



Understanding as Usability and Context-Sensitivity to Interests

Andreas Søndergaard¹

Received: 11 January 2023 / Revised: 6 November 2023 / Accepted: 7 November 2023
© The Author(s) 2023

Abstract

Is understanding subject to a factivity constraint? That is, must the agent's representation of some subject matter be accurate in order for her to understand that subject matter? 'No', I argue in this paper. As an alternative, I formulate a novel manipulationist account of understanding. Rather than correctly representing, understanding, on this account, is a matter of being able to manipulate a representation of the world to satisfy contextually salient interests. This account of understanding is preferable to factivism, I argue, mainly for simplicity reasons. While it explains the intuitive data about understanding as successfully as factivist accounts, it is simpler by virtue of reducing the value truth bestows on understanding to that of usability.

Keywords Understanding · Veritism · Assessment-sensitivity · Resource-bounded agents · Pragmatic encroachment

1 Introduction

Recently, the concept of understanding has experienced a resurgence in analytic epistemology. In this paper, I join in on this turn and inquire into the nature of understanding. More specifically, my focus is on whether understanding is factive. Must the vehicle of understanding be accurate in order for an agent to understand some subject matter? This issue is "one of the most hotly debated questions in the epistemology of understanding" (Hannon, 2021: 277). Factivists answer in the affirmative, while the most prominent alternative is manipulationism. Understanding, on such an account, consists in an ability to manipulate representations. In this

I am grateful to Asbjørn Steglich-Petersen for discussion of the ideas in this paper as well as the comments from all anonymous reviewers at *Philosophia* and another journal.

✉ Andreas Søndergaard
andreas.sondergaard@philosophy.su.se

¹ Stockholm University, Stockholm, Sweden

paper, I side with the manipulationist camp and formulate a novel version of the position. My main argument for this view is theoretical economy. Manipulationism is simpler in that accuracy is reducible to usability with respect to understanding.

To reach this conclusion, the paper will take the following structure. In Section 2, I lay out the contents of factivism, its motivations, main problems, and what I take to be the most plausible version of the view. In Sections 3–6, I flesh out my manipulationist alternative. In Section 3, I will lay out what a usability condition on understanding amounts to. In Section 4, I will elucidate the contextual nature of understanding and claim understanding to be relative to our interests. Then in Section 5 I will save manipulationism from tempting misunderstandings by responding to two objections. Finally, in Section 6, I will give an account of how we determine degrees of usability. In the rest of the paper, I argue that this version of manipulationism is preferable to factivism. In Section 7, I present my positive argument. I argue that accuracy is reducible to usability and so my manipulationist position can explain the value accuracy is thought to bestow on understanding. In Section 8, I respond to two objections to my reduction of accuracy to usability.

2 Factivism about Understanding

There is widespread agreement that knowledge is a factive epistemic state in the sense that truth is a necessary condition for correctly attributing knowledge: *S* knows *p* only if *p* is true.¹ A similar condition has frequently been held for understanding. Indeed, such a “condition has been implicitly or explicitly accepted by most writers on explanation and understanding” (de Regt & Gijsbers, 2017: 50). Factivists claim that understanding requires the vehicle of understanding to be accurate. Phlogiston theory, for example, does not help us understand combustion on this account since it is false. Rather than construing accuracy as a necessary condition, however, it has become more common to treat it as an evaluative criterion.² That is, a person understands a subject matter *better* the more accurate the vehicle of understanding is. At a first pass, this can be summed up in the following principle:

COVARIANCE

The degree to which *S* understands *O* by means of *V* varies in the same direction as the degree of accuracy of *V* (adapted from Rancourt, 2017: 386).³

Knowing what factivism claims, let us turn to ask what might motivate such a position. Many philosophers simply claim to intuit that there is some factivity condition on understanding (even one of the most prominent opponents of a factivity condition concedes this point: de Regt, 2015: 3782). This assertion is

¹ For a dissenting voice, see Hazlett (2010, 2012).

² Baumberger (2019) claims that an account of understanding *in general* must take this structure, since understanding is a graded concept.

³ I am not suggesting that anyone explicitly endorses COVARIANCE. It functions as a proxy position to set up the dialectic. Supposedly, most proponents of factivism would agree that it should be weakened.

sometimes given in very general terms and left at that. Some even claim this to be a near universal intuition among professional philosophers and “almost undisputed” (Bachmann, 2020: 78; see also Baumberger et al., 2017: 7). More persuasive, however, are intuitions about specific cases which seem to suggest a tight connection between understanding and accuracy. A factivity constraint, it is argued, can explain cases of misunderstanding, retractions of, and disagreement about understanding. The following is an example of the first of these:

REFRIGERATOR

Suppose that you open the refrigerator and notice that the light has gone off, so you put your hand inside and feel around, finding it warm to the touch. While looking around for an explanation eventually you notice something unusual: the cord has been unplugged. You now take yourself to understand why the refrigerator stopped working, but it should be obvious that you might be mistaken. It might have broken down as the result of a short circuit, and the chord might have been unplugged for fear of fire. In that case, despite your sense that you understand why the refrigerator stopped working, in fact you don’t understand. You’ve mischaracterized how things stand in the world (Grimm, 2006: 517–518).

In this case, supposedly, we attribute *mis*understanding rather than genuine understanding since the vehicle of understanding is inaccurate. At least, the subject would understand the refrigerator *better* had she known the *true* cause of its malfunctioning.

Despite these arguments in its favor, factivism has not gone entirely uncontested. To mention a few objections, it has been pointed out that superseded theories provide understanding though they have been shown false (de Regt, 2015), scientific development from one mistaken ontology to another may nonetheless leads to an increase in understanding (Elgin, 2007), contemporary science provides understanding though probably false on the grounds of the pessimistic meta-induction (see Laudan, 1981), contradictory theories can provide understanding (Zagzebski, 2001: 244; de Regt, 2015: 3791–3792), science is ripe with what Elgin (2004, 2017) terms ‘felicitous falsehoods’ such as idealizations (see also Reiss, 2012), and teaching myths are pervasive as a pedagogical tool to help students understand (Stewart & Cohen, 1997: 36–38).⁴

⁴ I am grateful to an anonymous reviewer who pointed out that a weaker, more plausible form of factivism does not hold the *vehicle* of understanding to be accurate. What matters instead is that the subject is able to *extract* enough truths about the object from this vehicle in order for the subject to understand the object. Rice (2021), for example, suggests that we can extract true modal information from inaccurate models or theories. This, however, does not save a narrow version of factivism that *only* focuses on truth as we can extract a wide range of truths that do not advance understanding. These are useless, trivial truths or truths that only apply to situations too far from our actual situation to be interesting. Knowing what would happen to his plants if sent in orbit around Mars after having soaked in coffee for three days does not help the farmer who just wants to make a living understand his plants any better. It is irrelevant information, he feels, and only complicates his representation of his plants. Thus, it seems to me, this weaker version of factivism does not save COVARIANCE either. Rice, supposedly, does not himself defend such a narrow version of factivism, as he allows that “in different contexts, some pieces of information will be more important (or salient) to one’s understanding” (2021: 4106). The point is that mere accuracy does not *suffice* to tell us everything about understanding. Other things influence understanding as well.

Rather than attempting to dismiss these counterexamples, another strategy that I will henceforth discuss is to weaken the factivist position to accommodate the cases. COVARIANCE implicitly assumes that accuracy is the *only* evaluative criterion of understanding.

MONISTIC FACTIVISM

Accuracy is the sole evaluative criterion for understanding.

Since accuracy alone evaluates understanding, the overall degree of any instance of understanding can be read directly off of the degree of accuracy. Weaker versions of factivism, however, allow for there to be *other* evaluative criteria as well.

PLURALISTIC FACTIVISM

Accuracy is one among several evaluative criteria for understanding.

If that is the case, accuracy is merely a *pro tanto* good that is defeasible by gains in other evaluative criteria. We need to weigh the different criteria to make an overall assessment and make any claims about the degree of understanding *all things considered*. On this conception, then, accuracy is a *ceteris paribus* condition for understanding.

PLURALISTIC FACTIVISM is actually consistent with radically inaccurate vehicles providing *better* understanding than their more accurate counterparts. Increased accuracy may be accompanied by a decrease in other evaluative criteria. If the decrement in the other evaluative criteria is severe enough or we are in contexts where those are weighed heavier than accuracy, the overall degree of understanding may decrease though the degree of accuracy increases. This gives the Pluralistic Factivist the resources to explain why we feel idealized models or fictional theories may provide better understanding than highly accurate ones. They may be accompanied by an extensive increase in ease of usability or other evaluative criteria. All that PLURALISTIC FACTIVISM claims is that *all else being equal*, increased accuracy leads to better understanding. This is consonant with the denial of COVARIANCE. So, it seems, our supposed counterexamples do not suffice to rule out all formulations of the factivist position.

PLURALISTIC FACTIVISM is explicitly endorsed by Wilkenfeld (2017) who defends the following account of understanding:

MULTIPLE UNDERSTANDING DIMENSIONS [MUD]

The quality of a state of understanding is evaluable along multiple orthogonal dimensions, including (but perhaps not limited to) *both* representational accuracy *and* intelligibility [emphasis original] (2017: 1276).⁵

Wilkenfeld's intelligibility criterion will be fleshed out in the following section. I prefer 'usability' to 'intelligibility', but this divergence is merely terminological. Wilkenfeld (2017: 1275) "count[s] as intelligible any understanding such that the

⁵ Other proponents of of PLURALISTIC FACTIVISM include Riggs (2003: 219) and Baumberger (2019: 378–379).

possessor of it can as a result do certain things or make certain (good) inferences.” ‘Usability’ better captures this, I think.

3 The Usability Condition

In the previous section, I rejected MONISTIC FACTIVISM as it wrongly rules out cases of genuine understanding. PLURALISTIC FACTIVISM is preferable to it since weaker. In the rest of the paper, I formulate a non-factive account of understanding and argue that it is preferable even to this weaker form of factivism.

Wilkenfeld’s MUD account of understanding took our practices of assessing understanding to be subject to (at least) two evaluative criteria: accuracy and usability. I suggest we drop the former altogether and retain only the latter criterion. This yields the following account:

MONISTIC ABILITISM

Usability is the sole evaluative criterion for understanding.

What do I mean when I claim understanding to be connected to usability? For ease of exposition, I will largely adopt Wilkenfeld’s own account and add my own points when needed. What is central to this component of understanding, according to Wilkenfeld, is “our ability to interact with the world” and “what understanding empowers one to do in the world” (Wilkenfeld, 2013: 1003).

What kinds of things, more precisely, do we need to be able to do to qualify for understanding? ‘Manipulation’, seems to be Wilkenfeld’s answer. Yet, this simple answer comprises two aspects. Both of these are present in the following quote: “understanding must [...] consist of an ability to manipulate some mental correlate of the understood object such that, in the absence of interfering factors, one would then be able to manipulate the target itself” (Wilkenfeld, 2013: 1003).

The first aspect is *representation* manipulability. Wilkenfeld points to the “ability to manipulate some *mental correlate* of the understood object.” What does this consist in? Two things, according to Wilkenfeld. The ability to correct mistakes in the representation (I will not say much more about this) and, more interesting for our purposes, the ability to apply it to cases and solve problems. To give an example, consider logical understanding: the ability to apply a logical theory to determine whether some formula is valid or being able to construct proofs of theorems is paradigmatic of understanding logic. A deeper understanding arises when I can figure out which validities, theorems, and meta-logical properties arise in different semantics or deduction systems. Another case in point is personal understanding. Understanding a person, Wilkenfeld claims, consists in forming a manipulable representation of that person to make inferences about her past or future actions (2013: 1005). Understanding my wife seems to amount to me having a representation of her psychological makeup such that I can explain why she performed certain actions or predict what she would do in a range of situations.

Having said something about the manipulation part of ‘representation manipulability’, what about the first component? Does any kind of representation do? More or less, I think. It is better to err on the side of inclusivity as various kinds

of vehicles can provide understanding. This is the main shortcoming of de Regt's influential ability account according to which one must possess an intelligible *theory* (e.g., 2009). Construing 'theory' narrowly, at least, this is too restrictive. Consider personal understanding again. It is controversial exactly which cognitive resources are needed for this accomplishment (for discussion, see Grimm, 2016, 2019; Khalifa, 2019). Sometimes, it seems, we understand another person by empathetically stepping into their shoes (this idea has a long history in the *Verstehen* tradition in continental philosophy) (for an overview, see Kögler & Stueber, 2000). We take on their perspective by imagining ourselves in their circumstances and adopting some of their psychological states, run what Goldman (2006: 32) calls a 'process-driven simulation' by engaging our own decision-making mechanisms, and see what we would do. It is not clear that such simulations are reducible to tacit theory. Even scientists often use models, equations, diagrams, images, etc. Rather than discriminating between kinds of representations, the central thing seems to be being able to get the right result. Was I to simulate my wife's reaction in a given situation (perhaps I wondered why she got angry after I reorganized the kitchen) or attempting to deduce it from a theory, but got a different result as output than what she actually did (I thought she would be happy), I would not have understood her.

The second aspect of manipulation central to understanding consists in utilizing the ability to manipulate representations to "be able to manipulate *the target itself*" (Wilkenfeld, 2013: 1003). Taking the case of understanding my wife again, it seems that my ability to manipulate my representation of her gives rise to an ability to act effectively around her. If I want to make her happy on her birthday, I know what gift to buy her and what plans to make for her. If I want to avoid her getting upset after a long day's work, I know what things to avoid saying and doing as well as what I can do to positively set her in a good mood. There is a tight relationship between understanding and effective agency, it seems.

On top of Wilkenfeld's account, I make two additions to flesh out 'usability' more in the rest of the paper. Firstly, in Section 4, I discuss *which* kinds of manipulations *S* should be able to do to understand and argue that *S* should be able to manipulate the object so as to satisfy the interests of an assessor of some understanding attribution. Moreover, in Section 6, I suggest three components that determine *degrees* of usability: manageability, reliability, and scope of manipulations.

4 Contextual Sensitivity to Interests

Can we say something more about the tasks that understanding is thought to give us the ability to solve? Precisely which tasks should one be able to perform before one can be said to have understanding? "It all depends", is the most plausible answer, I think. More specifically, different tasks are salient in different contexts and it depends on our interests to what degree someone can be said to understand an object in any given situation.

It is near ubiquitous *that* understanding attributions depend on contextual factors to some extent (for a small sample, see Bachmann, 2020; Baumberger, 2019; Dellsén, 2020; de Regt & Dieks, 2005; Kelp, 2015; Rancourt, 2017;

Wilkenfeld, 2015). A frequently employed argument for this conclusion concerns cases of faultless disagreement (see de Sa, 2011). Let us borrow another case from Wilkenfeld (2013: 1007):

VARIANCE

The understanding of someone who got a 4 on her AP calculus test might clearly count as understanding when being evaluated for one job opportunity (e.g. high-school summer intern), clearly not count when being evaluated for another (e.g. professorship in an elite math department), and neither clearly count nor not count when being evaluated for yet a third (e.g. admission as a student to an elite math department).

These are contradictory claims about whether the agent understands mathematics, yet none of them seem infelicitous. The subject, vehicle, and object, however, are held fixed across all three contexts, so understanding is not exhausted by these. Understanding attributions, then, must be relativized to an additional parameter about the context not explicitly stated in everyday language. Thus, truth about understanding is not a monadic property of propositions as it appears on the surface. In fact, the contextualist claims, truth is a two-place relation between an understanding attribution and some feature(s) of the context.

Having noted the contextual nature of understanding attributions, the next question becomes how exactly this context dependence should be fleshed out. Recently, Bachmann (2020) has complained that no extensive treatment of the supposed context-sensitivity of understanding has been developed. In the following, I will outline a preliminary sketch of such a contextual account.

To give a contextual treatment of understanding attributions, we need to specify *what* exactly these are relativized to. On my account, understanding attributions are relativized to our interests. I think this reflects our use of ‘understands’ in everyday language quite well:

COMPUTER PROBLEMS

Your computer has broken down. You consider to have Ted who is a software developer or Ned the computer technician help you. You reason as follows: “Ned has a better understanding. I better call him.” The next day at work, you take part in a project to develop an app for the company at which you work. You make the following suggestion: “I have a friend, Ted. No one understands computers better than him. Let us hire him to work on the project.”

This point has sometimes been made before yet often only in passing (see Wilkenfeld, 2013: 1007–1008, 2015: 6, 2017: 1267; de Regt, 2015: 3791; Zagzebski, 2019: 125).

Relativizing understanding attributions to interests is still quite vague. Let us make four points to get more precise. Firstly, interests should not be confined to merely practical such. We also have intellectual interests. Philosophical understanding is a case in point. Understanding understanding by giving an account of it, for example, is not immediately practically relevant in that it neither feeds me tonight nor helps me fix my chair that has been broken for a while. But

it may become relevant for other *intellectual* tasks if, for example, we attempt to give a characterization of epistemic value that incorporates understanding (as has been the focus of much of the recent interest in the concept of understanding) or attempt to analyze warranted assertion in terms of it (see Carter & Gordon, 2011).

Secondly, de Regt and Gijsbers (2017: 56) restrict their account of understanding to *scientific* goals. They motivate this choice as follows:

[I]n the Soviet Union during the middle decades of the 20th century, applications of Lysenko's theory of heritability led to a far greater social success than applications of Darwin's theory of natural selection, as Darwinian geneticists were imprisoned and even killed. But it would be perverse to say that in this context, Lysenkoism led to more understanding than Darwinism.

As a general account of understanding, this restriction seems unmotivated, since we clearly possess understanding outside scientific contexts as well. I can be said to understand the local metro system though I only want to visit my mother in the other end of town and not utilize my representation of it for any scientific purposes at all. Perhaps de Regt & Gijsbers intend only to capture a special, scientific kind of understanding and analyze ordinary understanding differently. While an account of scientific understanding is valuable, I think a general account is preferable since it is more unified. For that reason, it is better to relativize understanding to all our interests without specifying a specific subclass, if possible. Fortunately, I do not think de Regt & Gijsbers's case undermines this possibility. To appreciate this point, note that a similar case can show that satisfaction of scientific interests is not sufficient for understanding either.

LUCKY SCIENTIST

Cindy the capricious scientist is a firm believer in fairies. But each time she attempts to use this theory to make predictions or useful practical applications, she is very careless in her work. She mixes up words, confuses her tools and uses a telescope when she meant to be using a microscope, and makes numerous, grave miscalculations. As it turns out, she is extremely lucky and makes extraordinary scientific discoveries each time she attempts to do any scientific work at all.

Yet, to borrow de Regt & Gijsbers's words, it would be perverse to say that fairyism leads to the best understanding of the world. This suggests an alternative diagnosis of the case if understanding is to be relativized to interests at all. Rather than blaming the kind of interests at play, the fault lies with the relation between the representation and the satisfaction of the interests. It is *how* the theories are used to satisfy the interests that better explains why they do not provide understanding. In both cases, it is not the contents of the theories that lead to success, but matters entirely external to them. In **LUCKY SCIENTIST**, it is Cindy's extreme luck and in the case of Lysenkoism, it is the wider sociopolitical climate. If Lysenkoism *itself* led to practical applications to satisfy our non-scientific interests, we would be much more prone to grant

it understanding-conferring properties. That would be the case if, for example, Lysenkoist ideas could be used in agriculture to lead to a radical increase in crop yield as Lysenko himself claimed. These results failed to appear, however, which is a more plausible explanation of why we attribute better understanding to Darwinists. This, however, does not reduce understanding to accuracy as inaccurate representations can also have successful, practical applications in and of themselves.

Thirdly, immediate interests are not all-encompassing. Though Ned the computer technician is not well-versed in software design, it is wrong to think that he has *no* understanding of computers *whatsoever* in a context where we want to develop an app. Rather than relativizing understanding to a single purpose of ours, like developing an app, it may be better to postulate a one-to-many relation between contexts and the set of all our interests which specifies a weight distribution on all our interests. Thus, though being able to fix a broken computer is not immediately relevant, the weight of satisfying this desire is not zero, wherefore Ned can correctly be said to possess *some* understanding of computers. I leave it an open question just how we weigh these interests.

Fourthly and finally, we may want to know *whose* interests we are considering. On Bachmann's (2020: 81) account, it is the *utterer* of the understanding attributions who determines which tasks are relevant. Wilkenfeld (2013: 1003, 2014: 3383, 2019: 2809) makes the same suggestion though he is open to correction. I think such correction is preferable. Recently, MacFarlane (2005a, 2012, 2014) has developed a framework for relativist semantics that postulates a context of assessment distinct from the context of utterance. The context of assessment is a situation where a person *evaluates* whether an utterance is true. MacFarlane applies his relativism to attributions of taste (2007), knowledge (2005b, 2011b), future contingents (2003, 2008), epistemic modals (2011a), and oughts (Kolodny & MacFarlane, 2010). I suggest that assessment sensitivity should be extended to understanding attributions as well. Understanding is relative to the interests of the *assessor* of an attribution of understanding, not the utterer. This puts us in a better position to explain cases of disagreement and retraction. Disagreement, at least sometimes, depends on a difference in interests. A farmer and a botanist may disagree about whether the evolutionary history of wheat provides understanding without either of them being wrong. The farmer takes it as irrelevant for efficient crop growth, while the botanist may take it as understanding conferring since it is important for her intellectual interests. Retractions, too, are well explained by our assessment sensitive semantics.

COFFEE APPRECIATION

A child, like so many other children, is not particularly fond of coffee. He knows it is a black fluid that comes out of the coffee machine and that his parents get cranky if they do not drink it in the morning. This leads him to claim "I understand coffee." Being told that brewing using soft water with low levels of calcium ions and bicarbonate results in acidic tasting coffee does not really increase his understanding of coffee, he feels. It is completely irrelevant information. 15 years later, when the child has grown

to love coffee, this piece of information becomes very important. Being reminded of his claim to understand coffee on the basis of knowing its origin and influence on mood, he says “I take that back. I did not understand coffee at all.”

Contextualism, however, rules out retraction as *both* the previous attribution and subsequent denial are true. On contextualism, we should not retract understanding attributions. But the coffee lover’s retraction clearly seems felicitous. So much worse for contextualism.

In this section, I have claimed understanding to be relative to all contextually salient interests of the assessor of an understanding attribution. This is only a provisional account as much more needs to be said about the contextual nature of understanding.

5 Two Objections

At this point, it is helpful to consider two objections to a tight relationship between understanding and abilities. Our responses to these will help us get rid of some potential misunderstandings and get a firmer grasp on what I take the relation between understanding and abilities to be. The first objection goes against treating abilities as sufficient for understanding. It may seem that usability is too far detached from cognition to be the right place to locate understanding. Ylikoski (2009: 102), for example, seems to separate understanding from specific cognitive states. According to him,

[w]hen we evaluate someone’s understanding, we are not making guesses about his or her internal representations, but about the person’s ability to perform according to set standards. The concept of understanding allows that the ability can be grounded in various alternative ways, as long as the performance is correct.

To appreciate the shortcomings of this approach to the relation between understanding and abilities, consider cases where a student scores perfectly on a test but only does so by pure luck or after having copied all of her neighbor’s answers. Intuitively, in such cases, the student does not possess understanding, though the abilities to do all the right things are present. For this reason, I think it is better to agree with van Camp (2014: 99) that understanding should not be *identified* with abilities, but rather that abilities demonstrate an underlying state. What is this state? I think it is *some* sort of representation of the object to be understood as we have pointed out. I claimed that understanding attributions typically include (or implicitly presuppose) some vehicle that provides the understanding. These considerations show that such a vehicle is actually necessary. Wilkenfeld’s account gives the same result as he takes understanding to consist in *representation* manipulability.

Secondly, it may be thought that usability is not even a necessary condition for understanding. Kelp uses the following case to drive this point home.

SUPERMIND

[C]onsider a cognitive agent, *O*, who is omniscient. *O* knows everything there is to know. Intuitively, *O* will also be omni-understanding in the sense that *O* understands every phenomenon there is to understand. It is conceivable that *O*, while omniscient, is entirely passive, an unmoving knower, as it were. In particular, he does not draw inferences, make predictions, or manipulate his representations in any other way (2015: 3804).

Kelp takes this case to show that abilities to manipulate representations and solve problems are not as central to understanding as we may have thought.

What to think about Kelp's argument? I want to make two points. Firstly, if we discount Kelp's intuition, there is no problem for MONISTIC ABILITISM. Secondly, however, even if one *does* share Kelp's intuition about SUPERMIND, it is not detrimental to my position. While Kelp may be correct that abilities do not figure in *ideal* understanding, it does not follow that they neither do so in *non-ideal* situations. I think abilities to manipulate representations are much more valuable and indicative of a higher degree of understanding than mere knowledge-collection for non-ideal agents. There is only so much information a cognitively limited agent can possess, but being able to manipulate information to acquire new information enables the agent to have a wider range of information ready at hand and act efficiently in a broader class of situations. This is analogous to the famous saying: "give a man a fish and you feed him for a day; teach a man to fish and you feed him for a lifetime." Ideally, the man may prefer to satisfy his nutritional needs by being given a fish every day. Then he would not have to do any work. But in the non-ideal situation where that is not an option, the ability to catch fish himself is preferable to being given a single fish though the latter closer approximates his ideal preferences.

These considerations bring us to the heart of two broad approaches to normative theories in general and epistemology in particular: should our norms govern ideal or resource bounded agents? (for a brief overview, see Caton, n.d.). Generally, I think it is more interesting to pursue theories of normativity for non-ideal agency. After all, none of us are ideal agents and we probably never will encounter any such. It may even be doubted whether such cognitively ideal agents are possible at all (e.g., Grim, 2007). I take it that understanding attributions should be relativized to the agent. Some vehicle V_1 may provide *S* a better understanding than another V_2 due to pragmatic factors entirely internal to *S*. *Given* the fact that *S* is cognitively bounded, V_1 may provide *S* a better understanding than V_2 though V_2 is more accurate and so would figure in ideal understanding – perhaps because V_1 is less complex, more intuitive, fits better with her view of the world, or similar reasons (Zagzebski (2001: 244); Hannon (2021) make similar suggestions). It may be suspected that it is the difference between these two approaches to epistemology that lies at the bottom of the disagreement between factivist and non-factivist conceptions of understanding.

6 Evaluating Degrees of Usability

Now that we have gotten clearer on what abilityism about understanding does and does not involve, let us turn to ask how we evaluate usability so as to be able to differentiate better and worse states of understanding. I think our assessment of a person's abilities are subject to at least the following three concerns (these are similar, though not identical, to the dimensions supplied by de Regt and Gijssbers (2017: 55–58)):

- 1) Manageability
- 2) Reliability
- 3) Scope

What do these factors involve? Manageability, as I take it, is to be thought of along the lines of ease-of-use. This factor is especially closely tied to our non-ideal approach to understanding as it does not seem to play any role for cognitively ideal agents. Supposedly, such agents possess unsurpassed reasoning abilities and so it would not matter much for them whether the information is simple and easy to compute or extremely complex and requires numerous inferential moves. For bounded agents, on the other hand, the story is another. If we are confined to extremely unmanageable representations, no matter how accurate, we will not get to solve very many problems. In fact, the representations may even be beyond our abilities to manipulate altogether such that we *never* get anything done. If the representation is particularly difficult to manipulate, moreover, the probability for mistakes to creep in rises. This naturally leads us to the second dimension.

Reliability with respect to usability is a relation between the agent and a type of task and is roughly a high probability of success. If we consider a task like deciding whether a given proposition is a tautology, an agent is reliable with respect to that task if he successfully completes that task in a range of situations. However, as de Regt & Gijssbers make us aware, probability of success should not be given a frequentist reading.

A theory that would fail in 99 percent of all likely applications could, by sheer coincidence, be applied to just the 1 percent for which it does deliver the right predictions. This makes the scientist lucky, perhaps, but doesn't make the theory reliably successful – in the same way that looking at a clock that has stopped at 2 o'clock doesn't become a reliable method of knowing what time it is just because you happened to look at it at 2 o'clock (2017: 58).

What we need, rather, is success in a range of situations that include counterfactual such. I think this explains some of the motivation to postulate that the agent should possess an internal vehicle to understand. If our student who copied his neighbor's answers, for example, was to be placed in a situation where he could not rely on his neighbor's competences, he would not be very successful

and so should not be deemed a reliable problem-solver. If he himself, however, was able to work out the solutions by manipulating an internal representation, factoring out the reliance on his neighbor would not undermine his reliability.

Finally, scope is a measure of how many types of tasks the agent can complete. While reliability concerned only a single type of task, scope states that an agent who can solve a broader range of tasks in which *O* plays a role understands *O* better. If a person can use his logical theory to find tautologies *as well as* identifying contradictions and valid and invalid arguments, he understands logic better than the person who can merely establish tautologies.

7 Justification of the Usability Account

My argument for favoring MONISTIC ABILITY over PLURALISTIC FACTIVISM takes two steps. Firstly, I will argue that the former can explain all the intuitive data as well as the latter. Thus, PLURALISTIC FACTIVISM does not perform better than MONISTIC ABILITY. Secondly, I will argue that MONISTIC ABILITY is simpler and hence preferable to PLURALISTIC FACTIVISM.

7.1 No Worse Explanatory Scope

The three considerations in favor of a factivity condition on understanding were cases of misunderstanding, retraction, and disagreement. These suggest that understanding is not exhausted by the psychological make-up of the agent, but is subject to an *objective* constraint, which factivists claim to be accuracy. My version of manipulationism, however, adds objective constraints on understanding as well, namely, reliability and scope. This gives us the resources to explain the cases. Recall REFRIGERATOR in which the agent misidentified the cause of the broken refrigerator. MONISTIC ABILITY correctly deems the agent to lack understanding in this case, for *were* he to attempt to get the refrigerator to work again, his goals would not be accomplished. Merely plugging the cord back in would not be enough. The agent is an unreliable problem-solver. Similar considerations apply to cases of retraction and disagreement. Our assessment sensitive semantics suggest an additional diagnosis not open to the factivist, namely, divergence in interests, as already pointed out. COFFEE APPRECIATION is a case where a broadening of interests means that a similar broadening of scope in abilities to solve tasks (namely to now also be able to brew well-tasting coffee) is necessary to understand.

7.2 Simplicity

MONISTIC ABILITY is at least as successful as PLURALISTIC FACTIVISM in accommodating the intuitive data. In the following, I will argue that MONISTIC ABILITY is actually preferable to PLURALISTIC FACTIVISM since simpler.

To state my view on the role of accuracy in understanding, I think that accuracy is valuable in the sense that an increase in accuracy often leads to an increase in

understanding. In most cases, a more accurate representation of the object gives rise to a *better* understanding of that object. This sounds an awful lot like PLURALISTIC FACTIVISM, does it not? To see where I depart from that position, let us invoke a common distinction about value claims, namely, that between instrumental and final value. Crudely put, something is finally valuable if it is valuable for its own sake, while something is instrumentally valuable by virtue of furthering some other value. In the context of understanding, I take values to be factors that determine *degrees* of understanding: they make understanding *better* or *worse*. Final value with respect to understanding, then, is anything that furthers understanding irrespective of its relation to other dimensions of understanding. This is exactly how Wilkenfeld thinks of his dimensions of understanding as “context [...] fixes how much we value [usability] [...] relative to representational accuracy” (2017: 1276). The idea being that each dimension should be treated as a dimension of understanding in its own right and when there are conflicts, we need to balance the two somehow.

While PLURALISTIC FACTIVISM takes accuracy to be a final value of understanding, I think it is merely instrumentally valuable. More precisely, accuracy is instrumentally valuable with respect to usability which, on my view, is the only final value of understanding. That is, the value of accuracy is *reducible* to that of the value of usability. Each time an increase in accuracy leads to better understanding, this can be explained by an increase in usability, but the converse does not hold. I will defend something along the lines of the following reductionist principle:

(R) If V_1 is more accurate than V_2 , then S 's abilities are better if she adopts V_1 rather than V_2 .

At first hand, this principle seems too strong. We have already seen counterexamples to it. If the abilities a vehicle confers are relativized to the agent doing the problem-solving, accuracy may be a hindrance. If, for example, V_2 uses concepts and presupposes background knowledge, S is not familiar with, then S does not become better at solving problems with the help of V_2 . It may instead paralyze her since she does not know how to manage it. Moreover, V_2 may be so complex that it becomes inefficient to solve the relevant task. These considerations suggest that a higher degree of accuracy may lead to a lower degree of manageability which, all else being equal, implies that S becomes a worse problem-solver.

The sense in which I *do* take (R) to be true, however, is with respect to the other two evaluative dimensions of usability. *Prima facie*, it seems clear to me that more accurate representations of the world make us more reliable in our dealings with the world. We are all too aware of situations in which false assumptions lead to the wrong results though our reasoning is impeccable. To give but a simple example:

BROKEN CAR

Your car will not start. You suspect it to be out of gas, so you fuel it. Nothing happens. As it turns out, the battery is flat.

This is an instance where a false belief led to a frustration of the agent's interests. The agent was less reliable in solving problems since her representation departed from the facts. If instead she had truly believed that the battery

was flat, she could have taken another line of action, namely, to charge it. Accuracy not only influences reliability but scope as well. While some inaccuracy may not be immediately relevant for a contextually prominent task, it can still affect our reliability in less salient matters. If our agent also believed the function of the handbrake was to start the wipers, that would not undermine her reliability with respect to getting the car to work in the specific situation. Yet, it would still decrease her overall understanding of the car by making her an unreliable driver.

So, we should slightly weaken (R), it seems. While accuracy may not increase manageability, it seems more plausible that it is accompanied by a higher degree of reliability as well as scope.

(R') If V_1 is more accurate than V_2 , then S is more reliable in a broader scope of tasks if she adopts V_1 rather than V_2 .

For this reason, increased accuracy is a *pro tanto* gain in usability by virtue of increasing reliability and scope, though an agent's abilities may be worse off *all things considered* if it is accompanied by an extensive decrease in manageability. Let me note that accuracy is not taken to be a *necessary* condition for increase in reliability and scope. False theories can be highly reliable as well in a broad range of situations as idealized or fictional theories evidence. Thus, usability does not also reduce to accuracy, wherefore understanding does not correlate with accuracy. The point about (R') is only to account for the intuitions about cases where it *does* seem that accuracy is a progress in understanding. In such situations, we do not need to postulate an independent evaluative criterion of understanding, however, as Wilkenfeld has it. Rather, the value of accuracy can be accounted for by the abilist as merely instrumental.

Finally, we must ask, if we all agree that accuracy is valuable with respect to understanding why does it matter whether it should be thought of as final or mere instrumental value? Because an account that postulates fewer final values is *simpler* and that is a reason to prefer that account. To elaborate a bit further, Ahlstrom-Vij and Grimm (2013: 342) have laid out the following desiderata for an axiology of some domain:

- A) That the position be *axiologically parsimonious*.
- B) That the position accounts for prevalent *pluralistic intuitions*.
- C) That the position accounts for prevalent *monistic intuitions*.

Let us evaluate the various positions with respect to these dimensions. Regarding (A), MONISTIC FACTIVISM and MONISTIC ABILITISM both score well which gives us reason to prefer either of these above PLURALISTIC FACTIVISM. The former positions postulate only one final value of understanding, while the latter invokes (at least) two.

This, however, is only a *pro tanto* reason to prefer monistic over pluralistic positions. For it may be that monisms postulate *too few* final values to account for our intuitions, thus violating (B). I think this is indeed the case for MONISTIC FACTIVISM, as we have seen. Sometimes we value vehicles that are less accurate

over more accurate ones. If accuracy can positively hinder understanding, the project of reducing other values to it seems a non-starter. MONISTIC ABILITISM, I have argued, fares better on this score and this is so for two reasons. Firstly, our evaluation of abilities itself comprises several dimensions. Manageability, reliability, and scope are all species of what it means to be a good problem-solver. Secondly, MONISTIC ABILITISM can accommodate our intuitions about the value of accuracy by treating it as instrumentally valuable for usability. PLURALISTIC FACTIVISM, of course, scores well on this dimension, too, since it postulates multiple final values.

Finally, dimension (C) may at first seem quite similar to (A). With respect to this condition, Ahlstrom-Vij and Grimm (2013: 345–346) claim that an axiology should respect “the intuition that all epistemic goods can be referred back to a single epistemic goal”. This serves to explain why some values are epistemic and others not. Applied to understanding, MONISTIC ABILITISM provides just such a principle. The ability to solve problems functions as

a concluding explanatory postulation of a more general phenomenon, providing a genus for all the [understanding-]values that we end up with when trying to account for the domain of [understanding-]value in terms of a limited number of goals (Ahlstrom-Vij & Grimm, 2013: 345).

PLURALISTIC FACTIVISM fails in this regard since it cannot explain why usability as well as accuracy should both constitute values of understanding. What do they have in common? Why not also add other values such as being formulated in beautiful language or psychological comfort, to borrow Ahlstrom-Vij & Grimm’s own examples? Why should those be ruled out?

If our arguments are correct, we should prefer MONISTIC ABILITISM to MONISTIC FACTIVISM, since the former and not the latter can account for *all* our intuitions about what constitutes good understanding. More interestingly, however, MONISTIC ABILITISM is also superior to PLURALISTIC FACTIVISM. It is simpler and gives a unifying picture of why we value understanding in the first place as well as which aspects about it are conducive to that goal at no explanatory cost. The fact that we are not forced to postulate factivity as an evaluative criterion of understanding to provide a plausible account is reason to refrain from doing so.

8 Further Objections

I have laid out MONISTIC ABILITISM and attempted to convince the reader of its theoretical superiority compared to PLURALISTIC FACTIVISM. In the final part of the paper, I will consider objections to my account. Though MONISTIC ABILITISM can of course be attacked from various directions, I will consider objections only to my reduction of the value of accuracy to usability. Let us consider two types of potential counterexamples. Both of these are situations where the antecedent of (R’) is satisfied, but the consequent is not. Without (R’) accuracy does not always lead to increased abilities to manipulate, wherefore there are situations where increases in accuracy make a *distinct* contribution.

As such, it should be an evaluative criterion *alongside* usability. Usability would not be enough in itself.

8.1 Useless Truths

The first kind of counterexample revolves around situations where our representations get more accurate, yet that makes *no* difference in usability considerations whatsoever. These situations may arise if we add absolutely useless truths to our stock of beliefs about the object we attempt to understand. Suppose you want to understand your daughter and so form the true belief that she is four years old. This belief may be relevant to solve tasks in some contexts. For example, with this piece of information you can predict that taking her to the playground would make her happier than going to the opera or the spa. Was she forty years old, the preferences would likely be reversed. Yet, you can increase the accuracy of your representation of your daughter by pursuing entirely trivial information. Perhaps you add “the age of my daughter to the seventh power is 16,384” or “if my daughter is four years old, then she is four years old or she is an astronaut” to your stock of beliefs about your daughter. But it does not seem that you become any more reliable in solving any tasks in which your representation of your daughter plays a role than you were before. In which contexts would those truths be relevant? Thus, accuracy is not sufficient for better abilities as (R') has it. The two do not change in tandem.

I take it that trivial truths do constitute a genuine counterexample to (R'). The appropriate response, I think, is to weaken (R') to accommodate our intuitions about useless truths. To see how this may be done, I want to point out that I do not intend to defend the reduction of accuracy to usability in all cases. It was introduced to deal specifically with the intuition that accuracy sometimes leads to better *understanding*. The following alteration of (R') captures this intent:

(R'') If V_1 is more accurate than V_2 and S understands O better by adopting V_1 rather than V_2 , then S is more reliable in a broader scope of tasks if she adopts V_1 rather than V_2 .

Now it becomes much more clear why useless truths do not undermine my attempt to reduce accuracy to usability *in the context of understanding*. For our case of useless truths certainly does not satisfy the antecedent of (R''). While your representation of your daughter does get more accurate upon adopting these beliefs, that does not help you to understand her better. It is hard to argue for this claim except by appealing to intuition. Does a life-long pursuit of useless information help one understand the world any better? The more plausible answer is that usefulness and the degree of understanding is close to being perfectly coordinated. The more trivial, the less we feel it helps us understand, but the more useful, the better we understand the target.

8.2 Counterproductive Truths

The second kind of counterexample we will consider is about increases of accuracy that lead to a *decrease* in usability. This is the more serious objection. In fact, this is

the main reason why Wilkenfeld introduced an accuracy condition in addition to the usability condition. He makes the following observation about the transition from geo- to heliocentrism:

The point about Ptolemaic astronomers is that, when Copernicus first proposed his heliocentric model, its predictions were actually less accurate than the Ptolemaic alternative. Given technological limitations, there was not much a 16th century astronomer could *do* with a theory of what revolved around what, so manipulation and control [...] were not really at issue. Thus, the greater predictive powers of the Ptolemaic system seems to imply that it was more intelligible. The problem for [the abilitist] is that this commits him to saying not only that the Ptolemaic astronomers understood why the planets exhibit retrograde motion, but that they actually understood that motion *better* than the Copernicans who correctly put the sun at the center of the solar system (2017: 1278).

Let us grant Wilkenfeld the historical details though de Regt and Gijsbers (2017: 60) qualify the claims somewhat. Thus, an increase in accuracy sometimes comes at the cost of predictive power and control implying a loss in usability as we have defined it.⁶

Does this case suffice to show we value accuracy in understanding irrespective of its contribution to problem-solving abilities? I am not entirely convinced. And the reason is the same as for useless truths. Intuitively, it seems to me, these situations exemplify *decreases* of understanding, and so the antecedent of (R'') is not obviously satisfied. The Copernican revolution led with it a host of unanswered questions and marked the mere beginning of a new research paradigm. Should a contemporary judge Copernicus to have understood the world better though he had not yet figured out the exact consequences of the position, would have to answer "I don't know" to quite a few questions, could not apply the theory, and routinely made wrong predictions? I do not find it outrageously wrongheaded to answer 'no' to this question. The superiority of the understanding the Copernican model provides compared to Ptolemaic astronomy seems to be judged in retrospect. The more natural description of the situation, I think, is that the Copernican theory marked a temporary setback in understanding for the sake of a long-term gain. This increase in understanding seems to have come along with an increased ability to incorporate the heliocentric position in the wider worldview, make predictions, and foster useful applications. de Regt and Gijsbers (2017: 60) take a similar strategy to respond to the alleged counterexample by biting the bullet and maintaining that Ptolemaic astronomers had a better understanding than their Copernican counterparts.

Perhaps not everyone will share these intuitions. Is there not *some* sense in which our understanding of the cosmos increased after Copernicus's discovery? Wilkenfeld, at least, thinks so. However, even if we grant this, it does not need to contradict MONISTIC ABILITISM. We just need to flesh out the position in more detail to figure out exactly what it claims. I can think of two ways of doing so that would enable

⁶ A similar worry stems from misleading truths (Rancourt, 2017: 392–394).

the abilitist to grant a higher degree of understanding to Copernicus compared to the Ptolemaic astronomer. Firstly, we do not need to read (R'') as a synchronic principle though we sometimes interpret conditional statements in such a way. "If you smoke, you die early", for instance, does not express death to be an immediate effect of smoking. Perhaps something similar can be said of (R''). If we treat it as a diachronic principle and all we claim is that accuracy is likely to lead to an increase in usability *in the future* or has the *potential* to do so, we can grant Wilkenfeld that Copernicus understood cosmos better than his contemporaries even before any of these gains were realized. In slightly different terminology, we can follow de Regt and Gijsbers (2017: 56) in letting the development of better science be a mark of a theory's usefulness. By this, they mean that "[a] representational device is more successful if it suggests more avenues of further research, and as those avenues lead to representational devices that themselves are more scientifically successful." On this parameter, Copernican cosmology was arguably more successful than the Ptolemaic theory. If so, we do not have a counterexample to (R'') anymore, since the antecedent is satisfied as well.

A second way to weaken (R'') that allows us to capture Wilkenfeld's intuition is to interpret it in a general rather than case-based way (see Lockard, 2013: 1713–1714). The important thing is not whether a *specific* true theory increases usability, but whether true theories do so *in general*. We have developed a practice of evaluating understanding with respect to its accuracy since this is typically the most efficient way to achieve our purposes. MONISTIC ABILITISM, on this view, is to be understood on the model of rule rather than act utilitarianism. So, even if there are some specific situations where accuracy is indifferent or even detrimental to usability – like the Copernican revolution, supposedly – (R'') can be read in such a way that following the general rule of preferring truth leads to an overall increase in usability. There is quite a bit of flexibility in MONISTIC ABILITISM to let us incorporate intuitions about the value of accuracy in understanding. I let the details be open for future research.

9 Conclusion

In this paper, I have rejected a factivity constraint on understanding. Strong versions of factivism where accuracy is the only evaluative criterion for understanding should be rejected as accuracy can hinder understanding. Highly idealized and fictional theories, for example, sometimes provide better understanding than more accurate counterparts. In its place, I have defended a version of manipulationism according to which understanding is a matter of manipulating a representation, and subsequently the object itself, to satisfy contextually salient interests. This position is preferable to weaker versions of factivism that treats accuracy as one among more evaluative criteria for understanding, I have argued, because it is simpler. When it comes to understanding, truth is reducible to usability. As our representations get more accurate, we become more reliable at solving a wider range of tasks. Thus, we do not *need* to invoke factivist dimensions in our evaluation of understanding. The abilitist conditions suffice.

Funding Open access funding provided by Stockholm University.

Declarations

Conflict of Interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ahlstrom-Vij, K., & Grimm, S. R. (2013). Getting it right. *Philosophical Studies*, 166, 329–347.
- Bachmann, M. (2020). The epistemology of understanding: A contextualist approach. *Kriterion*, 34(1), 75–98.
- Baumberger, C. (2019). Explicating objectual understanding: Taking degrees seriously. *Journal for General Philosophy of Science*, 50, 367–388.
- Baumberger, C., Beisbart, C., & Brun, G. (2017). What Is understanding? An overview of recent debates in epistemology and philosophy of science. In S. R. Grimm, C. Baumberger, & S. Ammon (Eds.), *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science* (pp. 1–34). Routledge.
- Carter, J. A., & Gordon, E. C. (2011). Norms of assertion: The quantity and quality of epistemic support. *Philosophia*, 39, 615–635.
- Caton, J. (n.d.). *Resource Bounded Agents*. Retrieved from *Internet Encyclopedia of Philosophy*: <https://iep.utm.edu/re-bo-ag/>. Accessed 20 Nov 2023.
- de Regt, H. W. (2009). Understanding and scientific explanation. In H. W. de Regt, S. Leonelli, & K. Eigner (Eds.), *Scientific Understanding: Philosophical Perspectives* (pp. 21–42). University of Pittsburgh Press.
- de Regt, H. W. (2015). Scientific understanding: Truth or Dare? *Synthese*, 192(12), 3781–3797.
- de Sa, D. L. (2011). The many relativisms: Index, context, and beyond. In S. D. Hales (Ed.), *A Companion to Relativism* (pp. 102–117). Wiley-Blackwell.
- de Regt, H. W., & Dieks, D. (2005). A contextual approach to scientific understanding. *Synthese*, 144(1), 137–170.
- de Regt, H. W., & Gijsbers, V. (2017). How false theories can yield genuine understanding. In S. R. Grimm, C. Baumberger, & S. Ammon (Eds.), *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science* (pp. 50–75). Routledge.
- Dellsén, F. (2020). Beyond explanation: Understanding as dependency modelling. *The British Journal for the Philosophy of Science*, 71(4), 1261–1286.
- Elgin, C. Z. (2004). True enough. *Philosophical Issues*, 14, 113–131.
- Elgin, C. Z. (2007). Understanding and the facts. *Philosophical Studies*, 132(1), 33–42.
- Elgin, C. Z. (2012). Understanding's tethers. In C. Jäger & W. Löffler (Eds.), *Epistemology: Contexts, Values, Disagreement* (pp. 131–146). Ontos Verlag.
- Elgin, C. Z. (2017). *True enough*. MIT Press.
- Goldman, A. I. (2006). *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford University Press.
- Grim, P. (2007). Impossibility arguments. In M. Martin (Ed.), *The Cambridge Companion to Atheism* (pp. 199–214). Cambridge University Press.
- Grimm, S. R. (2006). Is understanding a species of knowledge. *The British Journal for the Philosophy of Science*, 57(3), 515–535.
- Grimm, S. R. (2016). How understanding people differs from understanding the natural world. *Philosophical Issues*, 26, 209–225.

- Grimm, S. R. (2019). Understanding as an intellectual virtue. In H. Battaly (Ed.), *The Routledge Handbook of Virtue Epistemology* (pp. 340–351). Routledge.
- Hannon, M. (2021). Recent work in the epistemology of understanding. *American Philosophical Quarterly*, 58(3), 269–290.
- Hazlett, A. (2010). The myth of factive verbs. *Philosophy and Phenomenological Research*, 80(3), 497–522.
- Hazlett, A. (2012). Factive presupposition and the truth condition on knowledge. *Acta Analytica*, 27, 461–478.
- Kelp, C. (2015). Understanding phenomena. *Synthese*, 192(12), 3799–3816.
- Khalifa, K. (2019). Is *Verstehen* scientific understanding? *Philosophy of the Social Sciences*, 49(4), 282–306.
- Kögler, H. H., & Stueber, K. R. (2000). Introduction: Empathy, simulation, and interpretation in the philosophy of social science. In H. H. Kögler & K. R. Stueber (Eds.), *Empathy and Agency: The Problem of Understanding in the Human Sciences* (pp. 1–61). Westview.
- Kolodny, N., & MacFarlane, J. (2010). Ifs and Oughts. *The Journal of Philosophy*, 107(3), 115–143.
- Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, 48(1), 19–49.
- Lockard, M. (2013). Epistemic instrumentalism. *Synthese*, 190(9), 1701–1718.
- MacFarlane, J. (2003). Future contingents and relative truth. *The Philosophical Quarterly*, 53, 321–336.
- MacFarlane, J. (2005a). Making sense of relative truth. *Proceedings of the Aristotelian Society*, 105, 305–323.
- MacFarlane, J. (2005b). The assessment sensitivity of knowledge attributions. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology* (pp. 197–234). Clarendon Press.
- MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies*, 132(1), 17–31.
- MacFarlane, J. (2008). Truth in the garden of forking paths. In M. García-Carpintero & M. Kölbel (Eds.), *Relative Truth* (pp. 81–102). Oxford University Press.
- MacFarlane, J. (2011a). Epistemic Modals are assessment-sensitive. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality* (pp. 144–178). Oxford University Press.
- MacFarlane, J. (2011b). Relativism and knowledge attributions. In S. Bernecker & D. Pritchard (Eds.), *The Routledge Companion to Epistemology* (pp. 536–544). Routledge.
- MacFarlane, J. (2012). Relativism. In G. Russell & D. G. Fara (Eds.), *The Routledge Companion to Philosophy of Language* (pp. 132–142). Routledge.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Clarendon Press.
- Rancourt, B. T. (2017). Better understanding through falsehood. *Pacific Philosophical Quarterly*, 98, 382–405.
- Reiss, J. (2012). The explanation paradox. *Journal of Economic Methodology*, 19(1), 43–62.
- Rice, C. (2021). Understanding realism. *Synthese*, 198, 4097–4121.
- Riggs, W. D. (2003). Understanding “Virtue” and the virtue of understanding. In M. DePaul & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (pp. 203–226). Oxford University Press.
- Stewart, I., & Cohen, J. (1997). *Figments of Reality: The Evolution of the Curious Mind*. Cambridge University Press.
- van Camp, W. (2014). Explaining understanding (or Understanding Explanation). *European Journal for Philosophy of Science*, 4, 94–114.
- Wilkenfeld, D. A. (2013). Understanding as representation manipulability. *Synthese*, 190(6), 997–1016.
- Wilkenfeld, D. A. (2014). Functional explaining: A new approach to the philosophy of explanation. *Synthese*, 191(14), 3367–3391.
- Wilkenfeld, D. A. (2017). MUDdy understanding. *Synthese*, 194(4), 1273–1293.
- Wilkenfeld, D. A. (2019). Understanding as compression. *Philosophical Studies*, 176, 2807–2831.
- Wilkenfeld, D. A. (2015). *The contextual quasi-factivity of objectual understanding*. Manuscript.
- Ylikoski, P. (2009). The illusion of depth of understanding in science. In H. W. de Regt, S. Leonelli, & K. Eigner (Eds.), *Scientific Understanding: Philosophical Perspectives* (pp. 100–122). University of Pittsburgh Press.
- Zagzebski, L. T. (2001). Recovering understanding. In M. Steup (Ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (pp. 235–252). Oxford University Press.
- Zagzebski, L. T. (2019). Toward a theory of understanding. In S. R. Grimm (Ed.), *Varieties of understanding: New Perspectives from Philosophy, Psychology, and Theology* (pp. 123–135). Oxford University Press.