

FOLK ATTRIBUTIONS OF UNDERSTANDING: IS THERE A ROLE FOR EPISTEMIC LUCK?

DANIEL A. WILKENFELD, DILLON PLUNKETT AND TANIA LOMBROZO

dawilk@gmail.com

ABSTRACT

As a strategy for exploring the relationship between understanding and knowledge, we consider whether epistemic luck – which is typically thought to undermine knowledge – undermines understanding. Questions about the etiology of understanding have also been at the heart of recent theoretical debates within epistemology. Kvanvig (2003) put forward the argument that there could be lucky understanding and produced an example that he deemed persuasive. Grimm (2006) responded with a case that, he argued, demonstrated that there could not be lucky understanding. In this paper, we empirically examine how participants' patterns of understanding attributions line up with the predictions of Kvanvig and Grimm. We argue that the data challenge Kvanvig's position. People do not differentiate between knowing-why and understanding-why on the basis of proper etiology: attributions of knowledge and understanding involve comparable (and minimal) roles for epistemic luck. We thus posit that folk knowledge and understanding are etiologically symmetrical.

I. INTRODUCTION

Scientists, children, and adults often aim for more than merely having accurate, factual beliefs about the world – they aim for some kind of *understanding*. But what kind of epistemic achievement is understanding? Is understanding a special type of knowledge? Or is it something else? Here we begin to explore these questions empirically by investigating how people's attributions of understanding and knowledge respond to epistemic luck: cases in which beliefs are “accidentally” true, as when a stopped watch happens to tell the right time.

Precisely articulating the nature and boundaries of understanding, and in particular its connection to knowledge, is valuable for several reasons. First, understanding has at times been considered the primary target of inquiry (Zagzebski 2001), and even outside of academic philosophy, understanding is often invoked in discussions of education and development as a valuable end (e.g., Kember 1996). Second, if attributions of understanding track important features of our epistemic environment, then uncovering the basis for those attributions can help us “reverse engineer” our roles as epistemic agents, with implications for both philosophy and psychology (for an example, see Wilkenfeld *et al.* 2016). By investigating attributions of understanding and knowledge, we can potentially identify their unique epistemic roles.

As a strategy for exploring the relationship between understanding and knowledge, we focus on epistemic luck – that is, on the role of chance in the etiology of true beliefs. Since Gettier’s famous (1963) examples suggesting that epistemic luck can undermine knowledge, epistemic luck has been of great interest to epistemologists seeking an analysis of knowledge, and more recently to experimental philosophers investigating folk conceptions of knowledge (e.g., Starmans & Friedman 2012). Epistemic luck has also been at the heart of recent theoretical debates about the relationship between knowledge and understanding as part of a larger argument that understanding is not a species of knowledge. Specifically, Kvanvig (2003) argued that understanding differs from knowledge in its compatibility with luck, producing an example that he deemed persuasive. Grimm (2006) responded with a case that, he argued, demonstrated that there *could not* be lucky understanding¹ (except in one particular case – see footnote 11).

In this paper, we empirically examine whether people’s attributions of understanding line up with the predictions of Kvanvig or Grimm. While both philosophers are engaged in normative epistemology, not a descriptive characterization of “folk epistemology,” their arguments arguably presuppose the value of intuitions in establishing the connection between understanding and luck, with appeals to what “seems to be true” in hypothetical cases. Kvanvig (2003), for example, employs the method of “consideration of [i.e. intuitions regarding] particular cases” and discusses the results that emerge “when we think about knowledge” (2003: 197), which seems to invite appeal to how people actually consider cases and think about knowledge.² Grimm (2006) appeals to what “we’re tempted to say” (2006: 520) about cases and what the answer “seems to be” (2006: 521). Even if one maintains that the sort of philosophical reflection employed by Grimm and Kvanvig depends more on high-level theorizing than on untutored intuitions, characterizing untutored intuitions can help us understand whether the normative epistemic project involves descriptions of or revisions to our folk concepts.

The potential gap between philosophical and folk intuition is familiar from discussions of experimental philosophy and, we think, well addressed elsewhere (e.g., Alexander & Weinberg 2007). In the present case, however, an additional point is worth raising. One school of metaepistemology, owing to Craig (1999) and Weinberg (2006), suggests that the appropriate way to sharpen our epistemic concepts is to bring them into better alignment with our epistemic aims. However, we will argue in section 3 that our aims underdetermine whether understanding and knowledge should be responsive to etiology – there are different purposes one might have for each of the different ways we might

1 As will be discussed in the next section, Kvanvig mostly focused on “objectual understanding” (understanding an entire theory or entity – a principal example is understanding the Comanche dominance of the southern plains), whereas Grimm focused on “understanding why” (understanding why a particular event occurred – a principal example is understanding why a particular chestnut exploded). Our experiments follow suit in focusing on these sorts of understanding. Interestingly, neither Grimm nor Kvanvig really delves too deeply into questions of “understanding that.” We speculate that this is because understanding-that is pretheoretically assumed to behave somewhat like knowledge-that. We found some evidence for this hypothesis by comparing assessments of when someone “knows that” and of when someone “understands that” – using a variant of what we believe to be our study most sensitive to subtle differences – and finding no significant differences between assessments. See footnote 18.

2 Kvanvig also appeals to broader theoretical reasons for favoring his conclusions, but the appeal to intuitions about specific cases seems to stand as an independent and indispensable line of argument (see also footnote 5).

attribute or withhold understanding and knowledge. Since we cannot construct a theory *a priori* that predicts whether understanding and knowledge should be differentially sensitive to epistemic luck, it makes sense to examine *a posteriori* how they really are deployed, if only to understand the epistemic aims with which they ought to be aligned (see Lombrozo 2011 for a similar point regarding explanation). Even pursuing a more traditional conceptual analysis, any refinement of concepts must begin with the concepts people actually have, so there is at least that much value in exploring the folk notions.

The structure of the paper is as follows. In the next section, we briefly review Kvanvig's argument that knowledge and understanding can come apart, as well as Grimm's reply. In section 3, we explore the logical space of predictions one might make about people's responses to putative cases of knowledge and understanding that arise from deviant etiology generally (we will later explore the effects of different *ways* etiology can be deviant). Rather than cast aside some as pretheoretically unworkable, we will argue that, in fact, any of the four patterns we identify would be a sensible way of achieving some particular epistemic aim with our attributions of 'knowledge' and 'understanding'. It is thus of great interest to determine which of the four possibilities best describes people's judgments. This brings us to the data, presented in section 4. We argue that there is one relatively clear winner: that people are only weakly sensitive to etiology in attributing either knowledge or understanding,³ with no evidence that people are more willing to attribute understanding than knowledge in cases involving epistemic luck.

2. THE PHILOSOPHICAL DEBATE ABOUT ETIOLOGY OF UNDERSTANDING

2.1 *Kvanvig and the Comanche*

In *The Value of Knowledge and the Pursuit of Understanding*, Jonathan Kvanvig (2003) argues against the idea that understanding is a species of knowledge as follows:

- 1) Knowing a body of information "involves knowing a number of truths that make up that body of information." (p. 197)
- 2) Knowing truths requires "nonaccidental connections between mind and world." (p. 197)
- 3) Understanding does not require nonaccidental connections between mind and world.
- 4) → Understanding of a body of information does not require, and is thus not a species of, knowledge of a body of information.

Premise 1, while undefended, seems plausible enough. Premise 2 is a familiar consequence of the literature on Gettier cases (e.g., we would not consider the person with the broken watch as someone who *knows* what time it is, even when she is correct), though the general point can be spelled out in different ways that we will have reason to revisit in light of the data.

Kvanvig supports premise 3 with two arguments. The first stems from the theory of understanding propounded throughout Kvanvig's book. Kvanvig argues that while understanding is factive, what differentiates sets of factive beliefs that do and do not constitute

3 More specifically, they are sensitive to what Starmans and Friedman (2012) call "apparent evidence."

understanding is the additional requirement that the beliefs must possess the right *internal* structure. There is no theoretical motivation for an additional requirement in the form of a non-accidental connection between mind and world, which is an external consideration. A defense of Kvanvig's account of understanding is beyond the scope of this paper; for our purposes, the key point is the compatibility between understanding and epistemic luck.⁴

As there is no settled account of understanding, Kvanvig proposes to look at some data regarding which cases we would or would not classify as understanding.⁵ Kvanvig gives the example of someone's putative understanding of the Comanche dominance of the southern plains of North America (Kvanvig 2003). We know the person understands, because she is able to answer any questions we have about the Comanche, and by stipulation is not just guessing but retrieving the information from memory.⁶ Yet Kvanvig points out that the subject's true and sincerely held beliefs, while sufficient for understanding, might not rise to the level of knowledge if they could just as easily have been false. "For example, most history books might have been mistaken, with only the correct ones being the sources of understanding in question and with no basis in the subject for preferring the sources consulted over those ignored" (Kvanvig 2003).

Similarly, Kvanvig argues that the possibility of lucky understanding explains a case that had previously been a vexing puzzle. Richard Foley (1996) argues that a Swampman, assembled from stray particles by a lightning strike such that he possessed extremely comprehensive true beliefs, would have a great deal of knowledge. Foley argues that the possibility of such a creature brings into question the corpus of post-Gettier views that place etiology at the center of knowledge. Kvanvig, by contrast, agrees that there is something "epistemologically extraordinary" (2003: 199) about Foley's Swampman, but argues that this intuition is best explained by the Swampman's possessing understanding but no knowledge.

How can one determine whether there really is understanding in these cases? With an account of understanding already in hand, it would be appropriate at this juncture to

-
- 4 However, it should be noted that his point can be generalized beyond his own account. On the majority of accounts of understanding, there is no etiology condition. For example, on the influential account of De Regt and Dieks (2005), one understands a phenomenon when one can relate it to an intelligible theory, where the intelligibility of a theory is measured by what one can do with it (rather than any particular causal pedigree). On Morgan and Morrisson's (1999) account, one understands when one has a good model of the understood, but something can be a good model as a result of happenstance rather than any causal link to that which is modeled. (In fact, frequently the point of a model is to engage in surrogative reasoning about something that does not even exist, like a to-be-constructed building.) And on Wilkenfeld's (2013) representationalist account, one understands when one has a useful representation, but the usefulness of a representation does not depend on its etiology. In fact, to our knowledge, the only accounts of understanding that require that the content of understanding be acquired in a non-lucky way are those that presuppose that understanding is a species of knowledge. If the question under investigation is *whether* understanding is best thought of as a species of knowledge, it would seem that independent theoretical inquiry supports Kvanvig's contention that it is not.
 - 5 As we read Kvanvig, this is supposed to be an independent appeal to intuition, rather than a continuation of the theoretical argument. His exact words are "*Moreover*, consideration of particular cases of understanding suggests the same [as the theoretical point]" (Kvanvig 2003, emphasis added). This seems to be introducing a new line of argument, unmoored from theoretical considerations. It is this second line of argument that we take ourselves to be addressing.
 - 6 In fact, Kvanvig at one point suggests that "such an ability is surely constitutive of understanding" (2003: 198) but his point does not depend on anything so strong.

consult it. However, in the absence of such a theory or, as in this case, in the process of testing such a theory against the “data” of cases, there is little to do but consult intuitions.⁷

2.2 *Grimm’s Exploding Chestnuts*

Responding to Kvanvig, Grimm (2006) constructs a case that he argues is structurally analogous to the Comanche case. Nevertheless, he contends that it fails to manifest understanding without knowledge. Here is his discussion of such a case:

A further sort of problem turns on the relation that someone might bear toward his or her environment and parallels the problem Alvin Goldman ([1976]) made famous with his “fake barn” cases. Suppose, for example, that while wandering through a blacksmith’s shop Becky notices a chestnut lying on top of an anvil, and she pauses to watch as the blacksmith moves to strike the chestnut with his hammer. At the very instant that the hammer touches the chestnut, the chestnut explodes into fragments. Becky then concludes, naturally enough, that the chestnut shattered because of the blow from the hammer.

Here again, however, it is possible to introduce a Gettier twist. Thus we can suppose, first, that as a rule the blacksmith heats the anvil to extremely high temperatures, so that after a certain amount of time the chestnuts placed on the anvil eventually explode from simple heat stress. Moreover, we can imagine that the blacksmith enjoys testing his timing so that his hammer grazes the top of the chestnut at the precise moment that the chestnut is due to explode. In this case, however, as Becky is walking by, things don’t unfold normally. Either because he was in the mood for a change or simply because he forgot to heat up the anvil, as the hammer makes contact with the chestnut it is not on the verge of exploding from heat stress; so it is, in fact, the force of the hammer blow that shatters the chestnut. What then should we make of Becky’s conclusion that the chestnut shattered because of the blow from the blacksmith’s hammer? Does she genuinely understand why the chestnut shattered? Even though she has identified the genuine cause of the shattering in this case (the hammer blow), once again, the answer seems to be No. (Grimm 2006: 521)

Grimm thus contends that intuitions dictate the following: just as we would be loathe to attribute *knowledge* to Becky, we would be equally reticent to attribute *understanding*. We will argue that neither Kvanvig’s nor Grimm’s prediction is wholly accurate (at least as applied to *folk* intuitions): Kvanvig is right that attributions of understanding are relatively insensitive to luck (more specifically, to what we and others call “environmental luck”), but Grimm is right that attributions of understanding are *no less sensitive* to luck than attributions of knowledge.

2.3 *Disambiguating Epistemic Luck*

For clarity, we should set out what we mean by “epistemic luck.” As has been well documented (Blouw *et al.* Forthcoming), there are many different ways beliefs could be lucky while still being justified and true. Our main concern will be with barn-façade type cases (Goldman 1976) in which the fact that the state-of-affairs is as believed is really causally responsible for some true belief, but where one could easily have formed a false belief on

7 To be clear, “intuition” is here being used in as non-theoretically laden a manner as possible. Just one possibility is that what we are really probing is best understood as tacit (or so-called “Socratic”) knowledge (Cappelen 2014).

what would have appeared to be the same evidence; following Pritchard (2009), we refer to these as cases of “environmental luck.”⁸ This seems to be the structure of Kvanvig’s example, and Grimm explicitly states that there is neither knowledge nor understanding in barn-façade cases (2006: 521). However, our interpretation of both Grimm and Kvanvig suggests that they should make similar predictions in cases of even more suspect etiology, as when a belief is not even caused by its truth-maker at all (a kind of case, following Starman and Friedman 2012, that we refer to as “Apparent-evidence”).⁹ In Kvanvig’s case, we think this is suggested by the discussion of Foley’s Swampman, where the causal link to the truth is broken and he predicts understanding but no knowledge. For Grimm, this is simply a logical extension of his view: if understanding and knowledge are both absent when there is weaker luck, then they should both be absent when there is more radical luck.

Pritchard (2009: 16) differentiates between environmental and Gettier luck, arguing that while neither is compatible with knowledge, the former but not the latter is compatible with understanding.¹⁰ Given that we find no distinction between knowledge and understanding across cases that range from environmental luck to Foley’s Swampman, we take our results to be a *prima facie* challenge to Pritchard as well as Kvanvig.

3. POSSIBLE PATTERNS IN ATTRIBUTIONS OF KNOWLEDGE AND UNDERSTANDING

We’ve just considered Kvanvig’s and Grimm’s positions on whether epistemic luck undermines understanding; however, their positions are but two of four possibilities concerning the relationship between knowledge and understanding in cases of epistemic luck: people could use epistemic luck as a basis for withholding knowledge attributions, for withholding understanding attributions, for both, or for neither (see Table 1). We argue that, pre-theoretically, any of the four combinations is possible, and each way of marking out our concepts would serve a legitimate epistemic function.

The first possibility (P1) is that people withhold attributions of understanding *and* attributions of knowledge when the belief in question is true only on the basis of luck. This is the prediction of Grimm (2006), who argues that understanding seems absent in his exploding chestnut case in just the same way that knowledge seems absent in fake-barn cases.¹¹ Smith (2014) similarly seems to suggest that understanding and knowledge are both sensitive to etiology in this way, and Khalifa (In press, Chapter 9) suggests that

8 In Blouw *et al.*’s (Forthcoming) taxonomy, this would be a case of “Detection with Unsuccessful Threat.”

9 Blouw *et al.* (Forthcoming) would describe these as cases of “Failure to Detect.”

10 We thank an anonymous reviewer for pointing out the nuances of Pritchard’s view.

11 This is something of an oversimplification of Grimm’s view, as there are cases for which he would predict people would say that there is both knowledge and understanding. Specifically, he cites Hawthorne’s (2004) example of a circumstance where children are each given a different book, each of which lists a different capital of Austria. Hawthorne and Grimm would both say of the child whose book says “Vienna” that he knows the capital of Austria, and Grimm thinks there are similar cases for understanding. However, the template for the main sort of case we are considering is taken from Grimm’s discussion of barn-façade examples, of which he says there should be neither knowledge nor understanding.

Table 1 Summary of four possible positions concerning the compatibility of knowledge and understanding with the presence of epistemic luck.

Possible positions:	Is knowledge compatible with luck?	Is understanding compatible with luck?
P ₁ (Grimm's position)	No	No
P ₂ (Kvanvig's position)	Yes	No
P ₃	No	Yes
P ₄	Yes	Yes

deviant etiology is at least deleterious to understanding. This pattern is eminently sensible if it is generally desirable only to attribute epistemic virtuosity to individuals whose beliefs are sensitive to the truth – that is, for whom it is the case that if their (current) beliefs had been false, they would likely not have possessed them. This might be an important feature if we want to pick out not just who happens to be correct on this occasion, but who would likely be correct in closely related circumstances. The fact that Becky the chestnut-viewer and (the person we'll name) Rebecca the Comanche-reader could so easily have been mistaken might properly give us pause before we issue them the entitlements that come with an epistemic state as venerable as knowledge *or* understanding.

The second possibility (P₂) is that people would attribute understanding to agents whose beliefs are only accidentally true, but would not attribute knowledge. This is Kvanvig's picture. It would be sensible if knowledge were thought of as primarily an externalist success – that is, a state that demarcates who is appropriately epistemically situated given his or her actual position in the world. By contrast, we might want to use understanding to demarcate a more internalist success – that is, whether the agent's beliefs are sensible/rational/coherent by his or her own lights (though, as Kvanvig suggests, we might still hold them to the truth). If we want to use understanding attributions as a form of praise, it would make sense that we would want to abstract away from details concerning whether a belief happened to be formed by the proper etiology, as such facts are potentially inaccessible to the agent.¹²

The third possibility (P₃) is that people would attribute knowledge even to people whose beliefs are only accidentally true, but still withhold understanding from such people. As far as we know, no one actually advocates this position, but there would be a certain amount of sense to it. To the extent that understanding is about trust and deference – a connection espoused by Rosenberg (1981) and defended empirically in Wilkenfeld *et al.* (2016) – one might expect understanders to have a particularly robust relationship to the truth. If I am going to defer to you in a general area, it might not suffice that you happen to be correct now, in a particular case, but rather that you are correct for the right reasons. In contrast, if we merely want knowers to get the answer right, we might not care how the right answer came about. Surprisingly, given P₃'s lack of vocal

12 While the same logic might suggest we do away with a factivity requirement, that requirement might be forced on us as a consequence of intuitions about the extension of “understanding.” Kvanvig argues that understanding does in fact have to be factive, but we follow Kvanvig in thinking that understanding could be factive but compatible with epistemic luck.

advocates, we will see there is actually some evidence that people do attribute objectual knowledge, but not propositional knowledge, in accordance with this prediction (see Experiment 5 below).

The fourth possibility (P₄), which again has no explicit advocates, is that people might be insensitive to suspect etiology when attributing either knowledge or understanding. This would be sensible if what we really cared about is someone having the right answer – or, perhaps, having the right answer with appropriate internalist justification. This view is suggested in the knowledge literature by Starmans and Friedman (2012), who demonstrate that people are wholly insensitive to factors of *environmental* luck when attributing knowledge, though they are sensitive to luck that explicitly disrupts the particular causal chain by which the agent came to have the belief in question. So P₄ might arise as a natural prediction were one to endorse Starmans and Friedman’s stance regarding knowledge while maintaining that Kvanvig has identified an important truth about the insensitivity to etiology of understanding.

4. NATURAL LANGUAGE FOLK ATTRIBUTIONS OF KNOWLEDGE AND UNDERSTANDING

4.1 *Experiment 1*

In Experiment 1, we investigate whether folk attributions of knowledge and understanding are sensitive to epistemic luck. We do so by adapting a case used by Grimm (2006).

Participants read one of four vignettes about a woman who observes (or thinks she observes) a man destroying a chestnut with a hammer. In two baseline versions, the woman forms a belief about the destruction of the chestnut that is either simply true or simply false. The other two versions corresponded to variants of epistemic luck cases (in that they put forward true belief formation with suspect etiology). In one version (“*lucky environment*”), the belief is formed as a result of an unbroken causal chain that, under slightly different initial conditions, would have led to a mistaken belief. In the other version (“*veridical hallucination*”), the belief is formed completely by chance, where the truth of the proposition played no causal role in the woman coming to believe it (rather it was the result of the chance concordance between a hallucination and reality – see below). Participants were then asked either whether the woman *knew why* the chestnut exploded or *understood why* the chestnut exploded.

Experiment 1 tests the hypothesis that everyday attributions of knowledge and understanding are differentially responsive to the sorts of environmental epistemic luck that the barn-façade literature suggests should undermine knowledge attributions. In addition, this experiment lets us examine the robustness of the barn-façade and apparent-evidence intuitions generally, and whether they vary based on different ways epistemic luck could have affected the formation of the belief.

4.1.1 Method

Participants. One hundred and fifty-four adults (62 female, mean age 34) were recruited through the Amazon Mechanical Turk marketplace (MTurk) and participated in

exchange for monetary compensation.¹³ An additional 132 participants were excluded prior to analysis for incorrectly answering one or more reading comprehension questions, with one additional participant excluded for failing to answer all the questions.¹⁴

Materials and Procedure. At the start of the experiment, participants were randomly assigned to read one of four vignettes (see Table 2), all variations on the case of Becky in the blacksmith's shop from Grimm (2006). In each vignette, Becky sees John, a blacksmith, strike a chestnut with a hammer, and sees the chestnut explode. Unbeknownst to Becky, the anvil is hot enough to cause chestnuts to explode and John is playing a game of timing his blows based on when the heat would cause each chestnut on his anvil to explode. In all vignettes, Becky forms the belief that the chestnut exploded because it was hit by the hammer.

In the *normal belief* condition, Becky's belief is true. John usually strikes the chestnut just before it would explode from the heat, and he does so as she is watching. In the *false belief* condition, John times his blows so that he strikes *after* each chestnut explodes due to the heat, and he does this as she is watching. Accordingly, Becky's belief that the chestnut exploded due to the hammer is false. The *lucky environment*¹⁵ condition follows Grimm's blacksmith example most closely: John usually strikes each chestnut after it explodes due to the heat. But, on the one occasion Becky is watching, his timing is off and his hammer blow does cause the chestnut to explode. Thus, Becky's belief is true, but lucky in that she would have formed a false belief based on identical evidence if she had walked in a moment earlier or later. Finally, in the *veridical hallucination* condition, John usually strikes the chestnut before it explodes due to the heat, and does so as Becky is watching. In this scenario though, Becky accidentally ingests a hallucinogenic plant before entering John's shop. Becky's hallucinations just happen to correspond exactly to what she would have seen if she had not been hallucinating. As in the *lucky environment* condition, Becky's belief is true only by chance. In this case, however, her belief was not causally connected to its object.¹⁶

13 In Experiments 1–3, participants were given \$0.30 in exchange for an estimated five minutes of work. In all experiments, participation was restricted to users with an IP address within the United States and an approval rating of at least 95% based on at least 50 previous tasks.

14 Though these exclusion rates are high, the subtlety of the questions needed to make sure participants were tracking the small changes between vignettes make such rates not wholly unexpected. More importantly, in all of Experiments 1–4, including in analyses those participants who were excluded for incorrectly answering at least one reading comprehension question does not change the qualitative results. Including all participants in Experiment 5 does cause some results to fall below the threshold of significance, but still does not have any impact on the general pattern.

15 There is some potential for terminological confusion here. At one point (2006: 520) Grimm discusses “bad environments” as akin to “veridical hallucinations” in being cases of non-knowledge. However, he later introduces the notion of a “bad information environment,” as opposed to a “bad information source,” where he claims the former is compatible with both knowledge and luck (see footnote 11). Since our case is explicitly adapted from the same location as the first distinction, that is the sort of case on which we focus. We do not invoke the second sense of “bad environment” anywhere in this paper. In order to differentiate more sharply from Grimm's notion of “bad information environment,” we refer to the relevant scenario in our vignettes as “lucky environment”.

16 Kareem Khalifa has pointed out that we are actually changing at least two variables between our *lucky environment* and *veridical hallucination* conditions, as in the latter not only is the causal connection broken but there is something flawed about the mechanism by which the agent forms the belief. However, given that even with two changes we found no asymmetry between knowledge and

Table 2 Overview of the narrative in each condition used in Experiment 1.

Case	What usually happens	What happens this time	What appears to happen	What Becky believes happened	The Basis of Becky's Belief
Normal Belief	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	Accurate perception
Lucky Environment	Chestnut explodes from heat	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	(Accurate, but lucky) perception
Veridical Hallucination	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	(Accurate, but lucky) hallucination
False Belief	Chestnut explodes from heat	Chestnut explodes from heat	Chestnut explodes from hammer blow	Chestnut explodes from hammer blow	(Inaccurate) perception

For example, in the lucky environment condition, participants read the following vignette:

If chestnuts are placed on a hot anvil, they explode after exactly one minute.

As a game, John the blacksmith likes to put chestnuts on a hot anvil one at a time, swing a hammer at them, and time his blows so that he always connects *a moment after* each chestnut explodes from the heat. He is very good at this game, and almost always gets the timing exactly right. As a result, chestnuts on John's anvil usually explode because of the heat of the anvil, not due to the hammer.

Becky walks in to John's shop just after he's placed a chestnut on the anvil. Moments later, she sees John swing the hammer. This one time, John's swing is a few milliseconds too early, and the chestnut explodes because it is hit by the hammer, with the heat playing no causal role.

Becky, on the basis of her observation, forms the belief that the chestnut on John's anvil exploded because it was hit by the hammer. This belief is true; the chestnut that she saw explode did so because it was hit by the hammer.

After reading one of these four vignettes, participants were randomly assigned to judge either Becky's *knowledge* or her *understanding* by reporting their agreement with one of the following two questions on a seven-item Likert scale ranging from "Strongly Disagree" to "Strongly Agree."

Becky knows why the chestnut exploded.

Becky understands why the chestnut exploded.

understanding, we doubt that other perturbations would have found any either. See the General Discussion for more on this.

Participants were then asked eight true/false reading comprehension questions to ensure that they had understood the vignette. Examples included, “The chestnut on John’s anvil exploded from the heat of the anvil,” and “Becky thinks the chestnut on John’s anvil exploded because it was hit by a hammer.” Participants who made one or more errors on the comprehension questions were excluded from further analysis (though see footnote 14). Finally, participants answered demographic questions.

4.1.2 Results

Responses were analysed with an ANOVA with etiology condition (4: *normal belief*, *lucky environment*, *veridical hallucination*, *false belief*) and attribution (2: *knowledge*, *understanding*) as between-subjects factors. This revealed a main effect of etiology condition, $F(3, 146) = 130.07$, $p < 0.001$, $\eta_p^2 = 0.728$, and no significant main effect of attribution, $F(1, 146) = 0.71$, $p = 0.40$, nor interaction between etiology condition and attribution (see Figure 1). In the *normal belief* condition, participants attributed high levels of knowledge and understanding to Becky. The *lucky environment* had no negative impact on either attribution, with Tukey post-hoc tests revealing ratings that were not statistically different from the *normal belief* condition, $ps > 0.99$. However, in the *veridical hallucination* condition, Tukey post-hoc tests revealed that participants attributed significantly less knowledge and understanding to Becky than in the *normal belief* condition, $ps < 0.001$, or in the *lucky environment* condition, $ps < 0.001$. Finally, in the *false belief* condition, Tukey post-hoc tests revealed that participants attributed minimal levels of knowledge and understanding to Becky, significantly below even the *veridical hallucination* condition, $ps < 0.001$. There were no significant differences in participants’ ratings of

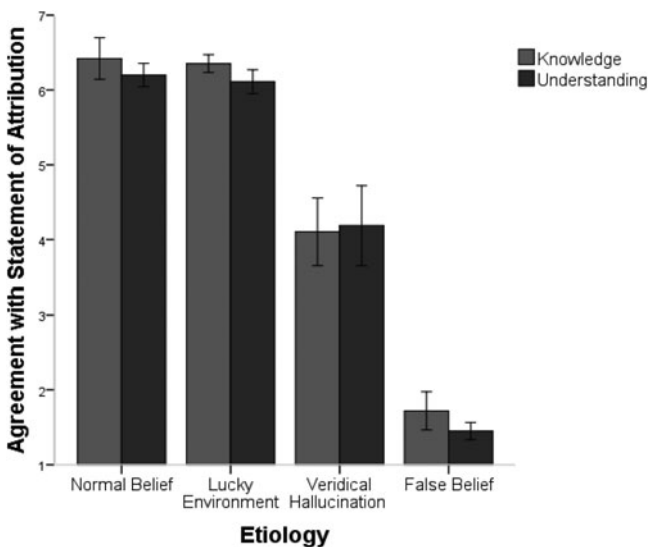


Fig. 1. Experiment 1 ratings of how much participants agreed with a sentence asserting that Becky knew/understood why the chestnut exploded, presented as a function of etiological condition. Ratings were made on a 7-point scale from completely disagree (1) to completely agree (7). Error bars are 1 SEM in each direction.

Becky's knowledge and of her understanding in any of the four conditions ($ps > 0.23$). Ratings in the veridical hallucination condition did not differ significantly from the scale midpoint, all $ps > 0.67$.

4.1.3 Discussion

Experiment 1 found that participants were completely insensitive to epistemic luck in the typical barn-façade/lucky environment case, and, though they were somewhat sensitive to it in the veridical hallucination case, they were not at all *differentially* sensitive depending on whether they were asked to attribute knowledge or understanding. That is, participants' assessments of understanding and of knowledge rose and fell in lockstep.

The finding that folk attributions of knowledge are relatively insensitive to epistemic luck is not new; in fact, our findings are consistent with prior work in experimental epistemology (e.g., Starmans and Friedman 2012).¹⁷ However, this is the first experiment (to our knowledge) to investigate the effect of epistemic luck on attributions of understanding, and also the first to compare attributions of knowledge with those of understanding.

Our pattern of results is consistent with P1 and P4, but challenges P2 and P3. In particular, the results in the *veridical hallucination* case are best accounted for by P1, and the results in the *lucky environment* case are best accounted for by P4. The lack of any differential sensitivity to knowledge and understanding challenges both P2 and P3. In the general discussion, we return to the implications of all of our findings for all four possible patterns.

4.2 Experiment 2

Experiment 1 failed to find evidence that judgments concerning knowledge and understanding diverge when it comes to cases of epistemic luck. However, it remains a possibility that our test was simply insufficiently sensitive to some difference that does exist. In particular, it could be that people do differentiate between understanding and knowledge when faced with the possibility of both attributions simultaneously (as readers of Grimm and Kvanvig's original pieces were), but that they gravitate to the same fixed point when judgments are made in isolation. To test this hypothesis, we reran Experiment 1 using a within-subjects design. We reasoned that this design would help participants home in on any subtle differences between knowledge and understanding. We did have the concern that by asking both questions one after another, we might create a pragmatic implicature that participants were *expected* to provide different ratings, thus risking a misleading result; however, this concern proved unfounded.

17 However, Nagel *et al.* (2013) reach different results. Weinberg *et al.* (2001) famously find cultural variability, and Colaço *et al.* (2014) have results that mostly accord with Starmans and Friedman (2012), but reveal a surprising variability based on participant age. Wright (2010), Cullen (2010) and Turri (2013) raise some general methodological worries regarding how precisely one should go about constructing such experiments. In general then, people's willingness to ascribe knowledge in cases of etiological luck should be viewed as neither shocking nor obvious.

4.2.1 Method

Participants. Forty-six adults (21 female, mean age 31) were recruited through MTurk. An additional 38 were excluded for answering at least one reading comprehension question incorrectly, with two more excluded for failing to answer every question.

Materials and Procedure. Experiment 2 followed the same procedure as Experiment 1, with two changes. First, we included only the two epistemic luck cases, eliminating the baseline conditions. Thus all participants were randomly assigned to either the *lucky environment* condition or the *veridical hallucination* condition. Second, whereas we had varied the attributed state *between* participants in Experiment 1, in Experiment 2 each participant was asked *both* about Becky's knowledge-why and about her understanding-why (in random order). The text of the vignettes and the two attribution questions were unchanged from Experiment 1, as were the reading comprehension and demographic questions.

4.2.2 Results

Responses to the two attribution questions were analysed with a mixed ANOVA with etiology condition (2: *lucky environment*, *veridical hallucination*) as a between-subjects factor and attribution (2: *knowledge*, *understanding*) as a within-subjects factor. This revealed a significant main effect of etiology condition, $F(1, 44) = 16.34$, $p < 0.001$, $\eta_p^2 = 0.271$, and no significant main effect of attribution, $F(1, 44) = 0.0043$, $p = 0.95$, nor an interaction between etiology condition and attribution ($p = 0.798$) (see Figure 2). Consistent with Experiment 1, in the *lucky environment condition*, participants attributed both knowledge and understanding to Becky at high levels, and in the *veridical hallucination* case, participants provided equally low attributions of knowledge and

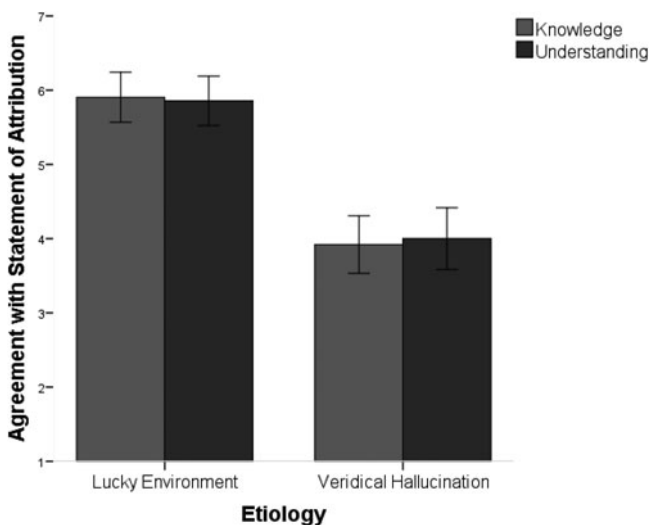


Fig. 2. Experiment 2 ratings of how much participants agreed with a sentence asserting that Becky knew/understood why the chestnut exploded, presented as a function of etiological condition. Ratings were made on a 7-point scale from completely disagree (1) to completely agree (7). Error bars are 1 SE in each direction.

understanding to Becky, with the average response for both attributions very close to the scale midpoint.

4.2.3 Discussion

The results from Experiment 2 support the same conclusion as those from Experiment 1: participants were in general insensitive to environmental luck, moderately sensitive to the sort of luck in *veridical hallucination*, and were no more sensitive to luck when attributing knowledge-why than when attributing understanding-why.¹⁸ This lack of difference was observed despite potential task demands suggesting that different answers were expected for the knowledge and understanding attributions. Of course, it's possible that participants instead felt a demand to be consistent; if so, however, this supports our conclusion insofar as it suggests that participants took the understanding and knowledge questions to be asking about the same thing. Moreover, the absence of order effects suggests that if an initial response somehow constrained the response that followed, the effect was symmetrical, further supporting a correspondence between knowledge and understanding.

In sum, the findings from Experiments 1 and 2 suggest that participants do not treat understanding and knowledge judgments differently with respect to their compatibility with suspect etiology, whether attributions are made in isolation or in conjunction. As with Experiment 1, the results of Experiment 2 provide some support for P1 and some support for P4, but challenge P2 and P3.

4.3 Experiment 3

Experiments 1 and 2 failed to find evidence that people distinguish between understanding and knowledge on the basis of etiology. However, the stimulus materials in both experiments were adapted from cases proposed by Grimm (2006), who would contend that there are no such differences. Perhaps his examples just happen to be cases for which knowledge and understanding coincide, but in reality other cases would reveal important differences. To examine this possibility we tested participants' intuitions regarding a case Kvanvig calls "Foley's Swampman." Kvanvig claims that this case illustrates a circumstance in which one can possess understanding but not knowledge, and that this is because the beliefs in question do not have the right etiology (2003: 199). An additional reason to test Foley's Swampman is that, as opposed to the recent empirical work on both apparent evidence and barn-façade cases, there has been (as far as we know) no empirical work investigating Foley's Swampman. It was thus not only of interest how people *differentially* rated knowledge and understanding in this case, but how people rated knowledge at all.

18 As suggested by a reviewer, we also ran a variant comparing knowledge-that with understanding-that. The setup and dependent variables were identical, save for the word "that" being substituted for "why" after "understands" and "knows." This resulted in four key ratings: *lucky environment*/knows ($M = 6.62$), *lucky environment*/understands ($M = 6.50$), *veridical hallucination*/knows ($M = 4.22$), and *veridical hallucinations*/understands ($M = 4.33$). A mixed ANOVA with etiology condition (2: *lucky environment*, *veridical hallucination*) as a between-subjects factor and attribution (2: *knowledge*, *understanding*) as a within subjects factor revealed the predicted main effect of etiology, $p < 0.001$, but no main effect of attribution or interactions ($ps > 0.004$). This suggests that knowledge- and understanding-that are no more likely to exhibit differences than knowledge- and understanding-why (see also footnote 1).

4.3.1 Method

Participants. Eighty subjects (35 female, mean age 33) were recruited through MTurk. An additional 24 participants were excluded for answering at least one reading comprehension question incorrectly.

Materials and Procedure. At the start of the experiment, participants were randomly assigned to read one of two vignettes. In the *normal* vignette, participants were presented with a story of someone (“Richard”) born in the usual way, who learned in the usual way, and formed ordinary true beliefs about various scientific topics, including the observed perturbations in Mercury’s orbit. In the matched *Swampman* vignette, participants read about a version of Foley’s Swampman. The *Swampman* vignette was intended to portray another way in which one’s beliefs could form from suspect etiology, and it read as follows:

A lightning strike randomly rearranges some particles in a swamp, by chance forming a person. Let’s call this person “Richard.” Richard happens to be formed in such a way that he has an array of beliefs¹⁹ about almost everything, and can correctly answer any question you put to him about most scientific topics. For example, if asked, he will tell you a lot of true things about the orbit of Mercury. He can explain how small perturbations in Mercury’s observed orbit are a necessary consequence of General Relativity. Everything he says about Mercury and General Relativity is true, and he sees all the connections among the things he says. However, he has never actually seen Mercury or any of the other experimental evidence for General Relativity. It’s interesting that Richard ended up with precisely these beliefs: they resulted from the way lightning happened to strike the swamp at a particular time, purely by chance. Had the lightning struck slightly differently, he would have had false beliefs or no beliefs at all.

After reading one of these two vignettes, participants were randomly assigned to judge either Richard’s *knowledge* or his *understanding* by reporting their agreement with one of the following two questions on a seven-item Likert scale ranging from “Strongly Disagree” to “Strongly Agree.”

Richard knows why Mercury has the observed orbit it has.
Richard understands why Mercury has the observed orbit it has.

Participants were asked four reading comprehension questions to ensure that they had understood the vignettes, including “Richard was taught about Mercury’s observed orbit by experts,” and “Richard has never seen Mercury.” Finally, participants were asked demographic questions.

4.3.2 Results

Responses were analysed with an ANOVA with etiology (2: *normal*, *Swampman*) and attribution (2: *knowledge*, *understanding*) as between-subjects factors. This revealed a main

19 It is less than ideal to refer to Swampman as having beliefs at all, since this begs important questions against accounts that individuate belief content by (past) causal role. However, other locutions philosophers sometimes employ (“expressions of thoughts”?) were judged too cumbersome, and unlikely to produce relevantly different results.

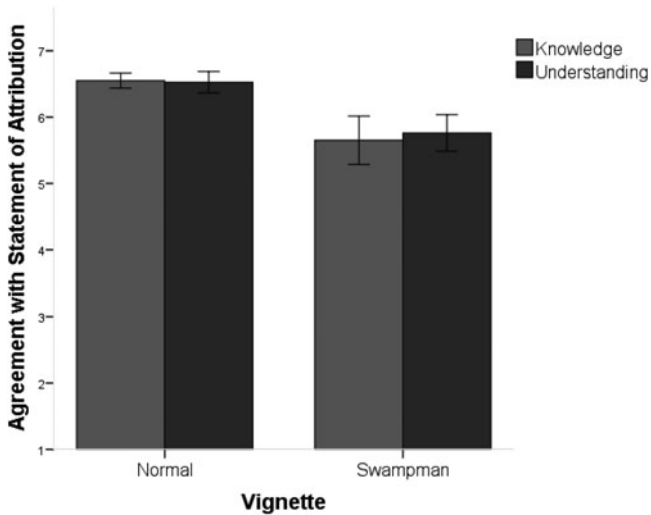


Fig. 3. Experiment 3 ratings of how much participants agreed with a sentence asserting that Richard knew/understood why Mercury has the observed orbit it does, presented as a function of etiological condition. Ratings were made on a 7-point scale from completely disagree (1) to completely agree (7). Error bars are 1 SEM in each direction.

effect of etiology, $F(1, 76) = 10.95$, $p = 0.001$, $\eta_p^2 = 0.126$, and no significant main effect of attribution, $F(1, 76) = 0.033$, $p = 0.86$, nor an interaction between etiology and attribution, ($p = 0.788$) (see Figure 3). In the normal condition, participants attributed high levels of knowledge and understanding to Richard, with a slight (but significant) decrease in their willingness to attribute either in the Swampman condition, $t(54.12) = 3.41$, $p = 0.001$.

4.3.3 Discussion

Participants were only minimally sensitive to whether someone's beliefs were the result of traditional learning versus random lightning strikes, and not at all *differentially* sensitive when it came to attributing knowledge versus understanding. Folk intuitions are thus consistent with Foley's original assertion that Swampman possesses knowledge, and consistent with Kvanvig's assertion that Swampman possesses understanding. However, they are not consistent with Kvanvig's additional claim that in this particular case, Swampman possesses understanding *but no knowledge*. Consistent with both of the previous experiments, these results provide support for either P1 or P4 (depending on how much weight one puts on the small, but significant, effect of the apparent-evidence-type deviant etiology), but continue to challenge P2 and P3.

4.4 Experiment 4

Experiments 1–3, which included cases from both Grimm and Kvanvig, consistently revealed no differentiation between understanding attributions and knowledge attributions. However, all of our experiments probed knowledge and understanding of particular propositions, rather than of whole objects. Kvanvig's original example, however, is a bit

more general, as it compares understanding an entire history (that of Comanche dominance of the southern plains) with knowing that history.²⁰ It is possible that only so-called “objectual” knowledge and understanding are differentially sensitive to etiology, and that this explains why Experiments 1–3 did not yield any interaction between etiology and attribution type.

In Experiment 4, we thus turned to Kvanvig’s main example, which is someone’s putative knowledge/understanding of Comanche dominance of the southern plains in the 18th century. This knowledge/understanding could be framed in either propositional terms – X knew/understood why the Comanche dominated the plains – or objectual terms – X knew/understood the history of Comanche dominance of the southern plains.²¹ Testing both knowledge and understanding of both objects and propositions also has the potential to provide indirect evidence for a proposal by Brogaard (2005) to the effect that judgments about knowledge and understanding may (falsely) appear to diverge because people tacitly compare objectual understanding with knowledge-why. For example, one might think that one’s *understanding* of quantum theory surpasses one’s *knowledge why* particular particles behave as they do, and be thus tempted to say that understanding is more valuable than knowledge. Brogaard, however, argues that a more appropriate comparison would be between *understanding quantum theory* and *knowing quantum theory* (Brogaard 2005: 4–5).

4.4.1 Method

Participants. Eighty-two participants²² (31 female, mean age 34) were recruited through MTurk.²³ An additional 24 participants were excluded prior to analysis for incorrectly answering at least one reading comprehension question, with one more excluded for not answering every question.

Materials and Procedure. At the start of the experiment, participants were randomly assigned to read one of two vignettes. In the *normal* vignette, participants were presented with an ordinary story of someone (“Rebecca”) coming to learn about the Comanche dominance of the southern plains in the 18th century from one among many reliably researched and written history books. In the *lucky library* vignette, Rebecca happened to pick the one reliable history book from amid an array of unreliable books; this is an example of environmental luck. To make the possibility of error particularly salient,

20 Notice that in this case, “knowing that” is not propositional, since the object is a whole history rather than a proposition.

21 We chose to test knowledge and understanding of “the history of Comanche dominance” rather than just “Comanche dominance” because we worried that saying X “knows Comanche dominance” might carry an implicature that X has experienced it firsthand, which is a potential difference independent of anything suggested by Kvanvig. Kareem Khalifa suggested that adding the word “history” might not alleviate this problem, as while “history” can refer to a narrative of past events it can also refer to the past events themselves, which participants might have then thought of as being known by acquaintance. We take it to be a step in the right direction, however, while minimally deforming Kvanvig’s case.

22 The number of participants per condition was roughly 20 in Experiments 1–3, but roughly 40 in Experiments 4 and 5. This change was incidental; Experiments 4 and 5 were originally descended from a different (but related) package of experiments (many of which have been reported in Wilkenfeld *et al.* 2016), and follow their precedent of including roughly 40 participants per condition.

23 In Experiments 4 and 5, participants were compensated \$0.25 for an estimated four minutes of work.

participants were told that she was about to grab one of the unreliable ones but happened to get a reliable one by chance. The full *lucky library* vignette read as follows:

The Comanche are a Native American tribe that dominated the southern plains of North America in the 18th century. Rebecca has a lot of beliefs about the details of Comanche dominance during this time. Suppose that if you asked Rebecca any questions about this matter she would answer correctly. Assume further that Rebecca is answering from stored information; she is not guessing or making up answers, but is honestly answering what she confidently believes the truth to be based on what she read in a history book.

However, it turns out most of the history books are mistaken about the Comanche in the 18th century. Rebecca happened to pick up one of the correct ones, but she had no basis for preferring it to any of the inaccurate books. Her beliefs about the Comanche are true, but if she had picked up any other book she would have had entirely false beliefs about the Comanche.

In fact, she almost did grab one of the many inaccurate books off the shelf but, just as she was reaching for it, someone else took it. Rebecca settled on the closer one that was uniquely accurate (though she had no way of knowing that). Had that other person not taken the other book at just that moment, Rebecca would have had entirely false beliefs about Comanche dominance.

All participants were then assigned to rate Rebecca's knowledge, knowledge-why, understanding, and understanding-why (order randomized) by expressing their agreement or disagreement with the following claims on a seven-item Likert scale ranging from "Strongly Disagree" to "Strongly Agree."

Rebecca knows why the Comanche dominated the southern plains of North America in the 18th century.

Rebecca understands why the Comanche dominated the southern plains of North America in the 18th century.

Rebecca knows the history of Comanche dominance of the southern plains of North America in the 18th century.

Rebecca understands the history of Comanche dominance of the southern plains of North America in the 18th century.

Participants were asked two reading comprehension questions to ensure that they had understood the vignette: "Rebecca's beliefs about Comanche dominance of the southern plains of North America in the 18th century are mostly false," and "Most of the history books are mistaken about Comanche dominance of the southern plains of North America in the 18th century." Finally, they answered demographic questions.

4.4.2 Results

Responses were analysed with a mixed ANOVA, with within-subject factors of attribution (2: *knowledge, understanding*) and form (2: *objectual, propositional*) and a between-subjects factor of etiology (2: *normal, lucky library*). This revealed a main effect of attribution, with participants registering significantly more agreement with knowledge attributions than with understanding attributions, $F(1, 80) = 12.39$, $p = 0.001$, $\eta_p^2 = 0.134$. There was also a main effect of form, with people more likely to agree with attributions of objectual knowledge/understanding than propositional attributions (i.e., that Rebecca knew or understood why), $F(1, 80) = 8.29$, $p = 0.005$, $\eta_p^2 = 0.094$. There was no

main effect of etiology, $p = 0.419$, and there were no significant interactions ($ps > 0.185$) (see Figure 4). All ratings differed significantly from the scale midpoint, $ps < 0.001$.

4.4.3 Discussion

For the first time in our series of experiments, we found a significant difference between knowledge and understanding, but it was in the opposite direction from that predicted by Kvanvig: participants were more inclined to attribute knowledge than understanding. This suggests that understanding is actually more restrictive than knowledge, a suggestion we explore more fully in Wilkenfeld *et al.* (2016). For present purposes, the important result of the experiment is that there was once again no interaction between whether participants were asked to attribute knowledge or understanding and how sensitive they were to environmental luck – in this case because they showed no hesitation to attribute understanding or knowledge in the *lucky library* etiology at all. As in Experiment 2, comparisons between knowledge and understanding attributions were made within subjects, and so should have been maximally sensitive to any possible differences.

These results certainly challenge P2 (Kvanvig's prediction), and on the surface, may seem to support P3: the possibility that people attribute knowledge, but still withhold understanding, to people whose beliefs are only accidentally true. However, because the basis of differential attributions for knowledge and understanding seems not to have been the difference in etiology, we take them as stronger support for P1 or P4: that epistemic luck – if it does undermine attributions – undermines knowledge and understanding equally.

Experiment 4 was also the first in our series to investigate objectual understanding, and we found that people were slightly more inclined to attribute both objectual knowledge




Fig. 4. Experiment 4 ratings of how much participants agreed with a sentence asserting that Rebecca was in the positive epistemic state (knowledge, knowledge why, understanding, or understanding why) regarding the history of Comanche dominance, presented as a function of etiological condition. Ratings were made on a 7-point scale from completely disagree (1) to completely agree (7). Error bars are 1 SEM in each direction.

and understanding than propositional knowledge or understanding. This suggests that people are unlikely to be falling prey to the mistake Brogaard suggests, where their (hypothesized) disproportionate willingness to attribute propositional mental states makes people think objectual understanding is a more impressive achievement than knowledge(-why). A direct comparison of objectual understanding and propositional knowledge did not reveal a significant difference, $t(81) = -1.47$, $p = 0.144$.

The present results seem to go against Kvanvig's prediction. However, there is an obvious reply available: the lack of differentiation between people's attributions of understanding and knowledge is a consequence of limitations with our stimuli. Specifically, because we failed to find significantly lower attributions in the *lucky library* condition relative to the *normal* condition, one could argue that the experiment was not an adequate test of Kvanvig's predictions. Kvanvig could either acknowledge that he made his case too weak to produce the difference between understanding and knowledge that he predicts, or argue that the case was right all along but that people overlook the epistemic luck due to performance errors. In order to establish that knowledge does not outstrip understanding with respect to epistemic luck cases even for the sort of case Kvanvig puts forward, it is necessary to ramp up the potential difference by shifting from a case of environmental luck to one of apparent evidence.

4.5 Experiment 5

Experiment 4 seemed to show that, in the very case Kvanvig constructed to illustrate how considerations of epistemic luck disproportionately undermine knowledge relative to understanding, people's actual attributions were insensitive to etiology. There was, thankfully, an easy solution to this concern: people do seem responsive to veridical hallucination cases, as we saw in Experiments 1–2. If Kvanvig's general claim is true – that in the case of understanding attributions, but not knowledge attributions, etiology is irrelevant – he should also predict that people will attribute understanding, but not knowledge, in cases of veridical hallucination involving knowledge or understanding of the Comanche.

4.5.1 Method

Participants. Ninety-eight subjects (50 female, mean age 32) were recruited through MTurk. An additional 27 participants were excluded for incorrectly answering at least one reading comprehension question, with one more excluded for failing to answer every question.

Materials and Procedure. At the start of the experiment, participants were randomly assigned to read one of two vignettes. In the *normal* vignette, participants were presented with a standard story of someone (“Rebecca”) coming to learn about the Comanche dominance of the southern plains in the 18th century from a history book that was reliably researched and written. (This was similar to the control condition in Experiment 4, but modified slightly to match the *hallucinating writer* vignette more closely.) In the *hallucinating writer* vignette, the history book was instead written by an author in the grips of a hallucination. Yet, by chance, the content of his hallucination happened to be true claims about the Comanche. (We made the author of the book have the hallucination rather than Rebecca to match Kvanvig's original case more closely; this manipulation proved

sufficient to elicit a change in how people perceived the case.) The full *hallucinating writer* vignette reads as follows:

The Comanche are a Native American tribe that dominated the southern plains of North America in the 18th century. Rebecca has a lot of beliefs about the details of Comanche dominance during this time. Suppose that if you asked Rebecca any questions about this matter she would answer correctly. Assume further that Rebecca is answering from stored information; she is not guessing or making up answers, but is honestly answering what she confidently believes the truth to be based on what she read in a book.

It turns out that the book Rebecca read was written by someone under the influence of extremely potent psychotropic drugs. The author hallucinated knowing a great deal about a Native American tribe from the southern plains of North America in the 18th century called the Comanche. However, none of the beliefs he had as a result were causally connected to the truth of the actual Comanche.

However, quite by chance, the author of the book's hallucinations happened to line up exactly with reality. For example, the author wrote that a tribe called the Comanche spent three hours a day training, and the Comanche really did spend that long training. This coincidence was completely by luck – he could have just as easily hallucinated that they trained for 10 hours a day or not at all; improbably though, his hallucinations always lined up with the facts.

Rebecca, who read the book he wrote, ends up with a lot of true beliefs about the Comanche. However, all that information is really based on a series of incredibly lucky coincidences.

All participants were then assigned to rate Rebecca's objectual knowledge, knowledge-why, objectual understanding, and understanding-why, by expressing agreement or disagreement with the same sentences and scale as used in Experiment 4.

Participants were asked two reading comprehension questions to ensure that they had understood the vignette: "Rebecca's belief about Comanche dominance of the southern plains of North America in the 18th century are mostly false," and "The author of the book Rebecca read based his claims on sound research." Finally, participants answered demographic questions.

4.5.2 Results

Responses were analysed with a mixed ANOVA, with within subjects factors of attribution (2: *knowledge, understanding*) and form (2: *objectual, propositional*), and a between subjects factor of etiology (2: *hallucinating writer, normal*) (see Figure 5). As with Experiment 4, there were significant main effects of attribution type (people registered more agreement with attributions of knowledge than of understanding), $F(1, 96) = 16.61, p < 0.001, \eta_p^2 = 0.147$, and of form (people registered more agreement with claims of objectual than propositional knowledge/understanding), $F(1, 96) = 4.30, p = 0.041, \eta_p^2 = 0.043$. As expected, there was also a main effect of etiology, with participants less inclined to attribute either knowledge or understanding in the *hallucinating writer* (apparent evidence) case: $F(1, 96) = 14.10, p < 0.001, \eta_p^2 = 0.128$. Interestingly, people showed differential sensitivity to the difference in etiology, with knowledge-why, understanding-why, and objectual understanding all undermined ($ps \leq 0.002$), but no significant impact on objectual knowledge ($p = 0.124$). As a result, there were significant interactions between attribution and etiology, $F(1, 96) = 5.16, p = 0.025, \eta_p^2 = 0.051$, between form and etiology, $F(1, 96) = 4.80, p = 0.031, \eta_p^2 = 0.048$, between attribution and form,

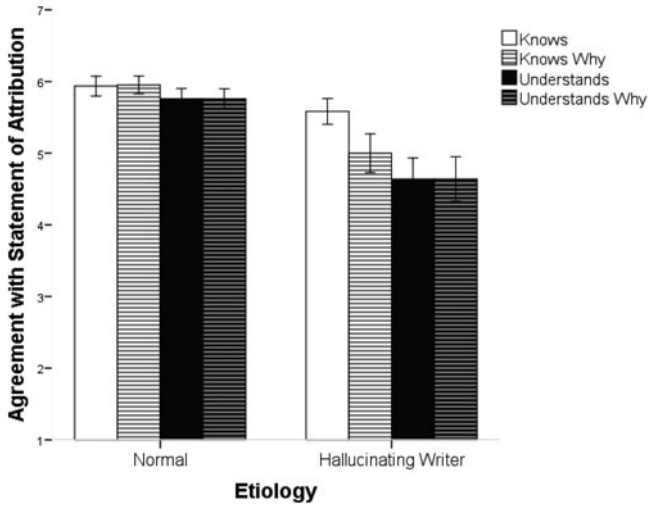


Fig. 5. Experiment 5 ratings of how much participants agreed with a sentence asserting that Rebecca was in the positive epistemic state (knowledge, knowledge why, understanding, or understanding why) regarding the history of Comanche dominance, presented as a function of etiological condition. Ratings were made on a 7-point scale from completely disagree (1) to completely agree (7). Error bars are 1 SEM in each direction.

$F(1, 96) = 6.53, p = 0.012, \eta_p^2 = 0.064$, and between attribution, object, and etiology, $F(1, 96) = 7.29, p = 0.008, \eta_p^2 = 0.071$.²⁴ All ratings differed significantly from the scale midpoint, $ps < 0.050$.

4.5.3 Discussion

This experiment reinforced the picture that began to emerge in Experiment 4: surprisingly, it is easier to undermine understanding than knowledge on the basis of etiological concerns. Specifically, participants were more inclined to attribute knowledge than understanding across all of our cases. Unlike Experiment 4, however, the manipulation of epistemic luck successfully lowered attributions: participants were generally less inclined to attribute knowledge or understanding in the *hallucinating writer* condition. There was also an interaction between etiology and attribution type, but it was driven almost entirely by the fact that objectual knowledge claims appeared uniquely immune to concerns about deviant etiology, even in the relatively efficacious veridical hallucination

24 At the suggestion of Wesley Buckwalter, we also ran a forced-choice version of the task (based on Turri 2013) in which, rather than express agreement on a Likert scale, participants ($N = 107$) were asked whether Rebecca “really [knew/understood]” or “only thought she [knew/understood]” (for both objectual and propositional knowledge and understanding). The results were much the same. People chose at chance for the *hallucinating writer* vignettes ($ps > 0.571$) but were much more likely to attribute real knowledge and understanding in the normal vignette ($ps < 0.001$). Moreover, while responses to every question in the *hallucinating writer* condition differed significantly from responses to every question in the *normal* condition ($ps < 0.033$), none of the responses to questions within the *hallucinating writer* condition ($ps > 0.427$) or within the *normal* condition ($ps > 0.168$) differed significantly from each other.

cases. Seemingly, when you want to figure out who knows history, as opposed to why things happened in history, it is sufficient that they have latched onto the right answer somehow or other.

When we consider only attributions in objectual form, these findings provide some evidence for P₃: that it is more important for understanding attributions than knowledge attributions that they be caused in the right sort of way. When we instead consider attributions in propositional form, the findings provide some evidence for P₁: that both knowledge-why and understanding-why are incompatible with luck. The findings thus offer qualified support for these two positions, while challenging P₂ and P₄.

4.6 General Discussion

In five experiments, we explored the extent to which people's attributions of knowledge and understanding are undermined by epistemic luck. We found that such attributions are surprisingly resilient, and, most importantly, equivalently resilient. Across exploding chestnut, Swampman, and Comanche cases – precisely the cases taken to be at issue between Kvanvig and Grimm – people attribute knowledge and understanding with equal insensitivity to environmental luck, and equal sensitivity to luck caused by apparent evidence.

We do not claim to have tested the full gamut of possible types of epistemic luck. Other examples include cases where the causal chain between truth and belief is broken, but not in a way that reflects badly on the believer's cognitive mechanisms. It would also be desirable to consider more realistic versions of our scenarios. However, while we anticipate that such variations would affect overall agreement with attributions of knowledge and understanding, we see no reason to expect such variations to have *differential* effects, with attributions of understanding outstripping attributions of knowledge.

Looking back at our four possible patterns, we see that none of them explains all the data. Some challenges to etiology – those that are based on hallucinations (what Starmans and Friedman (2012) call “apparent evidence”) – undermine *both* knowledge(-why) and understanding(-why), in accordance with P₁, which states that both understanding and knowledge are sensitive to etiological concerns. Other challenges to etiology – those that are more analogous to traditional barn-façade cases – undermine *neither* knowledge nor understanding, as predicted by P₄, which states that neither understanding nor knowledge is sensitive to etiological concerns. If we expand our attention from propositional knowledge to objectual knowledge, however, we find that people consider etiology a defeater of objectual understanding but not objectual knowledge, as predicted by P₃ (that understanding but not knowledge is sensitive to etiological concerns). Thus, which pattern people exhibit depends on both the type of etiological deviance and on the type of understanding and knowledge under consideration. However, we did not find any empirical support for P₂: the view we associate with Kvanvig, that attributions of knowledge, but not those of understanding, are sensitive to etiological concerns.

People's willingness to attribute knowledge and understanding in barn-façade-style cases (lucky environments) is potentially surprising to someone steeped in traditional epistemology, but it is not wholly unexpected given other results in experimental epistemology. As already discussed, Starmans and Friedman (2012) found that people were quite willing to attribute knowledge in lucky environments (but not for apparent evidence), which accords perfectly with the results from Experiments 1, 2, and 5 (but see footnote

17).²⁵ Our data thus provide some additional support for Starmans and Friedman’s claim that most people are simply insensitive to epistemic luck, unless there is a clear break in the causal connection between truth and belief.

Overall, our data largely support Grimm’s position at the expense of Kvanvig’s. While people do seem to intuit that Becky understands the exploding chestnut (i.e., they don’t take understanding to be undermined by mere luck in a bad environment), the more general point that people’s attributions of understanding track their attributions of knowledge seems to be correct. Of course, it remains possible that a more sensitive or targeted experiment would generate the “correct” response to the epistemic luck cases in regard to knowledge while leaving understanding untouched, but that does not seem probable given the present results.²⁶

While we did not find any evidence that knowledge is more demanding than understanding, Experiments 4 and 5 began to suggest ways in which the reverse may be true. Experiment 5, in particular, suggests that at least some sorts of knowledge attributions (i.e., objectual knowledge attributions) are less demanding than corresponding claims about understanding. There is thus some limited indication that understanding might require knowledge (and hence be at least as susceptible to epistemic luck), but not conversely, though we acknowledge that the evidence is far from conclusive.

Finally, it’s worth acknowledging potential limitations of the current work. Our sample sizes were not large, and one of our main conclusions – the equivalence of attributions of knowledge and understanding – rests on null results. The number of our studies and their consistency, however, mitigates this concern. Another worry is that our experiments involved explicit attributions of knowledge and understanding, without a measure of how these attributions affected “epistemic behaviors,” such as deference, trust, or information gathering, which are relevant to reverse-engineering the functions of folk epistemic concepts. We think that an investigation of such epistemic behaviors is an important step for future research.

5. CONCLUSION

In this paper, we have argued for a robust equivalence between people’s sensitivity to epistemic luck when attributing knowledge and when attributing understanding. This is not to say that understanding and knowledge are identical, even in folk conception. Elsewhere (Wilkenfeld *et al.* 2016), we do find empirical evidence of systematic differences in their attribution – this alleviates the potential concern that our concepts of knowledge and understanding are simply too coarse-grained to ever come apart. However, along the particular dimension explored by Kvanvig and Grimm, we find almost no difference (and, where there is a difference, not the one anyone predicted).

25 In somewhat more intermediate results, Colaço *et al.* (2014) find that people are somewhat sensitive to barn-façade cases. Since our main target was not absolute level of response to those cases, we take our results neither to confirm nor to conflict with theirs.

26 In another experiment, not reported here, we explored whether merely inserting the word “really” before “knows” or “understands” had any differential impact on knowledge and understanding attributions. We found that attributions were depressed in the various cases of etiological deviance, but not differentially.

Our findings are only a first step towards characterizing the folk concept of understanding and how it relates to knowledge, but already they rule out some possibilities and point to promising directions for further study. To the extent that we can read off our epistemic aims from our patterns of attribution, the findings rule out some possibilities – most notably that we use knowledge attributions, but not understanding attributions, to demarcate a particularly admissible etiological history. Our findings also suggest that we might do well to treat understanding and knowledge as roughly of a kind, and to look for ways in which understanding is *more* restrictive, not *less* restrictive, than knowledge.²⁷

REFERENCES

- Alexander, J. and Weinberg, J. M. 2007. Analytic epistemology and experimental philosophy. *Philosophy Compass*, 2: 56–80.
- Blouw, P., Buckwalter, W. and Turri, J. Forthcoming. ‘Gettier Cases: A Taxonomy.’ In R. Borges, C. de Almeida and P. Klein (eds), *Explaining Knowledge: New Essays on the Gettier Problem*. Oxford: Oxford University Press.
- Brogaard, B. 2005. ‘I Know. therefore I Understand.’ <https://sites.google.com/site/brogaard/> (accessed 1 January 2013).
- Cappelen, H. 2014. ‘Replies to Weatherson, Chalmers, Weinberg, and Bengson.’ *Philosophical Studies*, 171: 577–600.
- Colaço, D., Buckwalter, W., Stich, S. and Machery, E. 2014. ‘Epistemic Intuitions in Fake-barn Thought Experiments.’ *Episteme*, 11: 199–212.
- Craig, E. 1999. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.
- Cullen, S. 2010. ‘Survey-driven Romanticism.’ *Review of Philosophy and Psychology*, 1: 275–96.
- De Regt, H. W. and Dieks, D. 2005. ‘A Contextual Approach to Scientific Understanding.’ *Synthese*, 144: 137–70.
- Foley, R. F. 1996. ‘Knowledge is Accurate and Comprehensive Enough True Belief.’ In J. L. Kvanvig (ed.), *Warrant and Contemporary Epistemology: Essays in Honor of Plantinga’s Theory of Knowledge*, pp. 87–96. Savage, MD: Rowman and Littlefield.
- Gettier, E. 1963. ‘Is Justified True Belief Knowledge?’ *Analysis*, 23: 121–123.
- Goldman, A. I. 1976. ‘Discrimination and Perceptual Knowledge.’ *Journal of Philosophy*, 73: 771–791.
- Grimm, S. R. 2006. ‘Is Understanding a Species of Knowledge?’ *British Journal for the Philosophy of Science*, 57: 515–35.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kember, D. 1996. ‘The Intention to both Memorise and Understand: Another Approach to Learning?’ *Higher Education*, 31: 341–54.
- Khalifa, K. Forthcoming. In *Understanding: An Epistemology of Scientific Understanding*. Cambridge: Cambridge University Press.
- Kvanvig, J. L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lombrozo, T. 2011. ‘The Instrumental Value of Explanations.’ *Philosophy Compass*, 6: 539–51.

27 In addition to our grant sources (the Varieties of Understanding Project funded by the Templeton Foundation, a James S. McDonnell Foundation Scholar Award in Understanding Human Cognition to Tania Lombrozo) and home institutions, we would like to thank participants at the Society for Philosophy and Psychology annual meeting and the *Varieties of Understanding* mid-point conference, many of whom provided essential comments. We’d like to give particular thanks to Kareem Khalifa for extensive written feedback.

- Morgan, M. S. and Morrison, M. 1999. *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge: Cambridge University Press.
- Nagel, J., San Juan, V. and Mar, R. A. 2013. 'Lay Denial of Knowledge for Justified True Beliefs.' *Cognition*, 129: 652–61.
- Pritchard, D. 2009. 'Knowledge, Understanding and Epistemic Value.' *Royal Institute of Philosophy Supplement*, 64: 19–43.
- Rosenberg, J. 1981. 'On Understanding the Difficulty in Understanding Understanding.' In H. Parret and J. Bouveresse (eds), *Meaning and Understanding*, pp. 29–43. Berlin: W. De Gruyter.
- Smith, R. 2014. 'Explanation, Understanding, and Control.' *Synthese*, 191: 4169–200.
- Starmans, C. and Friedman, O. 2012. 'The Folk Conception of Knowledge.' *Cognition*, 124: 272–83.
- Turri, J. 2013. 'A Conspicuous Art: Putting Gettier to the Test.' *Philosophers' Imprint*, 13: 1–16.
- Weinberg, J. 2006. 'What's Epistemology For? The Case For Neopragmatism in Normative Metaepistemology.' In S. Hetherington (ed.), *Epistemological Futures*, pp. 26–47. Oxford: Oxford University Press.
- , Nichols, S. and Stich, S. 2001. 'Normativity and Epistemic Intuitions.' *Philosophical Topics*, 29: 429–60.
- Wilkenfeld, D. A. 2013. 'Understanding as Representation Manipulability.' *Synthese*, 190: 997–1016.
- , Plunkett, D. and Lombrozo, T. 2016. 'Depth and Deference: When and Why We Attribute Understanding.' *Philosophical Studies*, 173: 373–93.
- Wright, J. C. 2010. 'On Intuitional Stability: The Clear, the Strong, and the Paradigmatic.' *Cognition*, 115: 491–503.
- Zagzebski, L. 2001. 'Recovering Understanding.' In M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, pp. 235–51. Oxford: Oxford University Press.

DANIEL WILKENFELD is a post-doc in the History and Philosophy of Science department at the University of Pittsburgh. He is working on the nature of explanation and understanding, from both theoretical and empirical perspectives.

DILLON PLUNKETT is a graduate student in the Department of Psychology at Harvard University researching the neuroscience of high-level reasoning and conscious cognition.

TANIA LOMBROZO is an Associate Professor of Psychology at the University of California, Berkeley. Her current research focuses on explanation, causal reasoning, social cognition, and folk epistemology.
