what grasping involves, so it’s hard to know what it is literally doing or metaphorically pointing to (Gordon n.d.).

Grimm (2011, 2014) starts in a similar place, arguing that understanding (which is the same as non-propositional knowledge of causes) consists of “seeing” or “grasping” a modal relationship (as opposed to a proposition). The verb “seeing” is explicitly acknowledged as a metaphor, but what work is it doing here? He is more explicit:

What the metaphor of “seeing” seems to involve, then, is something like an apprehension of how things stand in modal space (...) Just as, in seeing with one’s eyes, one takes in or apprehends how things stand in the physical terrain, so too the basic idea here seems to be that in “seeing” with the eye of the mind, one takes in or apprehends how things stand in the modal terrain: one apprehends what cannot be otherwise, or how certain changes will lead, or fail to lead, to other changes. (Grimm 2014, p. 334)

The “seeing” here seems to be a bodily metaphor for a correlate in the mind’s eye. But if one can passively witness changes in physical terrain, then can’t one also witness how certain changes lead to other changes just as passively? What makes subject A understand better than subject B, is not that subject A has a mental sensation of a relationship that subject B doesn’t have, but that (s)he can exploit such a relation.¹⁰ Furthermore, if “seeing a relation” is read as synonymous to “exploiting that relation,” we are already moving away from putting representations center-stage in favor of abilities. The passivity is addressed elsewhere, where he points out the manipulist nature of this “grasping.”

[In the manipulist] sense, mentally to grasp (...) a structure would therefore seem to bring into play something like a modal sense or ability—that is, an ability not just to register how things are, but also an ability to anticipate how certain elements of the system would behave, were other elements different in one way or another. (Grimm 2011, p. 89)

We now have “modal ability,” consisting not merely of “apprehending” or “registering,” but also “anticipating.” What is this modal ability?

On our proposal, “seeing” or “grasping” would count as a kind of ability, because the person who sees or grasps [modal relations] will characteristically have the ability to answer a variety of what James Woodward (2003) has called “What if things were different?” questions. (Grimm 2014, p. 339)

Is the modal ability the ability to answer what if questions (it doesn’t seem that way) or is this ability a symptom of (or characteristically related to) another modal ability that occurs inside the mental realm? In trying to demarcate understanding, it’s unclear what work the occurrence inside the mental realm is doing and what work is done by the ability to answer what if questions. It’s not entirely clear whether we’ve made the concept of understanding more explicit or the metaphor more complex.

¹⁰ This is also the reason that de Regt (2009) invokes a skill condition.

Hills (2009, 2015) uses the metaphor of “grasping,” calling the metaphor “an extremely important one” (2015, p. 4), but also “not very clear” (ibidem). She explains the meaning in similar to Grimm:

When you grasp a relationship between two propositions, you have that relationship under your control. You can manipulate it. You have a set of abilities or [intellectual] know-how relevant to it [incl. intellectual know-how], which you can exercise if you choose. (Hills 2015, p. 4)

What that intellectual know-how entails, she spells out as concrete abilities (which we’ll spell out later), although she doesn’t say whether they mark understanding or are merely a symptom of it. She explicitly leaves it open: “Is understanding why (partly) constituted by these abilities? Or is it the ground of these abilities? I favour the former account, but it is one of the questions that I will leave open here” (Hills 2015, p. 5).

Khalifa (2013) argues that grasping an explanation is central to understanding and necessarily entails true beliefs (of the form q explains p) that must be the result of exercising reliable cognitive abilities, which will involve evaluating (or discriminating between) explanations. This needs further explication, by describing what it means to have a belief and what exactly is involved in evaluating an explanation, which is related to the question of what makes it a cognitive ability. What sort of modifier is “cognitive” in “cognitive ability”? Does it specify abilities as (1) an ability that can be categorised as primarily cognitive (as opposed to, say, tennis, which can be categorised as a primarily physical ability), as (2) produced by a cognitive entity (as opposed to provided by the environment), or as (3) private performances taking place within the secret realm of the mental (as opposed to potentially¹¹ taking place as public acts)? Wilkenfeld (2013a) could be attributed with the third, for he says; “Understanding is a cognitive achievement (...) [which] must instead consist of an ability to manipulate some mental correlate of the understood object” (p. 1003). Pritchard (2014) is not as clear on this point. He explains that an archer has an achievement (hitting the bullseye) if she exercised her ability (archery ability) and it was a success (hitting it) mainly due to that exercise (as opposed to due to other factors). So it also goes with the cognitive: someone has a cognitive achievement (understanding) if she exercises her cognitive ability (reliably forming a true belief) and it was a cognitive success (a true belief was formed) due to that exercise (and not, for instance, by trusting the testimony of an expert). Pritchard’s use of the term evokes the first kind (it distinguishes the cognitive ability to form true beliefs from the physical ability to hit a bullseye) and the second (it serves to distinguish the credit of the success due to the agent’s ability or something external), but it is unclear whether he supports the third - because the ability of having or forming a true “belief” is itself in need of clarification. Greco (2007), who uses the term “intellectual ability” in a similar way, clarifies the nature of “ability” as “dispositional properties that display a characteristic structure” (p.

¹¹ Just to be clear: no one is denying that people can keep their abilities to themselves, but this is a radically different thing from claiming that their abilities are inherently *located* in the realm of the private.

68), which clearly places abilities outside of the mind, while Pritchard's account shifts the issue to the nature of belief.

It seems uncontroversial to suggest that understanding involves beliefs. Most philosophers make reference to the notion of belief in some way (see e.g. Grimm 2011; Khalifa 2013; Kvanvig 2003; Hills 2009; Pritchard 2014). But belief is not quite such an uncontroversial term itself. It's far from obvious that beliefs exist at all, and if they do, it's far from obvious what interpretation is the most useful one (see Dennett 1990, c5 for a rundown on the interpretation of the term and their problems). Once again, we are faced with a term that, in the context of characterising understanding, moves the target rather than marking it. Is a belief a mental state, as is suggested by the traditional notion of propositional attitudes (which is still very open for interpretation as to what sort of entity a proposition is, where we find it, and how we take an attitude towards it) or a mental representation (with the problems that come with it) (Schwitzgebel 2019), or does the notion of belief involve a claim about performance, as is suggested by the dispositional or interpretationist reading? As we see it, in case of the latter, an ability approach (which also rests on performance) might have saved us a lot of conceptual legwork, and provide us with a more steady target.

4 The Mark Question II: Defending an Ability Account

As suggested, an alternative proposal to mark understanding is to place a premium on the presence of abilities (e.g. Avigad 2008 for mathematical understanding; de Regt and Dieks 2005; Hills 2009 and Ylikoski 2009 for scientific understanding), and this is the approach we will be defending as the most sensible and fruitful one. We will start by roughly characterising our approach to understanding:

'S understands X' corresponds to 'S possesses sufficient abilities appropriate to X in context C.'

There is still a lot to unpack here. Every word in this characterisation requires further elaboration: which "abilities" are appropriate, how many are "sufficient," how one can "possess" them and which types of entities "S" can do so, as well as what abilities are "appropriate," how to discern them and which role the "context" of the attributor and the circumstances of the subject do play in this. For the purpose of defending the ability-approach as a mark of understanding, in this paper we will focus on clarifications as they pertain to demarcating understanding. What is of prime importance here is that if we characterise understanding through abilities, then the successes of an understanding attribution will need to come from the subject's appropriate acts or performances. Conceptualising such an ability account will take us beyond the mental (which will force us to focus on performances, with the virtues that that entails), beyond single performances (which will force us to consider the concept of ability and we'll conceptualise as multi-track behavioural profiles), and even beyond a single ability (which will force us to consider the notion of the appropriate set of abilities and which we'll conceptualising as an appropriate behavioural profile).

4.1 Benefits of an Ability Account

Under the ability account, attributions of understanding stand or fall with how a subject acts (i.e. performs). As will become clear, this does not involve reducing everything to stimulus-response descriptions, but it does involve that any appropriate attribution of understanding bottoms out in claims about potential outward performance, and not anything that lies behind those performances. There are a number of benefits or virtues to this approach, and we will go over some of them here.

Firstly, we are sidelining (but not discarding) the role of feelings we associate with understanding without necessarily discarding them as epistemologically irrelevant (see 3.1). When we are specifying what marks understanding, we can refer to appropriate abilities without reference to the emotions or sensations that accompany, motivate or implement them. At the same time, this does not discard the role of feelings outside of the purview of epistemology. To the extent that feelings do play a role in the implementation of a subject with abilities, either as a goal in the subject's pursuit of them or a guide in the process of attaining them, they are also (with varying degrees) relevant to epistemology.

Secondly, by putting a premium on abilities, we are also avoiding some of the other problems that plagued the accounts we discussed earlier (see 3.2). If a subject's mental states are inferred by the acts of a subject (including speech acts about the subject's own phenomenology), then those mental states can only be discerned indirectly, if at all. But a subject's abilities are literally comprised of such acts, which means that the acts they display serve as direct (though incomplete) evidence of a subject's abilities. Next, it is also easier to pinpoint the appropriate abilities because we can refer directly to the appropriate acts themselves, as opposed to looking for the acts which correlate with the appropriate mental state. Additionally, determining which abilities are appropriate is more to the point because abilities can have intrinsic values (e.g. we value prediction for its own reward). Furthermore, justifications of understanding tend to boil down to abilities when asked to justify them (even if they are initially phrased in the mental state narrative). For instance:

Suppose I tell you that my friend Paolo understands group theory, and you ask me to explain what I mean. In response, I may note that Paolo can state the definition of a group and provide some examples; that he can recognize the additive group structure of the integers, and characterize all the subgroups; that he knows Lagrange's theorem, and can use it to show that the order of any element of a finite group divides the order of the group; that he knows what a normal subgroup is, and can form a quotient group and work with it appropriately; that he can list all the finite groups of order less than 12, up to isomorphism; that he can solve all the exercises in an elementary textbook; and so on. What is salient in this example is that I am clarifying my initial ascription of understanding by specifying some of the abilities that I take such an understanding to encompass. On reflection, we see that this example is

typical: when we talk informally about understanding, we are invariably talking about the ability, or a capacity, to do something. (Avigad 2008, p. 321)

Another benefit of the ability approach is that the notion of tacit understanding is given more room to flourish. People have been shown to respond appropriately to situations without being able to articulate (or even register) what they are doing. For instance: people are able to detect subtle variations in human appearance and behaviour without being able to tell what they detected, people are able to extract temporal patterns that underlie a train of events without even having noticed they did so, people are able to learn without understanding how they learn, or even knowing that they have learned (i.e. implicit learning) (Reber 1989). And people are able to form judgements about the quality of an explanation without knowing how they do so (Lipton 2009). Here “the contrast between our ability to make [these judgements] and our inability to describe them on the basis on which we make them is particularly stark” (p. 60). Such epistemic competences have been called tacit. Whether they deserve the label of understanding (or knowledge) is where internalists and externalists (be it about knowledge or understanding) differ. Externalists believe reliable competence is enough, but internalists claim that one cannot understand unless one can also articulate, justify or explain one’s understanding. Zagzebski (2001) explicitly makes the latter claim:

Understanding, (...) is a state that is constituted by a type of conscious transparency. It may be possible to know without knowing that one knows, but it is impossible to understand without understanding that one understands. (...) understanding is a state in which I am directly aware of the object of my understanding, and conscious transparency is a criterion for understanding. (Zagzebski 2001, p. 246)

Under a state account, this seems intuitive, because how would we know there was a “mental state” if its presence or constitution wasn’t in some way “accessible” to the subject.¹² Nonetheless, conscious transparency seems too strong a requirement for understanding. Mathematical competences, for instance, are valuable whether they are consciously deliberated or subconsciously brooded. Pritchard (2009) also thinks Zagzebski takes the transparency condition too far, but agrees with her that when one has understanding, as opposed to knowledge, “it should not be opaque to one that one has this understanding—in particular, one should have good reflectively accessible grounds in support of the relevant beliefs that undergird that understanding” (p. 39). It’s not clear exactly why this should be so, nor that it should be so. The suggestion seems to be that understanding, unlike knowledge, requires more competence, and the internalist condition is supposed to ensure this. But the internalist condition thus also discredits all other forms of competences, which is unfortunate. Grimm (2016) argues against “articulacy” as a necessary condition for understanding by appealing to understanding in children and animals. We can make similar points about experts. While mathematicians can prove and recognise a proof when they see one, it is notoriously difficult to say what it is that

¹² Unless it is inferred, in which case the premium once again shifts to abilities.

experts do while proving or what it is that they recognise when they see a proof.¹³ Consider this quote by the late mathematician Reuben Hersh:

When you're a student, professors and books claim to prove things. But they don't know what's meant by 'prove'. You have to catch on. Watch what the professor does, then do the same thing. Then you become a professor, and pass on the same 'know-how' without 'know what' that your professor taught you. (Hersh 1997, p. 50)

While this may be cause to say that understanding of proof is lacking or incomplete, we'd have a strange account of understanding if it entails that mathematicians don't understand the concept of proof. Under an ability account, it becomes clearer that the absence of the ability to articulate or explain one's understanding does not have to discredit other abilities; rather it marks a lower degree of understanding. Furthermore, what is missing is yet more valuable abilities, namely the ability to articulate what you understand (and other related abilities that may rely on that ability). This distinction may be clarified with the aid of terms like tacit versus theoretical understanding (see Ylikoski 2009) or implicit versus explicit understanding (see Hills 2015).

The ability account also sidesteps the problem of excessive or infinite encoding. Allow us to explain. If one subscribes to the idea that there is a grasping relation (e.g. an attitude) towards an object of understanding (e.g. propositions), then a complete understanding will, for instance, involve an attitude towards all the relevant propositions, "the ideal understanding text" (Van Camp 2014, p. 108), which "consists of a complete framework of all possible propositions about the phenomena, including their various relationships" (ibidem). This ideal understanding text may grow exponentially large with each relation, or infinitely large due to Carrollesque recursion,¹⁴ placing some strain on the physical or mental space that needs to encode it all. The ability account, by contrast, doesn't require explicit encoding of all the relevant propositions that can be constructed, as long as one can keep responding appropriately.

There are further issues with the idea of encoding states that the ability account avoids. How specific or compartmentalised does a state need to be for understanding to be attributable? Once someone understands, is there a mental or physical state which is permanently present as long as the person can be said to understand? Does that person lose the understanding as soon as she's no longer in the previous state? Is there a compartment of that person's whole mental or physical state that covers the particular understanding (and how many state-slots can there be before someone's whole state is "full")? Or is the state transient and do you stop understanding as soon as the mental or physical state has changed? Or do we have to refer to potential states (making it even more difficult to ascertain whether someone understands or not)? What it means to possess a mental representation or have access to the state is just another extra nasty question we can do without. What subjects with understanding share (or overlap in) is their abilities, not their states.

¹³ See also Thurston (1998) and Davis and Hersh (1998).

¹⁴ See Lewis Carroll's (1895) *What the Tortoise Said to Achilles*.

In spite of its many virtues, there are still some causes for skepticism towards the ability account, and these do need to be addressed. E.g., one may think that it has no way to talk about competence without performance (i.e. dormant competences), or one may object that a single performance is no guarantee for true competence. Both of these worries, and more, we think can be addressed (which we'll not be doing here in detail), but not without further developing the notion of "ability," so let's.

4.2 Beyond an Act and Ability

The ability account, we have already argued, demarcates understanding based on how a subject performs or acts, and whether those acts can be seen as appropriately successful. This is true for all abilities, not just epistemic ones. If Hamlet claims he is able to tell a hawk from a handsaw, then the validity of that claim rests on whether Hamlet correctly distinguishes one from the other. Likewise, whether a meteorologist is able to predict the weather stands or falls with how she fares in the predictions she makes. One reason why it may still seem appealing to refer to private occurrences beyond a person's performance is because it seems plausible for the same act to be performed with and without understanding, e.g. by sheer luck, by rote memorisation, or by blind rule-following. This seems to imply that the difference-maker for understanding lies not in the performance, but in something beyond it. But we may concede this without having to withdraw into a secret world.

Firstly, let us consider how our account goes beyond an act. "ability" is a modal predicate, concerned not just with how things are or have been, but how things could be. Breakable glass may never shatter, a meteorologist may never predict the weather and Hamlet may get through the entire play without ever distinguishing a hawk from a handsaw. And yet this alone is not sufficient to deny any of them with being able to shatter,¹⁵ predict the weather or distinguish a hawk from a handsaw. We can address this by opening up the range of circumstances under which one would perform appropriately.

[A]bilities in general are functions of success in relevantly close possible worlds. In other words, to say that someone has an ability to is to say that she would be successful in achieving X in a range of situations relatively similar to those in which she typically finds herself. (Greco 2000, p. 13)

So, it seems abilities involve successful performance under a range of circumstances, regardless of whether these circumstances actually obtain. Crucially, the assessment of a modal ability is not one where we literally peer into counterfactual worlds. Counterfactual worlds are not worlds we can discern with the naked eye. They are merely useful conceptual tools for us to make claims (i.e. explanations or generalised predictions) about the circumstances under which we expect the acts to be present. Those claims are warranted by their explanatory or predictive power with regards to the subject's actual acts. As such, counterfactuals are more instrumental than metaphysical claims. The task of assessing

¹⁵ One may wish to object that glass does not have the "ability" to shatter, but is merely disposed to. We'll address this concern later in this section.

understanding, then, becomes the task of identifying factual acts, and evidence of counterfactual ones. On the basis of these, we can derive generalised claims and predictions about the circumstances under which the appropriate acts would occur, or the defeaters that would prevent them. This means that when we ascribe someone as having an ability, we are using evidence of past and present performance to make an estimation of the quality and range of those performances under the salient circumstances. Making a lucky guess may be something one gets away with under the precise circumstances under which you got lucky, but luck quickly runs out if you are tested under different circumstances. The problem with sheer luck is not the success under a singular circumstance, but the range of failings under others. Similarly, masked abilities (e.g. placing tape over someone’s mouth) may be something which prevents the appropriate acts from obtaining, but what warrants the understanding is that the subject, without tape over their mouth, would or will be successful. Therefore, what marks the subject’s masked understanding is not to be found beyond its abilities, but beyond its mask.

There is a worry¹⁶ that invoking modal claims puts our ability-approach on the same footing as mental state accounts. While it is true that both the ability account and mental state account can be accused of going “beyond observation” in a way that limits the justification of our understanding attributions, their limits are not equivalent. Firstly, the ability account allows us to observe constituents of what marks understanding, namely appropriate acts. The mental state account, by contrast, limits our access even to its constituents. Secondly, it is readily acknowledged as part of the ability account that, on the basis of observed acts, we make an estimation of the range and scope of understanding which we, crucially, would have to revise or explain when presented with further instances of acts (or lack thereof). So if we are confronted with a lack of appropriate acts, we would either need to explain why the appropriate acts were lacking, but may be expected elsewhere (be it under different circumstances or in different challenges), or be forced to adjust our estimation of the (degree) of understanding¹⁷—as we would expect from a mark. Conversely, any instance of an appropriate act warrants a degree of success, and any insufficiency in understanding can only be pointed out by estimating or discovering the lack of acts (be it under different circumstances or regarding different challenges)—as we would expect from a mark. (See Delarivière 2020b, c3 for an extensive analysis of such cases). Mental states, by contrast, can only be revised or explained through acts, which effectively makes them instrumental constructs based acts—but that’s just an expansion of the ability account without acknowledging it.¹⁸

Secondly, let’s consider how our account entails going beyond an ability. The use of the plural in “abilities” is not incidental. Understanding involves more than just having a single ability appropriate to the object of understanding. I can easily memorise a correct response to a certain question (or even a few of them) without

¹⁶ Thank you to the reviewers for pointing out this worry. .

¹⁷ For more on the degrees of understanding, see (Delarivière 2020b, c2).

¹⁸ For such an expansion of the ability account where we use mental properties as instrumental concepts, see (Delarivière 2020b, c4).

understanding what it is that I'm saying. The problem with such a single-track ability is not that it should be discredited, but that the scope of abilities is too narrow, way too narrow. Answering questions according to a set script may be something one can get away with in tests that happen to only involve the memorised questions (giving a misleadingly good impression), but this will quickly fail once tested for related abilities or other appropriate abilities (e.g. correct courses after a setback, explain in different words, apply it in a practical circumstance, make an analogy, instruct others, criticise incorrect practice, predict the outcome of observed lapses, answer what-if-things-had-been different questions,...).¹⁹ It is not the narrow success of rote memorisation, but the wide failing it entails that makes for a poor understanding. The same is true of blind rule-following, i.e. what Skemp (1976) calls "rules without reason" (p. 20). But it's not accidental that the adjective "blind" seems equally fitting. It's because any deviation or extension beyond where the rule leads, would leave the subject in the dark. It's not the narrow success of rule-following that is the problem, but the wide failing it entails. If we draw the scope of understanding to include all the relevant abilities, then it becomes increasingly difficult to motivate why a subject displaying the wide range of relevant abilities would not merit an understanding-attribution.

So far, we've only been talking about performance in the most neutral sense, as a happening or behaviour. When we think of abilities, however, we're not thinking of glass having the ability to break in the same way that a meteorologist has the ability to predict the weather. In the latter, we're thinking of a person having the skill to do something (i.e. the successful performance) and the option to engage that skill (i.e. the power or freedom to act). She will employ the skill only when she chooses to (Moore 1912) or tries to (Fara 2008). This may provide an appealing reason to retreat behind the act and look for the right kind of state that makes it a power over a mere reflex. However, the difference between a reflex and a power lies not *behind* the act, but *beyond* it. The common intuition, for instance, is that glass is disposed to break, but doesn't thereby have the power to break. While we can appreciate the distinction in this case, must we therefore also say that the ability to calculate is different for humans than it is for calculators? Furthermore, does that distinction really matter in itself, or is the distinction related to the valuable differences in performance between a mathematician and a calculator? When we are interviewing someone for a job based on her abilities, we are not interested in whether she is "able to X" in the sense that her doing X was preceded by some appropriate mental

¹⁹ Where one ability begins and another ends will largely depend on how wide a net one is attempting to cast. The ability "to multiply two single digit numbers" clearly casts a wider net than the ability "to multiply 5 with 6." Some abilities don't have an obvious net-size. Consider the ability "to prove that the square root of 2 is irrational." Is this ability composed of the act of giving a proof in the salient circumstances along with other acts (e.g. criticizing incorrect steps, predicting outcomes of observed lapses, ...) or do these other acts constitute different abilities? Providing an answer is not only difficult, but will largely depend on how you phrase the ability in question. We will not concern ourselves here with providing universal demarcation criteria for what constitutes a single ability. As long as one is with us in the claim that understanding casts a wide net, it doesn't matter to us whether it's because understanding captures many multiple abilities, or because abilities capture many acts.

state (or worse, some libertarian kind of non-determined free will),²⁰ we are interested in whether she will do X with the desired level of sensitivity and when it is appropriate. The same holds for doing something conscientiously. We are not peering behind the act to see whether it was performed with the appropriate mental state, but looking beyond the act, to see whether it was performed with the required level of sensitivity and quality. Acting conscientiously is not doing two kinds of things: being conscientious and acting, it's doing one kind of thing (acting) well (Ryle 1949/2000). For this reason, we won't distinguish between acts in the dispositional sense and acts in the power sense, because we believe they are distinct in degree, not kind.

Summing up, when we attribute people with understanding, we are not making untestable inferences to any secret phenomena which are forever out of our reach and judgement, but we are gauging the appropriate multi-track behavioural profiles of an understanding subject. So it is true that our understanding-attributions go beyond single acts, but this is not going beyond as in going behind them (to occurrences which are impractical or impossible to discern and don't themselves contribute anything of value), but beyond as in considering what people could and would do, namely their abilities, which we can discern, even if never fully (Ryle 1949/2000). We may, instrumentally, speak of mental states, but we do so on the basis of abilities, and not the other way around. Hence, it's the abilities that actually mark the understanding.

4.3 Brands of Abilities

Evidently, not just any ability is relevant to any understanding. The ability to make a good cup of tea is not indicative of understanding the general theory of relativity. Which brands of abilities are "appropriate" to understanding has a variety of candidates, depending on who you ask, what kind of understanding they are meant to capture, and what the object or field of understanding is. In this subsection, we'll be discussing several of the candidate brands or kinds of abilities offered up in the literature. Let's start with considering a popular distinction between three general kinds of understanding attributions:

- (1) Lindsay understands that there is a housing crisis.
- (2) Olly understands why there is a housing crisis.
- (3) Natalie understands the housing crisis.

These three kinds of attributions have often been distinguished in the literature, beginning with Kvanvig (2003), as (1) understanding-that or propositional understanding, (2) understanding-why or atomistic understanding, and (3) objectual understanding, respectively. Although the distinction is widespread in use, it's not without its issues. The first kind of attribution, propositional understanding, has been criticized (e.g. Pritchard 2010; Gordon 2012; Grimm 2016) as being indistinguishable from either attributions of propositional knowledge (if it only

²⁰ For a defense of free will that is compatible with determinism and does not rely on an inaccessible mental realm, see Delarivière (2016).

involves knowing a single proposition) or atomistic or objectual understanding (if it involves more). Furthermore, Grimm (2016) believes the distinction between the second and the third understanding is overstated, and largely due to a difference in focus or scope, rather than kind.²¹ We're inclined to agree. Nonetheless, it's worth bearing in mind which focus or scope an author has in mind when offering up the necessary abilities.

Even though it is unclear whether they should count as symptom or trait, Grimm's (2011, 2014) suggestion for the crucial ability involved in atomistic understanding (and thus also for objectual understanding) is that of being able to answer what-if-things-had-been-different questions (based on Woodward 2003). This suggestion is a powerful one that enjoys a lot of support. Hills (2009, 2015), focusing on understanding-why in particular, expands on Grimm. Although she left open whether abilities mark understanding or are mere symptoms of it, she does emphasise the ability to "to treat q as the reason why p , not merely believe or know that q is the reason why p " and spell out what that entails with a concrete list of abilities that cover not just answering what-if-things-had-been-different questions, but also a series of explanatory abilities, namely:

- (i) follow some explanation of why p given by someone else
- (ii) explain why p in your own words
- (iii) draw the conclusion that p (or that probably p) from the information that q
- (iv) draw the conclusion that p' (or that probably p') from the information that q'
(where p' and q' are similar to but not identical to p and q)
- (v) given the information that p , give the right explanation, q ;
- (vi) given the information that p' , give the right explanation, q' (Hills 2015, pp. 4–5)

Kuorikoski and Ylikoski (2015) and Ylikoski (2014), also take inspiration from Woodward's what-if-things-had-been-different questions, but they put the abilities front and center, calling them counterfactual inferences. They include instances of predictions, control and explanations under what-if circumstances. Wilkenfeld (2013b) pushes back on this characterisation by pointing out that counterfactual inferences do not help us with necessary truths, as is the case in mathematics, where "all (or at least almost all) of what one can understand involves necessary truths and the relations they bear to each other, and so there is no counterfactual dependence involved at all" (p. 101). This may be slightly too strong a pushback. There are multiple answers to "what if"-questions that don't rely on one to bend necessary truths. For example, in the case of the irrationality of the square root of 2, one can ask: what about the square root of 3? See Frans and Weber (2014) for more in depth examples of what-if-things-had-been-different questions and answers in mathematics. But the point may be taken that answering what-if-things-had-been-different questions or counterfactual inferences may not satisfy all needs.

²¹ "For both Kvanvig and Pritchard, "objectual" or "holistic" understanding has to do with our grasp of large chunks of information, especially as they relate to topics or subject matters. Understanding-why or atomistic understanding, by contrast, is focused on some particular state of affairs: understanding why the cup spilled, for example, or why Fred did poorly on his exam." (Grimm 2016, p. 254).

Speaking of mathematical understanding: Avigad (2008), focusing on mathematical understanding, takes what he calls a functionalist approach to understanding. Taking lessons from Wittgenstein, this involves characterising the abilities involved in understanding. For proof, he offers the following list of (kinds of) abilities involved:

- the ability to respond to challenges as to the correctness of the proof, and fill in details and justify inferences at a skeptic's request;
- the ability to give a high-level outline, or overview of the proof;
- the ability to cast the proof in different terms, say, eliminating or adding abstract terminology;
- the ability to indicate 'key' or novel points in the argument, and separate them from the steps that are 'straightforward';
- the ability to 'motivate' the proof, that is, to explain why certain steps are natural, or to be expected;
- the ability to give natural examples of the various phenomena described in the proof;
- the ability to indicate where in the proof certain of the theorem's hypotheses are needed, and, perhaps, to provide counterexamples that show what goes wrong when various hypotheses are omitted;
- the ability to view the proof in terms of a parallel development, for example, as a generalization or adaptation of a well-known proof of a simpler theorem;
- the ability to offer generalizations, or to suggest an interesting weakening of the conclusion that can be obtained with a corresponding weakening of the hypotheses;
- the ability to calculate a particular quantity, or to provide an explicit description of an object, whose existence is guaranteed by the theorem;
- the ability to provide a diagram representing some of the data in the proof, or to relate the proof to a particular diagram; And so on (Avigad 2008, pp. 327–328)

But the most extensive ability account (and most extensive understanding account in the literature of the philosophy of science, full stop) is the contextual account of scientific understanding provided by de Regt and Dieks (2005). de Regt starts by criticising the inadequacies of beliefs or theories if they can't be put to some use. As such, he places a skill condition on understanding. His focus is on scientific understanding, so his criteria involve explanations on the basis of a scientific theory. As criterion for understanding a phenomenon (CUP) scientifically, he offers:

CUP: "A phenomenon P is understood scientifically if and only if there is an explanation of P that is based on an intelligible theory T and conforms to the basic epistemic values of empirical adequacy and internal consistency." (de Regt 2009, p. 92)

Note that this criterion is not a pragmatic one, for it makes no reference to a subject that understands. The pragmatic element lies in another term, namely the "intelligibility" of a theory. For this, he offers a sufficiency criterion (at least for

disciplines which formulate theories in mathematical terms, such as the physical sciences):

CIT: “A scientific theory T (in one or more of its representations) is intelligible for scientists (in context C) if they can recognize qualitatively characteristic consequences of T without performing exact calculations.” (de Regt 2009, p. 102)

In short, one could say that CIT involves what it means to understand a theory, whereas CUP involves what it means to say a theory is scientific and can be understood.²² It may seem inviting to read that what understanding involves in his account is to “have an explanation,” because CUP requires merely that an explanation exists and de Regt doesn’t quite specify the link between the explanation and its use in understanding. However, he does say that a phenomenon is understood (scientifically) if one understands a scientific theory, and that understanding a scientific theory involves skill and abilities. This steers us clear of the notion that understanding would involve owning a scientific explanatory text or a scientifically accurate mental representation. Together, CUP and CIT suggest that phenomena can be understood by a scientist if there exists an explanation based on an empirically adequate and internally consistent theory of which the scientist is able to recognize qualitatively characteristic consequences without performing exact calculations. Wilkenfeld (2013b) gives a more technical interpretation of how to combine de Regt’s two criteria:

UD2: A phenomenon P is understood scientifically if there exists a scientist S in context C such that S (in C) can explain P with T and S in C can recognize qualitatively characteristic consequences of T without performing exact calculations and the explanation of P by T meets accepted logical and empirical criteria. (Wilkenfeld 2013b, p. 98)

A virtue of this particular ability account is that it allows intelligibility standards to vary along with the history of science and theories of scientific explanation. It leaves open “empirical adequacy,” which allows it to vary with the progress (or history) of scientific methodology and the variation of methodologies that come in different scientific disciplines. de Regt’s account has however not been without criticism. His theory requirement in particular has been attacked as too strong because we don’t always need a theory to understand (Kelp 2015). de Regt’s focus, however, doesn’t seem to be on a broader kind of understanding. His focus is meant to be on scientific understanding, not everyday understanding of science. de Regt’s focus on theories does constrain his account, but this constraint seems fair given that he’s not focused on attributions of individual scientists, but the progress of the sciences as a theory-generating discipline (even while not betraying the pragmatic nature of understanding). When talking about understanding in a broader sense, one would indeed like to cast a wider net.

²² Personally, we would read CUP not as a criterion for understanding, but as a criterion for scientific adequacy. CIT, on the other hand, seems to us the true criterion of understanding: namely that the intelligibility of a theory grants us an understanding of the phenomena, provided the theory in question is scientifically adequate (passes CUP).

Newman (2012) objects along similar lines as Kelp, saying that even if we have a theory, we don't always have the highly theoretical skills to use one, and that if it's only about qualitative consequences, then it's not clear what makes it scientific understanding (Newman 2012). We think he's reading de Regt either unfairly narrow or unfairly broad, without allowing anything in between. That being said, we have our own criticism of de Regt related to his theory-focus. In short, we have issues with "characteristic consequences" because he adds "without exact calculation" even though there is nothing particularly inappropriate about exact calculation. Presumably "without exact calculation" is added because one could calculate without knowing what one is doing, but the problem here, we think, is not in the exactness. Nevertheless, even if exact calculation is contestable in its appropriateness, recognising qualitatively characteristic consequences is not.

In sum, there are multiple candidates for multiple kinds of abilities in multiple fields, including: recognising qualitatively characteristic consequences (de Regt and Dieks 2005), making counterfactual inferences in the contexts of manipulation, prediction and explanation (Ylikoski 2009), give explanations (i.e. answering explanation-seeking questions) (Ylikoski 2009), answering what-if-things-had-been-different questions (Woodward 2003; Grimm 2014), being able to evaluate explanations (Khalifa 2013), relating knowledge to other knowledge (Van Camp 2014), controlling the phenomenon (Ylikoski 2009), specifying causal dependence (Ylikoski 2014), reliably tracking dependency relations (Grimm 2016), following an explanation given by someone else, explaining it in your own words, drawing the appropriate conclusion based on a given set-up, or vice versa (Hills 2009), responding to challenges as to the correctness of a proof (or theory, or label), identifying key features, identifying the nature of the objects and questions, mustering the relevant background knowledge, exploring the space of possibilities fruitfully, and so on (Avigad 2008).

Which (type of) candidate is the correct one, then? Our answer is that there is no single exhaustive one. We believe the scope of understanding to be quite wide (i.e. composed of, but not exhausted by any single set of candidates discussed in the literature) and contextually dependent on the many aims and values of the epistemic practice to which the understanding-attributions relate. Each field has its own objects of understanding, each with their own needs for what they are supposed to satisfy (needs which may, furthermore, change over time). Each kind of understanding (be it for objectual, atomistic or propositional understanding, or otherwise) has its own focus or scope. While the commonality here is abilities, it's difficult to find a clear characterisation of a single type of ability that covers all the needs of the fields of epistemology and kinds of understanding previously mentioned. Therefore, what we propose is that we keep constant that which stays constant and then we can look for systematic ways to talk about what varies. When it comes to the mark of understanding, however, there is but one sure commonality and that's the presence of the appropriate abilities.

5 Conclusion

In this paper, we focused on the *mark* of understanding, namely which systematic trait we find so philosophically or epistemically valuable about understanding and thus necessary for its attribution, regardless of who (i.e. which subject) it is attributed to or what the understanding is about (i.e. the object of understanding). We called this the “mark of understanding,” because it is what demarcates it. This mark of understanding needs a philosophically coherent and explanatory characterisation that can be applied consistently to various human subjects and across various objects with varying degrees of (contextual) quality. Furthermore, it needs to allow us to deal with the known philosophical problems of marks. With inspiration drawn from Ryle (1949/2000), we have argued that understanding-attributions always boil down to a particular set of appropriate abilities (of a subject), composed of acts (salient to the object for a certain context), and that this is the most coherent and useful conceptualisation of “understanding.”

One of the notable benefits of the ability approach is that we side-line, without discarding, the mistrusted role of feelings. We avoid some of the problems that plagued mental state-based approaches, such as locating its mark in an empirically unobservable realm, its explanatory redundancy and its requiring infinite encoding. Furthermore, the concept of implicit understanding is given more room to flourish.

Ability-based approaches do entail considering what lies *beyond* observable acts but not what lies *behind* them, as mental state-based accounts presumed were necessary. The ability-based approach does not, however, preclude us from using “internal” concepts such as beliefs, provided they are instrumental postulates that function as an explanatory interpretation derived from the way the subject acts (as is the case in interpretationist approaches to the mind). Finally, we briefly considered some candidate kinds (or brands) of abilities offered by the literature as the appropriate one(s), to indicate that we will consider none of them as the necessary or sufficient condition for understanding, but instead as what *composes* understanding.

References

- Adams F, Aizawa K (2010) Defending the bounds of cognition. In: Menary R (ed) *The extended mind*. MIT Press, Cambridge, pp 67–80
- Avigad J (2008) Understanding proof. In: Mancosu P (ed) *The philosophy of mathematical practice*. Oxford University Press, New York, pp 317–353
- Barnby P, Harries T, Higgins S, Suggate J (2007) How can we assess mathematical understanding? In: Woo JH, Lew HC, Park KS, Seo DY (eds) *Proceedings of the 31st conference of the international group for the psychology of mathematics education*, vol 2. PME, Seoul, pp 41–48
- Baumberger C, Beisbart C, Brun G (2016) What is understanding? An overview of recent debates in epistemology and philosophy of science. In: Grimm SR, Baumberger C, Ammon S (eds) *Explaining understanding: new perspectives from epistemology and philosophy of science*. Routledge, New York, pp 1–34
- Brogaard B (2005) I know. Therefore, I understand (unpublished typescript)
- Carroll L (1895) What the Tortoise said to Achilles. *Mind* 4(14):278–280

- Davis PJ, Hersh R (1998) The ideal mathematician. In: Tymoczek T (ed) *New directions in the philosophy of mathematics: revised and expanded*. Princeton University Press, Princeton, pp 177–184
- Dawson JW (2006) Why do mathematicians re-prove theorems? *Philos Math* 14(3):269–286
- de Regt HW (2004) Discussion note: making sense of understanding. *Philos Sci* 71(1):98–109
- de Regt HW (2009) The epistemic value of understanding. *Philos Sci* 76(5):585–597
- de Regt HW (2017) *Understanding scientific understanding*. Oxford University Press, New York
- de Regt HW, Dieks D (2005) A contextual approach to scientific understanding. *Synthese* 144(1):137–170
- Delarivière S (2016) Artificial free will: the responsibility strategy and artificial agents. *Apeiron Stud J Philos (Portugal)* 7(1):175–203
- Delarivière S (2020a) Collective understanding: a conceptual defense for when groups should be regarded as epistemic agents with understanding. *AVANT* 11(2)
- Delarivière S (2020b) *Characterising understanding and the understanding subject* (Doctoral thesis). Vrije Universiteit Brussel, Belgium
- Delarivière S, Van Kerkhove B (2017) The “artificial mathematician” objection: exploring the (im)possibility of automating mathematical understanding. In: Sriraman B (ed) *Humanizing mathematics and its philosophy*. Birkhäuser, Basel, pp 173–198
- Delarivière S, Frans J, Van Kerkhove B (2017) Mathematical explanation: a contextual approach. *J Indian Counc Philos Res* 34(2):309–329
- Dennett DC (1990) *The intentional stance*. MIT Press, Cambridge
- Dennett DC (1993) *Consciousness explained*. Penguin Books, London
- Dutant J (2015) The legend of the justified true belief analysis. *Philos Perspect* 29(1):95–145
- Elgin CZ (2007) Understanding and the facts. *Philos Stud* 132(1):33–42
- Fara M (2008) Masked abilities and compatibilism. *Mind* 117(468):843–865
- Ford KM, Glymour C, Hayes PJ (eds) (1995) *Android epistemology*. MIT Press, Cambridge
- Ford KM, Glymour C, Hayes PJ (eds) (2006) *Thinking about android epistemology*. MIT Press, Cambridge
- Frans J (2020) Unificatory understanding and explanatory proofs. *Found Sci*. <https://doi.org/10.1007/s10699-020-09654-4>
- Frans J, Weber E (2014) Mechanistic explanation and explanatory proofs in mathematics. *Philos Math* 22(2):231–248
- Friedman M (1974) Explanation and scientific understanding. *J Philos* 71(1):5–19
- Geist C, Löwe B, Van Kerkhove B (2010) Peer review and knowledge by testimony in mathematics. In: Löwe B, Müller T (eds) *Philosophy of mathematics: sociological aspects and mathematical practice*. College Publications, London, pp 155–178
- Gettier EL (1963) Is justified true belief knowledge? *Analysis* 23(6):121–123
- Gopnik A (1998) Explanation as orgasm. *Mind Mach* 8(1):101–118
- Gordon EC (2012) Is there propositional understanding? *Logos Episteme* 3(2):181–192
- Greco J (2000) Putting skeptics in their place: the nature of skeptical arguments and their role in philosophical inquiry. Cambridge University Press, Cambridge
- Greco J (2007) The nature of ability and the purpose of knowledge. *Philos Issues* 17(1):57–69
- Griffin J (1986) Well-being: its meaning, measurement and moral importance. Clarendon Press, Oxford
- Grimm SR (2011) Understanding. In: Pritchard D, Bernecker S (eds) *The Routledge companion to epistemology*. Routledge, New York, pp 84–94
- Grimm SR (2012) The value of understanding. *Philos. Compass* 7(2):103–117
- Grimm SR (2014) Understanding as knowledge of causes. In: Fairweather A (ed) *Virtue epistemology naturalized*. Springer, Dordrecht, pp 329–345
- Grimm SR (2016) Understanding and transparency. In: Grimm SR, Baumberger C, Ammon S (eds) *Explaining understanding: new perspectives from epistemology and philosophy of science*. Routledge, London, pp 251–271
- Gordon (n.d.) *Understanding in epistemology*. In: *Internet Encyclopedia of Philosophy*. Retrieved from <https://www.iep.utm.edu/understa/>.
- Hempel C (1965) *Aspects of scientific explanation, and other essays in the philosophy of science*. Free Press, New York
- Hersh R (1997) *What is mathematics, really?*. Oxford University Press, Oxford

- Hiebert J, Carpenter TP (1992) Learning and teaching with understanding. In: Grouws DA (ed) Handbook of research on mathematics teaching and learning: a project of the national council of teachers of mathematics. Macmillan, New York, pp 65–97
- Hills A (2009) Moral testimony and moral epistemology. *Ethics* 120(1):94–127
- Hills A (2015) Understanding why. *Noûs* 50:661–688
- Kelp C (2015) Understanding phenomena. *Synthese* 192(12):3799–3816
- Khalifa K (2013) Understanding, grasping and luck. *Episteme* 10(1):1–17
- Kitcher P (1989) Explanatory unification and the causal structure of the world. In: Kitcher P, Salmon WC (eds) *Scientific explanation*. University of Minnesota Press, Minneapolis, pp 410–505
- Kuorikoski J, Ylikoski P (2015) External representations and scientific understanding. *Synthese* 192(12):3817–3837
- Kvanvig JL (2003) *The value of knowledge and the pursuit of understanding*. Cambridge University Press, Cambridge
- Lipton P (2009) Understanding without explanation. In: de Regt HW, Leonelli S, Eigner K (eds) *Scientific understanding: philosophical perspectives*. University of Pittsburgh Press, Pittsburgh, pp 43–63
- MacKenzie D (1999) Slaying the Kraken: the sociohistory of a mathematical proof. *Soc Stud Sci* 29(1):7–60
- MacKenzie D (2004) *Mechanizing proof: computing, risk, and trust*. MIT Press, Cambridge
- Moore GE (1912) Free will. In: Moore GE (ed) *Ethics*. Oxford University Press, Oxford, pp 84–95
- Newman M (2012) An inferential model of scientific understanding. *Int Stud Philos Sci* 26(1):1–26
- Palermos O, Pritchard D (2016) The distribution of epistemic agency. In: Reider P (ed) *Social epistemology and epistemic agency: de-centralizing epistemic agency*. Rowman & Littlefield, London, pp 109–126
- Pritchard D (2009) Knowledge, understanding and epistemic value. *R Inst Philos Suppl* 64:19–43
- Pritchard D (2010) Cognitive ability and the extended cognition thesis. *Synthese* 175(1):133–151
- Pritchard D (2014) Knowledge and understanding. In: Fairweather A (ed) *Virtue epistemology naturalized*. Springer, Cham, pp 315–327
- Rav Y (1999) Why do we prove theorems? *Philos Math* 7(3):5–41
- Reber AS (1989) Implicit learning and tacit knowledge. *J Exp Psychol Gen* 118(3):219
- Riggs WD (2003) Understanding “virtue” and the virtue of understanding. In: DePaul MR, Zagzebski L (eds) *Intellectual virtue: perspectives from ethics and epistemology*. Clarendon Press, Oxford, pp 203–226
- Ryle G (2000) *The concept of mind*. Penguin books, London (Original work published 1949)
- Schwitzgebel E (2019) Belief. In: Zalta EN (ed) *The Stanford encyclopedia of philosophy*. Retrieved May 7, 2019, from <https://plato.stanford.edu/archives/fall2019/entries/belief/>.
- Sierpiska A (1990) Some remarks on understanding in mathematics. *Learn Math* 10(3):24–41
- Sierpiska A (1994) *Understanding in mathematics*. The Falmer Press, London
- Skemp RR (1976) Relational understanding and instrumental understanding. *Math Teach* 77(1):20–26
- Steiner M (1978) Mathematical explanation. *Philos Stud* 34(2):135–151
- Swart E (1980) The philosophical implications of the four-color problem. *Am Math Mon* 87(9):697–707
- Thurston WP (1998) On proof and progress in mathematics. In: Tymoczko T (ed) *New directions in the philosophy of mathematics: revised and expanded*. Princeton University Press, Princeton, pp 337–355
- Tollefsen DP (2006) From extended mind to collective mind. *Cognit Syst Res* 7(2–3):140–150
- Toon A (2015) Where is the understanding? *Synthese* 192(12):3859–3875
- Trout JD (2002) Scientific explanation and the sense of understanding. *Philos Sci* 69(2):212–233
- Trout JD (2005) Paying the price for a theory of explanation: de Regt’s discussion of trout. *Philos Sci* 72:198–208
- Van Camp W (2014) Explaining understanding (or understanding explanation). *Eur J Philos Sci* 4(1):95–114
- Vervloesem K (2007) *Computerbewijzen in de Wiskundige Praktijk* (MA thesis). Katholieke Universiteit Leuven
- Wilkenfeld DA (2013a) *Understanding as representation manipulability*. *Synthese* 190(6):997–1016
- Wilkenfeld DA (2013b) *Explaining and understanding*. (Doctoral thesis). The Ohio State University, Columbus
- Wittgenstein L (1953) *Philosophical investigations* (trans. Anscombe GEM). Oxford: Basil Blackwell

- Woodward J (2003) *Making things happen: a theory of causal explanation*. Oxford University Press, New York
- Ylikoski P (2009) The illusion of depth of understanding in science. In: de Regt HW, Leonelli S, Eigner K (eds) *Scientific understanding: philosophical perspectives*. University of Pittsburgh Press, Pittsburgh, pp 100–119
- Ylikoski P (2014) Agent-based simulation and sociological understanding. *Perspect Sci* 22(3):318–335
- Zagzebski L (2001) Recovering understanding. In: Steup M (ed) *Knowledge, truth, and duty*. Oxford University Press, New York, pp 235–251

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.