

0. Introduction¹

What is the relationship between understanding and knowing? Much of the recent literature argues that understanding is a distinctive cognitive state: unlike knowledge, understanding comes in degrees, it is immune to certain forms of epistemic luck, and it has a different relationship to testimony. In light of this, it has been suggested that we need to make room in our epistemic theorizing for a distinct notion of understanding.

My aim in this paper is to push back against this trend and to defend reductionism about understanding. I argue that knowledge is both necessary and sufficient for understanding. I consider the two main arguments against reductionism – the Argument from Epistemic Luck and the Argument from Testimony – and argue that neither is successful. Finally, I argue that knowledge provides us with all the resources we need to account for some of the central features of understanding.

1. Some Clarifications

Before we start, it will be helpful to have a more precise statement of what reductionism about understanding endorses and non-reductionism denies. As a first gloss, reductionism claims that knowledge is all there is to understanding; there is no need to stipulate a novel cognitive state that goes over and beyond knowledge. But note that understanding, just like knowledge, takes both *wh*-clauses as well as *that*-clauses. An agent can understand what happened or why it did; she can also understand that something happened. The reductionist that I shall defend maintains that instances of understanding reduce to the corresponding instances of knowledge. Consequently, my aim will be to show that corresponding instances of knowledge are both *necessary* and *sufficient* for the relevant instance understanding in question. Thus, I will argue that:

An agent understands why *p* *if and only if* she knows why *p*.

An agent understands what happened *if and only if* she knows what happened.

An agent understands that *p* *if and only if* she knows that *p*.

etc.

In what follows, I will for the most part focus on understanding and knowing *why p* but my main arguments should straightforwardly carry over to understanding what and understanding that. For simplicity, I will sometimes refer to the reductionist claims as the claim that ‘understanding requires knowing’. But it’s important to keep in mind that throughout, I take the reductionist to be committed to the view that instances of understanding reduce to the corresponding instances of knowing.²

¹ Various ideas that found their way into this paper were presented at the CamPoS seminar and the Roles of Knowledge Workshop at the University of Cambridge, the workshop on “Understanding and the A Priori” at the University of Birmingham, the University of East Anglia, MIT, and the Workshop on Religious Testimony in Oxford. I’m grateful to the audiences and organizers for their comments and feedback. Special thanks go to Emil Moeller, Alison Hills, Hallvard Lillehammer, Scott Sturgeon, John Maier, Katia Vavova, Jack Marley-Payne, Kenny Walden, Michael Bloome-Tillman, and Robbie Williams.

² I steer clear of the debate about the role of explanation in understanding. Even if reductionism about understanding is true, there is still a substantive further question as to whether understanding requires knowledge of an *explanation*. Philosophers of science sympathetic to the claim have predominantly focussed on understanding why. Plausibly an explanation why *p* is important to understanding why *p*. But it’s less clear that it’s central to understanding what happened or what to do. See Lipton [2009] for arguments that explanation is not required for understanding why. See Khalifa [2012] and [2013], Strevens [forthcoming] for criticism.

With reductionism on the table, we can give a first gloss of non-reductionism. Some non-reductionists deny that understanding requires knowing.³ All deny the sufficiency claim. In the next section, I make a positive case for the necessity claim. I then defend the sufficiency claim.

2. Is Knowledge Necessary for Understanding?

According to a number of influential non-reductionists, such as Kvanvig, Pritchard, and Hills one important difference between understanding and knowledge is its relationship to epistemic luck. The epistemic luck in question is of a particular kind: environmental luck. A paradigm example of environmental luck is the fake barn case. The agent formed a true belief that there's a barn in front of her. But she could have just as easily looked at a fake barn and hence come to have a false belief. Pritchard, Hills, and Kvanvig have all suggested that while the presence of epistemic luck is incompatible with knowing, it is compatible with understanding. And so it is possible for an agent to come to understand why something happened without at the same time coming to know why it happened. Knowing why *p* is not necessary for understanding why *p*.

Hills argues as follows:

Suppose that your school has been sent a set of extremely inaccurate textbooks, which have been handed out to your class. But you are very lucky because there is only one that is accurate, and by chance you have it. You read in your book that Stalin was responsible for the deaths of millions of people. You draw the obvious conclusion that he was an evil person.⁴

According to Hills, since you do not know that Stalin killed millions, you do not know that he was evil because he killed millions. And so, you don't know why he was evil. However, Hills argues that you nevertheless understand why Stalin was evil:

After all, you believe that he was evil because he killed millions of people, and that is correct, and you have—let us assume—the ability to draw the conclusion that he was evil from the reasons why he was evil and to do the same in similar cases. So it seems that you can have moral understanding why *p* without having knowledge why *p*.⁵

And so, Hills suggests, what an agent understands can come apart from what the agent knows.

Similarly, Pritchard argues:

Does epistemic luck of this sort [environmental epistemic luck] undermine one's understanding in the way that it would undermine one's knowledge? I don't think that it does, since one did indeed find out the relevant facts in the right kind of way...[O]ne can gain an understanding of why one's house burnt down by speaking to the fire officer—even though one could just have easily been misled by someone who isn't the fire officer...⁶

We can summarize the argument as follows: in environmental luck cases, an agent can come to understand why *p*. But because knowledge is not immune to environmental luck the agent cannot come to know that *p* because of *q* and hence cannot come to know why *p*. Thus, knowing why *p* is not necessary for understanding why *p*.⁷

³ For example, Kvanvig [2003], [2009], Pritchard [2008].

⁴ Hills [2009], p. 204. Hills' case is modeled after Kvanvig's Comanche case. See Kvanvig [2003].

⁵ Hills, *ibid*.

⁶ Pritchard [2008], p. ?. Pritchard calls this type of luck "environmental luck". He argues that while "environmental luck" is compatible with understanding but not knowledge, the kind of luck that is present in Gettier scenarios precludes both knowing and understanding.

⁷ Others have argued that understanding does not require knowledge on the grounds that understanding is not factive. See Elgin [2007]. The argument that follows addresses this view, too.

3. Defending the Necessity Claim

Before offering a response to the argument from epistemic luck on behalf of the reductionist, I want to first make a positive case for the reductionist's necessity claim: that there can be no understanding without knowing. I will then return to the epistemic luck cases.

According to the non-reductionist there is a whole class of cases in which we intuitively attribute to the agent understanding while also intuitively judging that they lack the corresponding knowledge. But if this was the intuitively plausible way to evaluate agents in these situations, we would expect that we could just straightforwardly conjoin the sentences "I understand why Stalin was evil" and "I don't not know why Stalin was evil". And yet, consider the following:

I understand why Stalin was evil but I don't know why Stalin was evil.

I understand that Stalin killed millions but I don't know that Stalin killed millions.

I understand that Stalin was evil because he killed millions but I don't know that Stalin was evil because he killed millions.

All of these seem infelicitous. Could this just be some special feature of the first-personal case that the non-reductionist can explain by appealing to pragmatic considerations? That seems unlikely, since the third personal case sounds just as bad. Consider the following case:

Journals: Jane is a surgeon who picks up a scientific journal from a stack and reads an article about a novel therapy against hepatitis B. She learns that the medication is highly effective – much more so than the established therapies – and that this is so because it targets a particular protein of the virus. As it turns out, this journal was the only accurate one in the stack. Had Jane picked up any of the others, she would have come to believe that this therapy was no better than the established alternatives.

As before, the non-reductionist should maintain that this is the case in which Jane intuitively neither knows that the medication is effective nor why it's effective. And yet, the following seem infelicitous:

Jane understands why the hepatitis medication is effective but she does not know why it's effective.

Jane understands that the new medication is effective because it targets a particular protein of the virus but she does not know that it's effective because it targets a particular protein of the virus.

Jane understands what the most effective therapy against hepatitis B is but she does not know what the most effective therapy against hepatitis B is.

Jane understands how to treat hepatitis B but she does not know how to treat hepatitis B.

Note that there is no obvious way to explain away their unassertability by appealing to pragmatic considerations. If there was, we would expect that there was some way to embed the sentence which would render it felicitous. But the following sound just as bad:

If Jane understands why the hepatitis medication is effective but she does not know why it's effective, then there is something that she does not know.⁸

This suggests that the defect here is semantic: the reason why these sentences sound infelicitous is that they are contradictory. This is because understanding *entails* knowing.

For the reductionist this datum is easily accommodated: understanding should entail knowing since, according to the reductionist, knowing is necessary for understanding. But it creates problems for non-reductionism. This is because if it is indeed so intuitive that in cases like *Journals* we judge the agent to understand without knowing, then we would expect the semantics for knowing and understanding to reflect that. After all, non-reductionists do not take themselves to be introducing a novel theoretical concept and stipulating its properties. Rather, they take themselves to describe an

⁸ Unlike the Moorean sentence "It's raining but I don't know that it's raining," which, in certain embeddings, does sound felicitous: "If it's raining but I don't know that it's raining, then there is something I don't know."

epistemic phenomenon that already plays a pervasive and important role in our ordinary epistemic practice (including our scientific practice) and that traditional epistemology in its single-minded focus on knowledge has hitherto neglected.

Thus, for example, Kvanvig argues that, in contrast to knowledge, understanding is characterized by its internal component, rather than a non-accidental connection to the external facts and that this is why understanding and knowing come apart in cases such as those above:

What is distinctive about understanding, once we have satisfied the truth requirement, is internal to cognition. It is the internal seeing or appreciating of explanatory and other coherence-inducing relationships in a body of information that is crucial for understanding. When we think about knowledge, however, our focus turns elsewhere immediately if we have learned our lessons from the Gettier literature: We think about the possibility of fortuitousness, of accidentality, of being right but only by chance. We focus, that is, on what kinds of further external connections there are between mind and world, beyond the fit required for the belief to be true.⁹

Kvanvig does not take himself to be making stipulations here; he takes himself to be articulating the differences in our intuitive grip on understanding and knowing.

Could the non-reductionist respond to this recalcitrant datum by arguing that cases like *Journals* are rare and that therefore situations in which understanding and knowing will come apart are going to be few and far in between? If so, then perhaps we may still often infer that an agent knows why *p* from the fact that she understands why *p* – it's just that the inference fails for a circumscribed class of very special and rare cases. But this will not do. The problem is that these kinds of situations are decidedly *not* rare. To see this, we need to have a closer look at the kind of epistemic luck at work here. The cases on which non-reductionists rely are cases in which an agent is lucky to have come to possess non-misleading evidence when she could have easily acquired different, misleading evidence, on which she would have formed a false belief. Thus, Jane is lucky to have picked up a journal which accurately reports how the study was conducted, the statistical results, the mechanism of the drug, etc. She could have, easily, instead picked up a journal which would have told her quite different things about the drug: perhaps, it would have misreported the statistics for its effectiveness or it would have left out crucial information about how it compares to other drugs on the market. Similarly, in the textbook case discussed by Alison Hills, the agent is lucky to have picked up a history text book which accurately reports Stalin's crimes against humanity – as opposed to one that leaves this information out and just praises Stalin for liberating the working classes and defeating the Nazis.

Cases in which agents acquire non-misleading evidence when they easily could have acquired misleading evidence are in no way an unusual or ecologically rare occurrence. We often find ourselves in exactly that epistemic predicament. Just consider the following examples:

- The police detective is about to declare the case done; everything at the crime scene incriminates the jilted lover with a history of violent assaults. On her final walk through the crime scene, a very subtle clue catches her eye – a clue she had overlooked multiple times and could have easily continued to do so, had a big spider not just walked over it. On the basis of this clue, she correctly concludes that the jilted lover is innocent and catches the real murderer instead.
- Upon examining the patient, the doctor notices a very faint rash on her finger. She almost overlooked it but the patient just happened to hold her hand in the right position. If she had not noticed the rash, the doctor would have diagnosed the patient with condition X, since this is what all the other symptoms were pointing to. However, the rash indicates that the patient suffers from disease Y.

The police detective is lucky to have seen the subtle clue – had the giant spider not been in the right place at the right time, she would have overlooked it. Had she overlooked it, she would have

⁹ Kvanvig [2003], p. 198

charged the wrong man. Similarly, the doctor is lucky to have noticed the faint rash – had the patient not positioned his hand just right, she would not have seen it. And in this case, she would have been misled by the patient’s other symptoms. Other examples are easy to come by: the cheated boyfriend, who accidentally comes across the incriminating love note, the biologist who decides to look at the moldy probe and finds a new antibiotic,...

According to non-reductionism, these are all cases in which what an agent understands comes apart from what she knows. But then the non-reductionist cannot plausibly claim that situations in which an agent’s knowledge and understanding will come apart happen only very rarely. Quite on the contrary, these situations represent a common epistemic predicament. And so, the non-reductionist does not seem to have a good explanation for why sentences such as:

Jane understands why the drug is effective but she does not know why it’s effective, are infelicitous. Equally seriously, the non-reductionist’s verdicts about the cases above seem at odds with our intuitive judgments: the police detective and the doctor may have been lucky that they stumbled across the crucial pieces of evidence but, once they have stumbled across it, they do seem to know who committed the murder or which disease the patient has. Contrary to Pritchard, Kvanvig, and others this kind of epistemic luck seems to be compatible with knowledge.¹⁰

This suggests a response on behalf of the reductionist to cases like *Journals*. The reductionist can plausibly maintain that since the scientist’s epistemic situation in *Journals* is relevantly similar to that of the police detective and that of the doctor. Since very few of us would be willing to deny that the police detective knows who committed the murder, it’s plausible that the scientist also knows why the new drug treatment is effective. After all, the scientist formed this judgment based on sound and reliable evidence about the mechanism of the drug and how it was tested. Of course, she was lucky to acquire that evidence – she happened to pick up the right journal edition, just as the police detective happened to look at the right spot on the crime scene floor. But given that she has acquired it and has formed her beliefs in response to it, she knows both that it’s effective and why it’s effective.¹¹

Does the reductionist thereby commit herself to the claim that even the subject in fake barn country knows that there’s a barn in front of her? I don’t think so. The reductionist simply observes that any plausible account of knowledge will have to allow that in *some* cases – like that of the police detective or the doctor – agents can be lucky to know, lest it dooms us to widespread revisionism about our ordinary knowledge attributions. She then notes that cases like *Journals* are more similar to those cases than to the fake barn case. Spelling out the difference between cases like *Journals* and barn-cases is then a task for an account of knowledge.¹²

¹⁰ Grimm draws a similar conclusion in his [2006].

¹¹ One way to make the argument more precise is by appealing to safety: Knowledge requires safety; one knows that P only if one could not easily have believed P falsely by a relevantly similar method. The main way in which epistemic luck can threaten knowledge is by threatening safety. But not all kinds of epistemic luck lead to beliefs being unsafe. Since the police detective knows who committed the murder, her belief is safe. This means that in this case, the right way to individuate methods is relatively fine-grained: the method is not just “scanning the room” (since this would be unsafe). Rather that it includes the specific piece of evidence that the police detective chanced upon – say, the piece of fiber from the murderer’s clothing. But the case of *Journals* is structurally analogous to the case of the police detective. Therefore, in *Journals* the “relevant method” also includes the specific piece of evidence that the scientist read about – say, the specific studies or information about the mechanism of the drug. And so, the scientist’s belief is safe, too. It just fails to be “safely safe”.

¹² See Khalifa [2013] for a careful discussion of the differences between the kind of epistemic luck present in barn cases and the kind of epistemic luck to which non-reductionists appeal.

Note that there's an alternative strategy available to the reductionist, too. Rather than arguing that in *Journals* the agent has knowledge, she could try to debunk the intuitive verdict that in that case the agent has understanding. There is a plausible way to develop this strategy, too. In particular, the reductionist can argue that we need to distinguish two questions: whether Jane has *medical understanding* and whether Jane *understands why* the medicine is effective. Insofar as we are willing to attribute understanding to Jane in the case, it's in the former sense rather than the latter. And the reductionist can allow that Jane has medical understanding: after all, Jane knows a great deal about medicine and so she is in a position to know many medical truths that may not be accessible to the layperson. But Jane's having medical understanding is compatible with her failing to understand why the medicine is effective in this particular instance.

In this section I have defended the reductionist claim that knowing is necessary for understanding: an agent does not, for example, understand why a therapy is effective unless she knows why it's effective. She doesn't understand what the best way to treat hepatitis is unless she knows what the best way to treat hepatitis is. I have given a general reason to think that understanding and knowing come together: sentences such as "She understands why the car broke down but she does not know why the car broke down," seem semantically defective. To say that someone understands what happened, how it happened, or why it happened *entails* that they know what happened, how it happened, or why it happened. I have then suggested two ways in which the reductionist can respond to the epistemic luck cases on which the non-reductionists rest their case. Since the reductionist takes knowledge to be necessary for understanding, she can deny the verdict that agents in these epistemic luck cases lack knowledge or she can deny the verdict that agents in these cases have understanding. While I favor the former, both are plausible ways to go; which one the reductionist pursues will likely depend on her other theoretical commitment and the details of the case.

4. Is Knowledge Sufficient for Understanding?

Not all non-reductionists reject knowledge as a necessary condition for understanding.¹³ But virtually all non-reductionists maintain that knowing is not *sufficient* for understanding. In particular, all non-reductionists deny the following:

If an agent knows why p, then she understands why q.

If an agent knows what happened, then she understands what happened.

If an agent knows how p happened, then she understands how it happened.

Non-reductionists argue that the claim that knowledge is sufficient for understanding is vulnerable to counterexamples. In particular, they argue, knowledge and understanding have different relationships to testimony. An agent may come to know why based on testimony of a reliable source. But such testimonial knowledge typically does not suffice for understanding. Thus, Pritchard presents us with the following case:

Faulty Wiring: Suppose that I understand why my house burned down, know why it burned down, and also know that it burned down because of faulty wiring. Imagine further that my young son asks me why his house burned down and I tell him. He has no conception of how faulty wiring might cause a fire, so we could hardly imagine that merely knowing this much suffices to afford him understanding of why his house burned down. Nevertheless, he surely does know that his house burned down because of faulty wiring, and thus also knows why his house burned down.¹⁴

Along the same lines, Alison Hills has argued:

¹³ See, e.g. Grimm [2006] as well as [forthcoming].

¹⁴ Pritchard [2009], p. ? Similar cases are brought up by Kvanvig [2003], Elgin [2004], Strevens [forthcoming].

If you are attempting to gain knowledge, testimony can serve as the justification for your own belief but it is not usually a good way of acquiring...understanding. Understanding why p will not – cannot – have the same relationship with testimony as knowing why p.¹⁵

Why do knowing why and understanding why bear such a different relationship to testimony? According to the non-reductionist this is because knowing why is ultimately a matter of having the relevant propositional knowledge – propositional knowledge that can be transmitted via testimony. Thus, insofar as the child comes to know that the house burned down because of faulty wiring, he knows why it burned down. Understanding why, however requires more than having the relevant propositional knowledge. What exactly this “more” is differs on different non-reductionist views. Some suggest that understanding why p should be assimilated to a kind of practical knowledge or an ability.¹⁶ Others suggest that understanding involves a different *object* than propositional knowledge. While propositional knowledge involves a relationship to a proposition, the object of understanding is rather “something like the modal relationships that obtain between the properties (objects, entities) at issue.”¹⁷ An alternative suggestion is that while knowing why p and understanding why p may both ultimately be propositional, they involve a different cognitive attitude to the proposition in question – in the same way as assuming that p, entertaining that p, and knowing that p involve different attitudes to the same proposition. Understanding differs from knowing in that it requires not just “assent” to but rather a “grasp of” the proposition in question. As Strevens argues:

These [testimonial] cases show that the sort of grasping needed for understanding requires a more intimate acquaintance with the structure of the explanation than sometimes accompanies mere knowledge. It is not enough to know that one or more parts of, or conditions for, a correct explanation hold; their holding must be directly mentally apprehended.¹⁸

In what follows, I will focus on Pritchard’s example of parent and child. In terms of this example, we can thus summarize the non-reductionist’s case against sufficiency as follows:

Premise 1: In cases like *Faulty Wiring* parent and child both know that the house burned down because of faulty wiring.

Premise 2: If parent and child both know that the house burned down because of faulty wiring, they have the same propositional knowledge about why the house burned down.

Premise 3: If they have the same propositional knowledge about why the house burned down, then both know why the house burned down.

Premise 4: Nevertheless, there is an epistemic asymmetry between two agents: the parent understands why the house burned down and the child does not.

C1: This epistemic asymmetry cannot be explained in terms of the parent’s and the child’s knowledge. (2,3,4)

C2: And so, knowing why p cannot sufficient for understanding why p.

5. Defending the Sufficiency Claim

It is hard to dispute that there is an epistemic asymmetry between the parent and child in *Faulty Wiring*, thus premise 4 is uncontroversial. I will also grant premise 1 for the sake of the argument. This means that the reductionist can get off the boat in two places: premise 2 and premise 3.

¹⁵ Hills [2009], p. ?

¹⁶ See Hills [2009].

¹⁷ Grimm [forthcoming], p. 9. Grimm defends a specific form of reductionism: he argues that understanding is a species of knowledge that’s not reducible to propositional knowledge.

¹⁸ Strevens [forthcoming], p.? This suggestion can also be found in Grimm. See his [forthcoming], p. 10.

To deny premise (3), the reductionist could argue that knowing why is a special kind of knowledge that does not, in turn, reduce to propositional knowledge.¹⁹ If successful, such a defense would support at best a *modest reductionism*: the view that all instances of understanding reduce to corresponding instances of knowledge, albeit knowledge of a special kind that may not reduce to propositional knowledge. I will not pursue this line of thought here. Instead my goal will be to challenge premise (2) – the claim that if both parent and child know that the house burned down because of faulty wiring, they have the same propositional knowledge about why the house burned down. This strategy allows for a more *ambitious reductionism*, on which all instances of understanding why, what, and how ultimately bottom out in propositional knowledge.

First, note that prior to hearing the testimony of the firefighter, the epistemic situation of parent and child are asymmetrical. The parent presumably knows that faulty wiring is a fire hazard. Plausibly he knows that electric wiring can be faulty in a number of ways: for example, there may be insufficient grounding or the insulation of some cables may be frayed. He also presumably knows that there are a number of ways in which faulty can give rise to fire: the material adjacent to the cable might overheat or the electric wires can generate sparks. He also knows which other things give rise to house fires. He knows that people put their house on fire in order to get insurance money. He knows that some fires are caused because people forget to put out candles, or because they fall asleep holding a cigarette. Because of this background knowledge – knowledge that seems propositional – the father can entertain a number of epistemic possibilities as to what might have caused the fire: arson, faulty wiring, an unattended electric heater, a forgotten candle, a careless smoker.

Contrast this with the child's epistemic situation. Plausibly the child knows something about what causes fires: he knows that some fires are deliberately caused by people – such as when you light a candle, or when you make a bonfire – and that other fires just happen – as when lightning strikes a tree. And so the child, too, may discriminate between different possibilities for what might have caused the fire. It discriminates between the possibility that the house was deliberately set on fire by someone and the possibility that the house (somehow) erupted in flames by itself. But there is a lot that the child does not know: he does not know that frayed cables can cause sparks, that electrical heaters can incinerate nearby curtains, that people commit arson for insurance fraud. Thus, even before finding out that the fire was caused by faulty wiring, father and son differ in what they know about house fires. Because they differ in what they know, they start out with different epistemic possibilities for what may have caused the fire. There is, of course, significant overlap: both may discriminate between the possibility that someone set the house on fire intentionally, that someone set the house on fire unintentionally through negligence, or that something about the house caused the fire. But for one, the father's epistemic possibilities are plausibly more fine-grained. Second, he plausibly excludes some epistemic possibilities that are open for the child – such as that the fire was caused by a witch.

Since father and son differ in what they know prior to learning that the house burned down because of faulty wiring, they also differ in what they learn when they both come to know that the house burned down because of faulty wiring. For one, they differ in which possibilities this information allows them to exclude. So, by coming to know that the fire was caused by faulty wiring, the parent also comes to know that the fire was *not* caused by someone trying to commit an insurance fraud, by an unattended electric heater, or by a carelessly placed cigarette. In contrast, the child may simply come to know that no one deliberately set the house on fire. Second, they differ in which epistemic possibilities are left open for them: the parent may still wonder whether the faulty wiring in question was a matter of a broken fuse or frayed cable insulation. For the parent, learning that fire was caused by faulty wiring may leave many more and different possibilities open: insofar as he does

¹⁹ Such a view has been defended by Grimm [forthcoming]. Grimm draws a parallel between this special kind of knowledge and a priori knowledge.

not know what exactly faulty wiring is, he may imagine that it could involve the television or vacuum cleaner spontaneously exploding or erupting into flames.

The reductionist then need not deny that both parent and child both know that the house burned down because of faulty wiring. They both divide a range of alternative possibilities as to what could have caused the fire and although there are differences – the child’s range may be more limited in some ways and more inclusive in others (magic spells, for example) – there is significant overlap. The testimony of the firefighter allows both of them to locate the actual world on the right side of the divide.²⁰ Nevertheless, they *differ* in what they know about the house fire. The parent knows *more* about what caused the house fire than the child. And this knowledge is propositional in nature: it’s a matter of a *difference in content*. And so, we should reject the non-reductionist’s Premise (2).

This still leaves the reductionist with a positive explanatory task: she needs to show that the asymmetry in propositional knowledge can account for the intuitive datum that parent and child differ in understanding. For this, let me first offer a sketch of how we might fill in the reductionist story about the relationship between knowing and understanding more fully. Recall that according to the non-reductionist, when two agents differ in whether they understand why p this is because they differ in the *cognitive attitude* that they bear towards the answer to the question why p. The understanding agent “grasps” the answer, while the merely knowing agent believes it or “assents” to it. In contrast, the reductionist maintains that agents differ in what they understand in virtue of differing in what they know. Better understanding is simply a matter of more knowledge.²¹ Differences in understanding thus are explained by differences in how much the agent knows about the matter at hand. And this difference, in turn, is a difference in the *content* of what is known.

A reductionist account of understanding can thus accommodate the natural thought that understanding comes in degrees: understanding comes in different degrees because what is known comes in different amounts. And so, it can explain why it’s so natural to describe the epistemic asymmetry between parent and child in terms of a difference in understanding. Since the parent knows more about why the house burned down, he understands better why it burned down.

Let me address two objections. The first objection is that the reductionist is not really capturing the epistemic asymmetry between parent and child that non-reductionists point to. After all, reductionists maintain that knowing why p is sufficient for understanding why p. But then it follows that, insofar as the child does know that the house burned down because of faulty wiring, the child does understand why the house burned down. Doesn’t this go against our intuition that the child doesn’t have *any* understanding why the house burned down?

I don’t think that this *is* our intuition. Insofar as it’s plausible to attribute to the child the knowledge that the house burned down because of faulty wiring, we do attribute some degree of understanding of why the house burned down to the child. And the more we fill in the case to make it plausible that the child does not understand why the house burned down, the less plausible it is that the child knows why the house burned down. Thus, suppose that the child has truly not the faintest clue what faulty wiring could possibly be – from the child’s perspective the house burning down because of faulty wiring is compatible with its being struck by lightning, its being set on fire by someone, and with spontaneously erupting into flames. In this case, it really is clear that the child does not understand why the house burned down. But neither does it know that the house burned down

²⁰ While I am not committing myself to any particular account of propositions, a natural way to make this precise is to appeal to a Stalnakerian account. See Stalnaker [1984], p. 64-65 for related discussion.

²¹ In this case, I’m appealing to an intuitive understanding of knowledge coming in ‘different amounts’. The question how to make the notion of different quantities of knowledge precise is a further substantive question beyond the scope of this paper. See Treanor [2013] for an in-depth discussion of this problem as well as a positive proposal.

because of faulty wiring. (It may, of course, know that it was something that the grownups call “faulty wiring” that caused the fire – whatever *that* is.)

Insofar as it seems fine to say that the parent understands why the house burned down but the child only knows why it burned down, this may well be a matter of pragmatics. Thus, in situations in which we are explicitly contrasting “understanding” with “knowing”, we pragmatically raise the standard for *how much* knowledge is required for the former.

A second objection is as follows: if reductionism is correct, knowing why *p* entails understanding why *p*. And yet consider the following:

John knows why the house burned down but he doesn’t understand why the house burned down.

John knows what happened but he doesn’t understand what happened.

If knowing entailed understanding, we might expect these sentences to be infelicitous. But they seem perfectly assertible. And this, the non-reductionist may argue, casts doubt on the reductionist claim that knowing is sufficient for understanding.

Here too, the reductionist has a response. She can argue that, at least in some contexts, to say that an agent knows why *p* can be ambiguous between:

The agent knows that some sentence ‘*p* because of *q*’ is true,
and

The agent knows why *p* – for example, she knows that *p* because of *q*.

These contexts may arise particularly in cases of testimony. Simply knowing that some sentence ‘*p* because of *q*’ is true is not sufficient for understanding why *p* – understanding why *p* involves first and foremost knowing something *about p* rather than knowing something about sentences. And it’s for this reason that the sentences above are assertible. In support of this, note that we can draw the following inference:

If John knows why the house burned down but he doesn’t understand why the house burned down, then there is something (about the fire) that John does not know.²²

As a matter of fact, these inferences spell additional trouble for non-reductionists. This is because according to the non-reductionist lack of understanding need not generally imply lack of knowledge: rather it may just be that the agent fails to ‘cognitively grasp’ the known content in the special way that’s required for understanding.

6. Conclusion

Let me end by pointing out some advantages of reductionism over non-reductionism.

One advantage of reductionism is that it gives us a parsimonious account of understanding. It does not postulate a novel cognitive attitude to a proposition (“grasping”), nor does it require us to introduce a novel epistemic attitude with an altogether different kind of object (a modal relationship, for instance). Of course, parsimony must always be weighed against other theoretical virtues – in particular the explanatory gain to be had by introducing these notions.

Even without postulating such notions, reductionism has ample resources to account for some of the central features of understanding. Thus, it gives us a plausible account of degrees of understanding. Degrees of understanding correspond to different amounts of knowledge. An agent understands better why *p* insofar as she knows more about why *p*. A full reductionist account of understanding then requires a principled account of how knowledge should be measured.²³ Such an account is of interest quite independently of the question about the relationship between understanding and

²² I’m grateful to Bernhard Salow for helpful discussion on this point.

²³ See Treanor [2013].

knowing – after all, it's very natural to attribute more knowledge to the expert than to the lay person or to compare how much we know now with how much we knew before.

Reductionism also gives us a nuanced account of the relationship between testimony and understanding. According to reductionism, if testimony is a source of knowledge, then it is also a source of understanding. More precisely, if an agent can come to know what happened or why it happened based on testimony, then she can come to know what happened or why it happened by testimony. And this is a good thing: after all, we do rely on testimony for much of our understanding of the world: we read textbooks, we go to lectures, we look to experts for advice. To deny that in doing so we can come to have understanding would fly in the face of our ordinary epistemic practice. At the same time, the reductionist can recognize that there are limits to the degree of understanding that may be transmitted in a single testimonial exchange. What exactly an agent comes to know through a testimonial exchange and how much she knows after the exchange are not just a function of what the speaker communicates; it's just as much a function of what the hearer already knows. And so, two agents receiving the exact same testimony about why *p* may nevertheless differ greatly in their degree of understanding of why *p*. This difference simply reflects the epistemic asymmetry between the speakers that was already present to begin with. Testimony is a source of understanding but not always of very deep understanding.

Third, as I have argued, reductionism gives us a better fit with the linguistic data. On the one hand, it is supported by the observation that attributions of understanding seem to entail attributions of knowledge. On the other hand, it can offer a plausible story as to why the reverse does not hold.

In contrast, the non-reductionist incurs a very substantial task of giving a positive account of what “grasping” a proposition is. Without such a positive account, the imagery of “grasping” a proposition may well be suggestive. But its explanatory benefits are at best unclear.

8. Bibliography

Elgin, Catherine (2007). Understanding and the facts. *Philosophical Studies* 132 (1):33 - 42.

Elgin, Catherine Z. (2004). True enough. *Philosophical Issues* 14 (1):113–131.

Fricker, Elizabeth (2006). Second-hand knowledge. *Philosophy and Phenomenological Research* 73 (3):592–618.

Grimm, Stephen (forthcoming). Understanding as Knowledge of Causes. *Synthese*.

Grimm, Stephen (2012). “The Value of Understanding”. *Philosophy Compass* 7 (2):103-117.

Grimm, Stephen R. (2006). Is understanding a species of knowledge? *British Journal for the Philosophy of Science* 57 (3):515-535.

Hills, Alison (2009). Moral testimony and moral epistemology. *Ethics* 120 (1):94-127.

Khalifa, Kareem (2012). Inaugurating Understanding or Repackaging Explanation? *Philosophy of Science* 79 (1):15-37.

Khalifa, K. (2013). The Role of Explanation in Understanding. *British Journal for the Philosophy of Science* 64 (1):161-187.

Kvanvig, Jonathan L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.

Lipton, P. (2009). Understanding without explanation. In H. W. de Regt, S. Leonelli, and K. Eigner (eds.), *Scientific Understanding: Philosophical Perspectives*, pp.43-63. University of Pittsburgh Press, Pittsburgh.

Pritchard, Duncan (2008). Knowing the answer, understanding and epistemic value. *Grazer Philosophische Studien* 77 (1):325-339.

Pritchard, Duncan (2009). Knowledge, Understanding and Epistemic Value. *Royal Institute of Philosophy Supplement* 84 (64):19-.

Stalnaker, Robert (1984). *Inquiry*. Cambridge University Press.

Strevens, Michael (forthcoming). No understanding without explanation. *Studies in History and Philosophy of Science Part A*.

Treanor, Nick (2013). The Measure of Knowledge. *Noûs* 47 (3):577-601.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.

Zagzebski, Linda (2001). Recovering Understanding. In M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.