

# Is Understanding Reducible?<sup>1</sup>

## Abstract:

Despite playing an important role in epistemology, philosophy of science, and more recently in moral philosophy and aesthetics, the nature of understanding is still much contested. One attractive framework attempts to *reduce* understanding to other familiar epistemic states. This paper explores and develops a methodology for testing such reductionist theories before offering a counterexample to a recently defended variant on which understanding reduces to what an agent knows.

**Keywords:** Understanding – Knowledge – Metaphysics of epistemology – Explanation

## 0. Introduction

The epistemic state of understanding plays an important role in epistemology and philosophy of science, and has become increasingly discussed within moral philosophy and aesthetics. However, there is a striking lack of consensus on the nature of understanding. In the search for a satisfactory account, a prominent theme has been the relationship between understanding and knowing. Although the initial resurgence in theorising about understanding was motivated by apparent differences between understanding and knowing (e.g. Zagzebski 2001; Kvanvig 2003; Elgin 2007), there has recently been a spate of work motivating knowledge-based accounts of understanding—accounts that explain understanding in terms of facts about what an agent knows (particularly Kelp 2014, 2015, 2017; Sliwa 2015, 2017; Riaz 2015; and see Lawler 2016 for discussion).

This paper takes a critical look at an attractive framework for theorising about understanding that has been spawned by these debates—namely ‘reductionist’ accounts that reduce understanding to other more familiar epistemic states. After situating this framework in relation to other prominent views of understanding, I develop a natural way to test such reductionist theories, before using this method to outline a counterexample to a recently influential view on which understanding reduces to a body of knowledge.

## 1. Knowledge, Understanding & Reductionism

‘Understanding’ features in various philosophically interesting locutions, with most recent work focusing on understanding-*why* (e.g. Jones understands why the war was lost) and *objectual* understanding (e.g. Juanita understands Polynesian culture).<sup>2</sup> Here, I focus primarily on *understanding*

---

<sup>1</sup> Forthcoming in *Inquiry*—please cite/email for the published version!

<sup>2</sup> Grimm (2012) summarises early work on understanding. On understanding-*why* see, inter alia, Sliwa (2015; 2017), Hills (2015), Pritchard (2010), Grimm (2006). On objectual understanding see, inter alia, Kvanvig (2003), Kelp (2017), Elgin (2007), Zagzebski (2001).

*why*. (When I use ‘understanding’ henceforth, I refer to understanding *why*). I will not be concerned with ‘linguistic understanding’, such as when one understands the meaning of a word or sentence.

The relationship between understanding and knowing has been a central topic in theorising about understanding. The earliest knowledge-based theories of understanding-*why*, theories that aim to account for understanding in terms of knowledge, defended the following view: *S understands-why p iff S knows-why p*.<sup>3</sup> This position, on which understanding is simply equivalent to some item of knowledge, proved unable to explain the *gradability* of understanding—the fact that, unlike an orthodox conception of knowledge, understanding comes in degrees.<sup>4,5</sup> We can readily conceive of two agents who both share knowledge-*why p* yet are in asymmetrical epistemic positions regarding *p*. A well-known example adapted from Pritchard (2010) exemplifies. Your house burns down due to faulty wiring: *S* is informed of this proposition through extremely reliable testimony, whilst *S\** deduces it from a raft of other beliefs after investigating the scene and applying their experience as an electrician. Plausibly, both know *why* the house burned down—on standard intellectualism, know *that* it burned down due to faulty wiring—yet *S\** possesses greater understanding. Equating knowing and understanding cannot explain such epistemic asymmetries.<sup>6</sup>

Recent work has outlined a more sophisticated knowledge-based approach, utilising the idea that understanding can be *reductively explained* in terms of knowledge.<sup>7</sup> This position offers resources with which to explain the gradability of understanding. To use an analogy due to Sliwa (2017): just as reductive physicalists about the mental can say that a subject experiences a mental state in virtue of a *complex* underlying physical state, the reductionist about understanding can attempt to explain the extent of an agent’s understanding by appealing to *a complex body of relevant and interconnected knowledge*. Such views permit the judgement that two agents can both know *why p* whilst understanding *why p* to different degrees—epistemic asymmetries can be explained by pointing at differences in other things the agents know. So, whilst two agents might know a house burned down due to faulty wiring, asymmetries in understanding can be explained by showing that the agent with greater understanding has more knowledge about (for example) how wires can be faulty, why houses are combustible, etc.

---

<sup>3</sup> E.g. Sliwa (2015), Riaz (2015).

<sup>4</sup> This is pressed effectively by Lawler (2016).

<sup>5</sup> It also prompted debate on the relationship between understanding and epistemic luck. See Pritchard (2010), Hills (2015), Sliwa (2017: 526-527) and Kelp (2017: 261-3).

<sup>6</sup> Grimm (2006; 2014) is also sympathetic to a sort of equivalence view, but places emphasis on knowledge of a special kind: knowledge of the “modal relatedness of the terms of the causal relata” (i.e. knowing that *p* causes *q* in such a way that one grasps the modal relationship between *p* and *q*). I won’t discuss further specifics of Grimm’s view here, but it seems like the same issue arises—two agents might possess such knowledge but be in asymmetrical epistemic positions. In response, it would be open for Grimm to appeal to a reductionist theory as I outline below.

<sup>7</sup> Most prominently, Sliwa (2015; 2017) and Kelp (2014; 2015; 2017).

In short, explaining the gradability of understanding is an essential requirement for any plausible theory and theories that reduce understanding to a body of knowledge have good claim to fulfil this requirement.

Although we have introduced the knowledge-based variant, the reductionist approach offers a more general framework that one might use to theorise about understanding without focusing on knowledge. Indeed, with varying levels of credibility, one could propose any epistemic state as the reductive base of understanding. Let  $\phi$  be a placeholder for some epistemic state, and we get the following very general gloss on reductionism:

**Reductionism:**  $\phi$  is all there is to understanding—the extent to which an agent understands why  $p$  can be explained wholly in virtue of  $\phi$  they possess.<sup>8</sup>

In addition to the knowledge-based view, there are a variety of other possible presifications of **Reductionism** that might have some plausibility. For instance, one might defend a theory on which understanding reduces just to a body of justified true beliefs; just to a body of true beliefs; just to a body of beliefs that have a sufficient degree of verisimilitude; and so forth.<sup>9</sup> These are all *univocal* presifications of **Reductionism**—they reduce understanding to only one type of epistemic state. Of course, on pain of implausibility, this is not to say that someone who defends a justified true belief view would claim that *only* justified true beliefs rather than items of knowledge can constitute one's understanding. It simply means that, to the extent that some piece of knowledge does figure in one's understanding, this is in virtue of the fact that knowing that  $p$  entails justifiably and truly believing that  $p$ . Furthermore, it is also perfectly consistent for there to be a *pluralist* presification of **Reductionism** on which understanding necessarily reduces to more than one type of epistemic state, even though it is harder to immediately see the motivation for such a view.

Crucially, although proponents of the knowledge-based variant of **Reductionism** propose it as an alternative to other dominant perspectives, **Reductionism** is actually compatible with a host of other views that feature prominently in recent literature on understanding. For example, two popular thoughts about understanding are the following:

---

<sup>8</sup> Kelp (2017: 252) points out the need for a 'well-connectedness' proviso in reductionist theories. For instance, an agent might have considerable knowledge but be unable to connect or draw upon it simultaneously, to the detriment of their understanding. I take this proviso as implied.

<sup>9</sup> Beyond the knowledge-based variant, which has been explicitly defended, it is less easy to confidently attribute other **Reductionist** views to particular theorists. This is because, as I discuss below, whether other extant views can rightly be called reductionist depends on whether other theoretical posits which they take to be relevant to understanding—such as grasping and ability—can themselves be reduced to some epistemic state. And extant theories often do not address this issue. However, other views emphasising the idea of understanding being grounded in particular epistemic states include Kvanvig (2003), who emphasises the connection between *objectual* understanding and a body of justified (viz. internally coherent), broadly factive beliefs.

**Grasping:** In order to possess understanding, it is necessary to stand in a particular cognitive relationship—that of *grasping*—towards the object of understanding.<sup>10</sup>

**Ability:** One possesses understanding (at least partly) in virtue of possessing certain *abilities* regarding the object of one’s understanding.<sup>11</sup>

It is (prima facie) possible to affirm or reject any combination of **Reductionism, Grasping** and **Ability**. For instance, one might hold that understanding is constituted by certain abilities. This is consistent with it also being the case that knowledge is all there is to an instance of understanding, just if one also holds that the abilities in question are themselves metaphysically dependent on the agent’s body of knowledge. The same thought applies to grasping. Of course, one might have particular views of grasping or the nature of ability that rule out some combinations of **Reductionism, Grasping**, and **Ability**. For instance, if you thought that abilities are irreducible to knowledge; then you would deny the conjunction of a knowledge based version of **Reductionism** and **Ability**.<sup>12</sup> In principle, however, **Reductionism** is a very general and ecumenical framework with which to investigate the nature of understanding.

## 2. How to test reductionism

Having sketched the nature of such theories, I now turn to develop a natural way to test presifications of **Reductionism**. In so doing I will be recommending a slight refocusing of the debate, but in a way that I suggest will afford us with a clearer grip on the underlying metaphysics of understanding.

Much extant discussion concerns ascriptions of *outright understanding*—i.e. involve cases about which we are asked to judge whether it is true or false that some subject *S* “understands why *p*” simpliciter. However, I propose is that it will be more fruitful to attend to *degrees of understanding* rather than judgements of outright understanding when testing various reductionist theories.<sup>13</sup> The easiest way to execute this approach will be to look at *comparative judgements* about whether one subject possesses a greater or lesser degree of understanding than another. The comparison subject can either be a separate person, or the same person across different times. Shifting the focus to degrees of understanding does not involve any controversial assumptions; it simply relies on the natural thought that if an agent’s understanding why *p* outright is constituted by (e.g.) their knowledge, then their degree of understanding why *p* is also constituted by (e.g.) their knowledge.

---

<sup>10</sup> The idea of grasping features widely in discussions of understanding. Bourget (2007) and Strevens (2013) provide among the most detailed accounts.

<sup>11</sup> Hills (2017) provides an influential statement of an ability-view of understanding why.

<sup>12</sup> This issue leads quickly into difficult territory; the plausibility of reducing abilities to knowledge will depend, in part, on one’s view of knowledge-how. The nature of knowledge-how and how it relates to propositional knowledge is a vexing issue in contemporary epistemology.

<sup>13</sup> The utility of focusing on degrees of understanding is also raised by Kelp (2017); I am indebted to this discussion.

But why should we prefer this approach? I propose that this focus has a number of independent advantages, whilst also providing us with a straightforward methodology with which to test reductionist theories of understanding. One obvious advantage is that the gradability of understanding has been widely noted as a distinctive feature of understanding that sets it apart from knowledge; focusing on degrees of understanding zeros in on this central feature. Moreover, as our earlier discussion of ‘equivalence’ theories illustrated, accounting for epistemic asymmetries between agents constitutes a crucial test for a satisfactory theory of understanding—focusing on comparative judgements ensures that we keep this desideratum firmly in mind. A second advantage is methodological. Attending to degrees of understanding rather than outright understanding allows us to examine very fine-grained differences between epistemic subjects. For, if we suppose that there will often be a considerable gulf between the epistemic position of someone who outright understands why something is the case compared with someone who lacks outright understanding why something is the case, then cases involving this comparison will involve numerous epistemic changes. This raises the *prima facie* worry that it will not always be obvious exactly what our intuitions are tracking when we consider such cases. Looking at degrees of understanding instead allows us to attend to smaller, incremental epistemic changes—thus allowing us to more confidently trace back our intuitive judgements about understanding to very specific changes in a subject’s epistemic position. Thirdly, focusing on degrees of understanding controls for unhelpful theoretical noise due to context-sensitivity influencing judgements about outright understanding. It seems plausible that the standards for, say, understanding why climate change is influenced by human behaviour might vary between a high-school class and an academic colloquium. Focusing on comparative judgements about degrees of understanding bypasses such complications. For example, whilst the threshold for being ‘tall’ might vary between a basketball court and a nursery school, judgements as to whether LeBron James is taller than a small child remain true across contexts. Similarly, whilst the standards for outright ascriptions of understanding why the climate change is influenced by human behaviour might vary between a high-school class and an academic colloquium, the comparative judgement about whether the high-school student or the academic has better understanding does not vary. By excluding contextual variables, the focus on comparative gradable judgements facilitates a cleaner focus on the underlying metaphysics of understanding. And, as a final advantage, the comparative approach I suggest allows us to issue instructive verdicts in cases where two agents *both* understand outright at some salient context, but possess different degrees of understanding.

In addition to these advantages, focusing on degrees of understanding allows us to identify some commitments of reductionist theories about understanding that have not been fully unpacked in the literature. Once we have noted these commitments, they can then be used to provide a useful testing procedure for the particular version of reductionism we are examining. The commitments I have in mind can be best appreciated if we consider a particular presification of reductionism, such as the

knowledge-based version. Consider: if the extent to which an agent understands is to be fully explained by what the agent knows, this entails a variety of theses about what can bring an agent to possess a *different degree* of understanding. In terms of Sliwa's metaphor—just as the physicalist claims that mental changes must be explained by some physical change, reductionists about understanding must hold that changes in understanding must be explained by whatever epistemic state is posited as the reductive base of understanding. This line of thought yields three principles that any reductionist theory of understanding will require a version of.<sup>14</sup> (Let  $\phi$  be a placeholder for some epistemic state, and  $\phi p$  be a placeholder for being in that state with relation to some proposition).

**IMPROVEMENT:** Only acquiring new (or better-connecting)  $\phi$  can directly improve your understanding; acquiring epistemic states other than  $\phi$  cannot directly improve your understanding.<sup>15</sup>

**CONVERSION:** If  $\phi p$  (partly) constitutes your understanding, your understanding will be impoverished if your relation towards  $p$  is replaced by some other relation that does not entail  $\phi p$ .

**IMPOVERISHMENT:** Only losing  $\phi$  (or connections between  $\phi$ ) can directly impoverish your understanding; losing epistemic states other than  $\phi$  cannot directly impoverish your understanding.

These principles can be used to test specific reductionist theories in the following way: substitute  $\phi$  for whichever epistemic state the reductionist claims is the basis of understanding. This will give us three principles that are constitutive of the particular reductionist theory in question. And if any of these principles turn out to be false, then so does the reductionist theory that entails them.

---

<sup>14</sup> An anonymous reviewer suggests that Kelp's view, on which understanding is measured against approximating maximally systematic knowledge, might avoid being committed to such principles. For instance, it might be the case that an agent who believes that  $p$  (where  $p$  is some true proposition) better approximates maximally systematic knowledge than an agent who has no doxastic attitude towards  $p$ . I think this is a promising route for a reductionist to explore. However, pursuing an approach that is not committed to the principles I outline risks abandoning the metaphysical ambitions of reductionism. For, if a knowledge-based reductionist theory of understanding holds that understanding can improve in virtue of acquiring a non-knowledgeable belief, then it is not clear that such a theory is offering a statement about the *metaphysical dependence* of understanding on knowledge.

<sup>15</sup> I include the qualification regarding 'direct' improvement (and below, impoverishment) of understanding to respect the fact that reductionists are interested in metaphysical rather than merely incidental relationships between understanding and other epistemic states. Thus, **Improvement** and **Impoverishment** are not violated by cases where a subject acquires or loses some epistemic state other than  $\phi$  and this brings about an *indirect or mediated change* in their understanding. For instance, an indirect impoverishment might occur when a subject loses an unrelated belief and this happens to start a psychological process (e.g. a crisis of confidence) that results in their losing further knowledge relevant to the object of understanding. We should regard the loss of the initial belief as an indirect cause and the loss of the knowledge as the direct cause of the change in understanding. Hence, the principle cannot be violated by the former.

### 3. Against knowledge-based reductionism

Having outlined a methodology to test reductionist theories, we can now turn to scrutinise the knowledge based variant. It is worth noting, just to set it aside, that one well-known but controversial argument against **URK** is that non-factive epistemic states can contribute to an agent's degree of understanding.<sup>16</sup> I will be developing a different criticism. The following theory will be our target:

**URK:** Knowledge is all there is to understanding—the extent to which an agent understands-why  $p$  can be explained wholly in virtue of what they know.<sup>17</sup>

And, substituting knowledge for our epistemic placeholder, we get the following three principles:

**K IMPROVEMENT:** Only acquiring new (or better-connecting) knowledge can directly improve your understanding; acquiring epistemic states other than knowledge cannot directly improve your understanding.

**K CONVERSION:** If knowing  $p$  (partly) constitutes your understanding, your understanding will be impoverished if your relation towards  $p$  is replaced by some other relation that does not entail knowing  $p$ .

**K IMPOVERISHMENT:** Only losing knowledge (or connections between knowledge) can directly impoverish your understanding; losing epistemic states other than knowledge cannot directly impoverish your understanding.

To summarise, I defend the following thought—whilst reductionism is compelling, the exclusive focus on knowledge is unwarranted. If we accept that the degree to which one understands why  $p$  (where  $p$  is some substantive fact or phenomena) depends on a network of other cognitive commitments, it is not plausible to hold that *every* commitment that contributes to an agent's understanding will be a known proposition. This thought seems especially compelling when it comes to a high level of understanding regarding some complex phenomenon. A historian's understanding of why the Roman Empire collapsed will be comprised by a vast range of her beliefs and assumptions. Each member of this vast network has a variety of aetiologies—inference, abduction, testimony of varying shades of reliability, and so forth. On pain of an overly liberal epistemology, I think we should reject the suggestion that each of these commitments must be known.

To press this objection I scrutinise an example of moral understanding recently discussed by Sliwa (2017), a prominent defender of **URK**. The instance in question is understanding <why eating meat is

---

<sup>16</sup> See Elgin (2007), Grimm (2012) and Rancourt (2017) for useful discussion. It is worth underlining that the method for testing reductionism I proposed can also be applied to reductionist theories that do not just focus on factive epistemic states.

<sup>17</sup> Again, I take a well-connectedness proviso as implied.

wrong>.<sup>18</sup> Sliwa claims that cases in which someone better understands why meat-eating is wrong than someone else are explained by that person having *more knowledge of the things that make meat-eating wrong*: for example, knowing that the meat industry involves maltreatment of animals, and/or knowing that is bad for the environment. This is a plausible idea—knowing more typically results in greater understanding. However, I think this reductionist move forces a broader perspective than its proponents have appreciated. If we accept that one’s degree of understanding why *p* is explained by certain epistemic relationships towards various relevant considerations, then there will turn out to be many such relevant considerations—and it will not turn out to be plausible to suppose that every epistemic relation towards these is that of knowledgeability. A case will elucidate.

### 3.1 Putative Counterexample

As we are focusing on comparative judgements about understanding, it will be helpful to have an initial reference point. I assume these characters to be epistemic duplicates: they start in the same epistemic position and have identical dispositions to respond to new information.<sup>19</sup>

*Annie, Brian & Carl:* Each of these characters are fairly unimaginative inquirers, but they do have some knowledge. For example, they know that the environment is fragile and sensitive to human behaviour, that methane is a gas that (somehow) is produced by certain industries such as petroleum manufacturers, that methane is bad for the environment because it (somehow) heats the globe, and that supporting practices (e.g. the petroleum industry) that harm the environment is morally wrong. They don’t currently think that the meat industry harms animals or the environment.

Now, consider one’s epistemic judgements should the following occur:

*Annie:* Annie sees a headline in the *New York Times* (a paper she knows is highly reliable) that reads—“MEAT INDUSTRY RESPONSIBLE FOR MORE METHANE THAN PREVIOUSLY THOUGHT”. Without reading on, she comes to believe <the meat industry produces methane> and on this basis infers that <the meat industry is morally problematic in the same way as the petroleum industry>.

*Brian:* Brian finds himself surfing a conspiracy website—featuring articles about fake Moon Landings and Paul McCartney dying in 1966. He sees a headline that says: “MEAT INDUSTRY METHANE DESTROYING THE PLANET: CLICK HERE FOR MORE”. Without clicking, he comes to believe <the meat industry produces methane>, and on this

---

<sup>18</sup> A salient question is whether moral understanding might be *sui generis* compared to understanding some non-evaluative state of affairs. Here, I follow the contours of the existing debate which discusses examples of moral understanding as evidence for general theses about understanding why. Addressing the difference between the epistemology of evaluative and non-evaluative domains is outwith the scope this paper.

<sup>19</sup> This is an artificiality, but not a vicious one. The same point could be made by taking one character and discussing counterfactual scenarios.



basis infers that <the meat industry is morally problematic in the same way as the petroleum industry>.

*Carl*: Nothing happens. Carl retains his ignorance regarding the ethics of eating meat, neither believing that it harms animals nor is bad for the environment.

### 3.2 Discussion

The case is set up such that the following is the case:

- (1) Brian's belief that <the meat industry produces methane> is not knowledgeable.

This is because Brian forms this belief as a result of trusting the headline, a classic piece of unreliable testimony (suppose the author of the article lacks knowledge, and the article itself contains many falsehoods). Thus, the belief-forming method Brian uses is not reliable, there are close worlds in which the headline would have contained falsehoods and been believed by Brian regardless, and Brian also benefits from environmental epistemic luck in believing the headline containing truth rather than directing his attention towards the various false headlines surrounding it.<sup>20</sup>

In addition, by stipulation it is true that:

- (2) All three started at epistemic parity.

Further, I argue it is intuitive that:

- (3) Annie and Brian *improved* their understanding of why eating meat is wrong whilst Carl did not.

Annie and Brian both gain the belief that the meat industry is morally problematic, and infer this belief from the truth that it produces environmentally harmful methane, and their background belief that supporting environmentally harmful practices is wrong. Both can now provide a rudimentary correct explanation of why eating meat is wrong. Carl, by contrast, has no such belief and can cite no reasons in favour of the morally problematic status of eating meat. Comparatively, Annie and Brian have better understanding than Carl. Furthermore, I suggest:

- (4) Annie and Brian understand why eating meat is wrong equally well.

If you accept (3), it is very difficult to deny (4). As they started at epistemic parity, share identical epistemic dispositions, and acquired beliefs that have *exactly the same content*, the natural assumption should be that they now understand why eating meat is wrong to the same extent. For instance, each possesses identical dispositions to explain why you should not eat meat and to justify actions performed

---

<sup>20</sup> See footnote 4 for references on understanding and luck. However, these discussions are pitched largely at early 'equivalence' theories; not the reductionist views we target here.

on the basis of their new moral belief. They represent the world in exactly the same way; the only discernible difference between them is in the aetiology of one of their beliefs.

Taken together, (1) – (3) spell trouble for **URK**. For, it seems like we have a case in which Brian’s understanding why *p* improved in virtue of acquiring a true belief that he does not know. In the case, Brian’s belief that <the meat industry produces methane> is central to his degree of understanding—without this belief, he would revert back to his former epistemic position and lose the increased understanding we attributed him. Indeed, Brian understands why eating meat is wrong *better* than somebody who simply takes it on reliable testimony that eating meat is wrong because it is bad for the environment, where this person lacks any beliefs about methane, noxious gases, the globe heating up, and so forth. The degree to which one understands why eating meat is wrong depends not only on whether one believes that it harms the environment, but also on why one believes that it harms the environment. Given this, it is hard to avoid the conclusion that Brian’s increased understanding is partly constituted by his non-knowledgeable true belief that <the meat industry produces methane>. Thus, we have found a putative counterexample to **URK**—the thesis that understanding reduces just to what is known—by showing that the principle of **K IMPROVEMENT** is false; one can increase one’s understanding (partly) in virtue of acquiring a non-knowledgeable true belief. Once we grant this, it is easy to see why **K IMPOVERISHMENT** is also false. Were Brian to stop believing that the meat industry produces methane, his understanding would deteriorate in virtue of losing a commitment that falls short of knowledge.

Moreover, (4) states that Brian is on epistemic parity with Annie, who *knows* the same content—<the meat industry produces methane>—that Brian merely truly believes. Once we grant this, **K CONVERSION** becomes dubious. **K CONVERSION** predicts that Annie’s understanding would deteriorate if she lost knowledge but retained belief in the fact that the meat industry produces methane. But this seems doubtful if we grant that Brian understands equally well despite not knowing this same content. **K CONVERSION** might occur, for instance, within any framework under which knowledge can be *defeated* or can be degraded via exposure to *high-stakes*.<sup>21</sup> For example, Annie could lose knowledge that the meat industry produces methane via *normative defeat*—when there is some additional belief she ought to have entertained that would either conflict with, or undermine, her initial belief.<sup>22</sup> Consider a case where Annie learns that the edition of the *New York Times* she consulted was full of errors, but pigheadedly continues believing regardless. On various theories of defeat, Annie would lose knowledge that the meat industry produces methane. But, I suggest, losing this knowledge (whilst maintaining true belief) would not impoverish her understanding of why eating meat is wrong.

---

<sup>21</sup> Hills (2015) provides discussion of the relationship between epistemic defeat and understanding.

<sup>22</sup> Some knowledge-first epistemologists argue against orthodox theories of defeat. See Lasonen-Aarnio (2014) for a good example. My argument stands or falls independent of my discussion of **K CONVERSION** and defeat—undermining **K IMPROVEMENT** or **K IMPOVERISHMENT** suffices to refute **URK**.

### 3.3 Objections Considered

Some proponents of **URK** might derive it from a broader knowledge-first framework and attempt to explain away our earlier judgements. There are a number of possible strategies, of which only some can be considered here. I will take it to be implausible to suggest that both Annie and Bob have—on par with Carl—no understanding whatsoever. Furthermore, insofar as Annie and Brian believe exactly the same content, and they intuitively understand why eating meat is wrong to the same degree, I assume we prefer a unified explanation of the epistemic improvement they have made.<sup>23</sup> As such, the best response will be to attempt to show that their improved understanding is explicable by **URK**—there is something they know that explains our judgements.

*Conditional Knowledge?* An initial objection might run as follows: the understanding we attribute can be fully explained by *conditional knowledge* of the form <if the meat industry produces methane, then it would be bad for the environment, hence ...>. One might press this objection by supposing that the testimony acquired by Annie and Brian only *better connects* knowledge they already possess, and by suggesting that just this conditional knowledge explains their ability to infer that eating meat is morally problematic.

However, this objection is untenable. We stipulated that all three characters have this conditional knowledge at the outset (consider their attitude towards the petroleum industry) and yet Annie and Brian improve their understanding without acquiring new conditional knowledge. The issue is that, by itself, conditional knowledge doesn't suffice to explain understanding-attributions regarding a *specific* phenomenon. To exemplify: imagine that, unbeknownst to anybody, domestic cats produced vast quantities of methane. You and I *know* that methane is bad for the environment, and *know* that if supporting the domestic cat industry produced lots of methane, then keeping cats would be bad for environment and hence morally deleterious. But if it were discovered tomorrow that domestic cats produce lots of methane, we would not say that 'we understood why keeping domestic cats is wrong all along'. At best, conditional knowledge only puts one in a position to understand.

*Knowledge-acquisition?* A second more powerful objection is that Annie and Brian improve their understanding because they both come to know something Carl does not. Namely, *they both come to know* either that <the meat industry is bad for the environment> (MIBE) or <eating meat is morally wrong> (EMMW). And, on this objection, only this knowledge (part) constitutes their understanding rather than any mere true belief.

In the vignette, I avoided the question of whether the characters acquire knowledge of (MIBE) or (EMMW). In the case of Brian, affirming this would rely on the dubious assumption that we invariably gain *knowledge* by making inferences from non-known premises. If Brian is to know (MIBE), this

---

<sup>23</sup> Although I will consider and reject a disjunctive explanation, at the end of this section.

would involve claiming that Brian knows <the meat industry is bad for the environment> via the use of a premise he knows <if an industry produces methane then it is bad for environment> and a (true) premise that he does not know <the meat industry produces methane>. What we should say about (EMMW) is substantially more complex and will depend on what we should say about (MIBE); as such, I will only focus on the latter.

Whether we can gain knowledge through mixed inferences where some premises are known and some are not is a difficult topic.<sup>24</sup> However, for our purposes, we can simply note that the schema this objection relies upon is not knowledge-producing *in general*. To see this, consider the following schema (let K = knowledge and let B = truly believed):

- (1) **K**: If X then Y
- (2) **B**: X

---

- (3) **K**: Y

It is not hard to appreciate that there are many instances of this schema where the subject does not acquire knowledge. For example:

- (1) I know that if the Butler's prints are on the knife, the Butler is the murderer.
- (2) I believe, spuriously, that the Butler's prints are on the knife.

---

- (3) I deduce that the Butler is the murderer.

Clearly, I do not know the Butler is the murderer here even if it happens to be true. So, it is far from obvious that we should grant that Brian *knows* the meat industry is bad for the environment or that eating meat is morally wrong. Certainly, given his credulity, his route to these beliefs does not seem to have the properties of being reliable, counterfactually robust, or excluding salient error-possibilities. The proponent of **URK** needs to show that there is no way to describe the case such that Brian fails to acquire knowledge that the meat industry is bad for the environment or that eating meat is morally wrong. This seems like a potentially quixotic battle, but I leave it as a challenge for proponents of that theory.

The second more fundamental problem with this objection is that it isn't clear that attributing some extra piece of knowledge suffices to defuse our initial worries. As we have pressed, the challenge is to explain epistemic asymmetries, and the strategy adopted by **URK** was to say that an agent's understanding is merely a matter of what they know. However, even if we accept that Annie and Brian both know either (MIBE) or (EMMW), we can conceive of further cases where some other agent has better understanding of why eating meat is wrong in virtue of having a more fine-grained appreciation of why the meat industry is bad for the environment. **URK** will again have to identify and appeal to

---

<sup>24</sup> There is a literature on this issue stemming from Warfield (2005). However, the cases found in that debate are structurally different from those at issue here, often focusing on false premises.

further knowledge to explain this. But the puzzle can be reiterated again, taking cases where an even more extensive and complicated body of belief is at issue. And so on. The worry is that the proponent of **URK** ends up being committed to insisting that every single one of the many beliefs relevant to explaining any conceivable epistemic asymmetry towards one target fact must count as knowledge, or they will be unable to explain such epistemic asymmetries. This is a dilemma: either **URK** cannot fulfil its explanatory ambitions, or it ends up licensing overly liberal knowledge-ascriptions.

*Proximity to knowledge?* A further objection might be that Annie's understanding improves because she knows, whilst Brian's improves because—by acquiring a true belief—he comes *closer to knowing* an important explanatory fact about the ethics of eating meat.

This objection is problematic insofar as it belies the intuition that both Annie and Brian understand equally well. However, even if one is willing to grant this on broader theoretical grounds, the proximity to knowledge objection is independently dubious. We can imagine agents using a variety of similar, yet differentially reliable processes, to form exactly the same belief. Such an objection would yield the verdict that understanding increases incrementally with the reliability of the belief-forming process used. I lack a knock-down argument against such a position, but register my doubt that it is plausible to modify ascriptions of understanding why on the basis that one of the many beliefs contributing to this understanding was formed using an incrementally more or less reliable process, with no difference in the content of what is believed. Such incremental differences in reliability would not be transparent to either the agent who understands or to those who attribute understanding to them.

*A disjunctive explanation?*<sup>25</sup> When introducing constraints on possible objections, I suggested that we want a unified explanation of the epistemic improvement made by Annie and Brian. However, one might question whether excluding disjunctive explanations was really well-motivated. For instance, demanding a unified explanation of Annie and Brian's epistemic improvement rules out responses like the following: Annie understands better in virtue of knowing *p*, whilst Brian understands better in virtue of knowing some different body of knowledge *q*, where *q* is reflexive knowledge such as: <according to a source he trusts, the meat industry produces methane> and <it is very likely on his evidence that the meat industry is morally problematic in the same way as the petroleum industry>. I'll end by briefly arguing against this sort of disjunctive explanation.

An initial reason to be sceptical of such an approach is that it is far from clear that appealing to reflexive epistemic attributions really provides much succour to the defender of **URK**. We might conceive of a case in which Brian *merely believes rather than knows* that his evidence supports his conclusions about the meat industry. For instance, we could suppose that Brian has just had a conversation with a sensible friend who warns him against trusting unreliable websites, but that Brian irrationally disregards this

---

<sup>25</sup> I am grateful to an anonymous referee for comments and suggestions that were invaluable for this section.

evidence when deciding whether or not to believe the conspiracy website about the meat industry. I think this would be bad news for our evaluation of Brian's intellectual character, but would not impoverish his object-level understanding of the meat industry.

Even if our intuitions go hazy when we consider cases involving higher-order beliefs about our evidence, there are more fundamental reasons to look for a unified approach, and avoid appealing to such reflexive knowledge. Namely: (i) intuitively, Annie and Brian have equal degrees of understanding; so we should expect that they know the same amount about the meat industry, and (ii) if it is legitimate to ascribe reflexive knowledge to Brian, it is surely legitimate to also ascribe it to Annie, in addition to her object-level knowledge about the meat industry. And so we are again faced with the burden of explaining why they appear to understand equally well, despite asymmetries in what they know. An anonymous referee suggests that, at this juncture, defenders of **URK** might attempt to reject the judgement that Annie and Brian understand equally well by appealing to the context-sensitivity of understanding attributions. One such attempt might try to identify a context at which the outright judgement 'A and B understand equally well' comes out false. A second such attempt might claim that gradable attributions can themselves be context-sensitive. For instance, in the context of discussing an unequal race, it might be permissible to make assertions like: "A is as far from the starting point as B and both are further than C" even if A is 20.4km away X whilst B is 19.7km away. Thus, one might wonder whether attributions of the same degree of understanding might be permissible in some contexts but not in others. These possibilities create the worry that the apparent equality in understanding might tacitly depend on assessing this statement at a less demanding context—perhaps at a more demanding context, this judgement might be false. However, if such contextual factors did explain our intuition about parity between Annie and Brian, one would have to suppose that the context at which we were evaluating Annie and Brian earlier involved less demanding standards of accuracy than other readily conceivable standards. But it is hard to see why this would be so. After all, the context of the present paper is an examination of the relationship between knowledge and understanding, where we are taking particular care to provide the most accurate description of the epistemic position of the characters we discuss. We are not in a situation in which loose-talk is permissible because subtle distinctions are unimportant (such as when identifying the likely winner of an uneven race). Rather, we were explicitly attending to the subtleties of the case. So I am not tempted by this alternative explanation. There may be further mileage in this objection, but I leave it to defenders of **URK** to take up this challenge.

#### **4. Concluding remarks**

This paper outlined a promising framework for investigating understanding on which it reduces to other more familiar epistemic states, situated it in relation to other prominent views, and then developed a way to test different variants of this theory. We found that a recently influential version of reductionism,

on which understanding reduces to a body of knowledge, was subject to a routine counterexample. However, the knowledge-based view is but one version of reductionism and, I suggest, the basic reductionist insight pressed by proponent of the knowledge-based view is a good one. Future work must investigate other ways to develop reductionism about understanding—here, I have provided a methodology that enables us to carry out these projects with greater precision.<sup>26,27</sup>

## References

- Bourget, David (2017). The Role of Consciousness in Grasping and Understanding. *Philosophy and Phenomenological Research* 95 (2):285-318.
- Elgin, Catherine (2007). Understanding and the facts. *Philosophical Studies* 132 (1):33 - 42.
- Grimm, Stephen (2006). Is understanding a species of knowledge? *British Journal for the Philosophy of Science* 57(3):515-535.
- (2012). The Value of Understanding. *Philosophy Compass* 7 (2):103-117.
- Understanding as Knowledge of Causes.” In *Virtue Scientia: Essays in Philosophy of Science and Virtue Epistemology*. Ed. Abrol Fairweather. Special issue of *Synthese*. (2014)
- C. Baumberger & S. Ammon (eds.), (2017) *Explaining Understanding*. Routledge.
- Hills, Alison (2016). Understanding Why. *Noûs* 49 (2):661-688.
- Kelp, Christoph (2014). Knowledge, Understanding, and Virtue. In Fairweather, Abrol (ed.) (2014). *Virtue Epistemology Naturalized: Bridges between Virtue Epistemology and Philosophy of Science*. *Synthese Library*, Vol. 366, Springer.
- (2015). ‘Understanding Phenomena.’ *Synthese* 192: 3799–816.
- (2017). Towards a Knowledge-Based Account of Understanding. In S. Grimm, C. Baumberger & S. Ammon (eds.), *Explaining Understanding*. Routledge.
- Khalifa, K. (2013). ‘Is Understanding Explanatory or Objectual?’, *Synthese* 190: 1153–1171.
- Kvanvig, Jonathan (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lasonen Aarnio, Maria (2014). Higher Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research* 88 (2):314-345.
- Lawler, Insa (2016). Reductionism about understanding why. *Proceedings of the Aristotelian Society* 116 (2):229-236.
- Pritchard, Duncan (2009). Knowledge, Understanding and Epistemic Value. *Royal Institute of Philosophy Supplement* 64 (64):19-43.
- (2010). *The Nature and Value of Knowledge: Three Investigations*. Oxford University Press.
- Rancourt, Benjamin T. (2017). Better Understanding Through Falsehood. *Pacific Philosophical Quarterly* 98 (3).
- Riaz, Amber (2015). Moral understanding and knowledge. *Philosophical Studies* 172 (1):113-128.
- Sliwa, Paulina (2015). IV—Understanding and Knowing. *Proceedings of the Aristotelian Society* 115 (1pt1):57-74.
- (2017). Moral Understanding as Knowing Right from Wrong. *Ethics* 127 (3):521-552.
- Warfield, Ted A. (2005). Knowledge from falsehood. *Philosophical Perspectives* 19 (1):405–416.
- Zagzebski, Linda (2001). Recovering Understanding. In M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press

---

<sup>26</sup> My thanks to Jessica Brown, Miguel Egler, and Matt McGrath for reading and commenting on earlier versions of this paper. I am also grateful to my friends and colleagues at the University of St Andrews for support and insightful discussions.

<sup>27</sup> This research was completed with the help of financial assistance from the Carnegie Trust for the Universities of Scotland.