

Brute Facts, Knowledge, and Senses of Understanding

1. Introduction

Having an explanation, whether the product of rigorous scientific investigation or in response to everyday explanatory requests, improves our overall epistemic position, and it is natural to think that explanations improve our epistemic standing by conferring understanding. My first aim in what follows is to draw out a knowledge-based sense of understanding and distinguish it from other senses of understanding, including explanatory understanding. My second goal is to show how these conceptual distinctions require reevaluating several extant discussions of explanation and understanding, including those that have been premised on the notion of a brute fact. While I shall not offer a positive account of explanatory understanding, the considerations that I advance will support the idea that there is a distinctive sense of explanatory understanding, and thus that there is a distinctive epistemic value associated with explanation.

2. Explanation, Understanding, and Knowledge

It is natural to think of the epistemic value of explanation in terms of *knowing why* something is the case rather than merely *knowing that* something is the case (see, for example, Kim 1994). This intuitive distinction between *knowing why* and *knowing that*, however, requires some comment.

The first thing to note is that even if *knowing why* can be explicated in terms of *knowing that* (say, in terms of knowledge that an explanatorily relevant causal relation obtains) we can nonetheless hold that there is a distinctive value associated with the relevant *knowledge that*. For even if having explanatory understanding amounts to having a certain kind of propositional knowledge, it may be that having this sort of propositional knowledge confers a certain kind of epistemic gain not conferred by other sorts of propositional knowledge.¹ It will be helpful to keep this point in mind in subsequent discussion, since I

¹ Moreover, it is not clear that drawing a connection between certain kinds of *knowledge that* and *knowledge why* should raise any special problems for the suggestion that there is a specific epistemic value associated with explanation. The story here, it seems, will conform to the following schema: having certain propositional knowledge amounts to having an explanation; and there is a conceptual connection between having an explanation and having explanatory understanding. The task of the theory of explanation will be to specify the relevant propositional knowledge. This is put forward, in part, as a response to some of the worries raised by Kim about the extent to which extant theories of explanation draw a substantial connection between explanation

will be emphasizing the significance of a knowledge-based sense of understanding and the distinction between this sense of understanding and explanatory understanding. The preceding remarks show that such a distinction will not be threatened by the idea that having explanatory understanding may require possessing certain kinds of knowledge, or at least certain kinds of true belief (see below).

That there *is* a knowledge-based sense of understanding, which we can refer to as understanding_K, can be brought out by noting that where a subject S comes to know a certain fact about the world, we are inclined to say that S better understands how things are than prior to obtaining that piece of knowledge. This is especially the case if, prior to coming to know a certain fact, S held a false belief about the subject matter in question: in replacing a false belief with a piece of knowledge, we think that there is an improvement in S's understanding of the world. However, in attributing this sort of understanding to a subject S, we are not thereby attributing explanatory understanding, or explanatory knowledge, to S. This is because in attributing understanding_K to S, we are not supposing that the knowledge that S possesses is explanatory (say, that it is knowledge of an explanatorily relevant causal relation). Further, while I shall primarily be concerned with knowledge-based understanding in what follows, it is plausible that there is also sense in which, other things being equal, a subject with *true beliefs* better understands the world than a subject who has false beliefs, regardless of whether the beliefs in question count as *knowledge*. We can refer to a sense of understanding associated with true belief as understanding_{TB}. Like understanding_K, understanding_{TB} does not imply explanatory understanding, since the true beliefs that may confer understanding_{TB} may not be explanatorily relevant (for example, they may not be beliefs about causal relations). Understanding_K will imply understanding_{TB} in virtue of knowledge implying true belief, but the converse will not hold, since knowledge requires more than true belief.

We might think that while knowledge-based understanding does not imply explanatory understanding, explanatory understanding implies knowledge-based understanding. After all, we saw above that explanatory understanding can be naturally put as *knowledge why* things are a certain way. It

and understanding (1994, 52 – 54). The present suggestion is that there may not be much more to say here beyond simply noting, as Kim does, that it seems to be part of the concept of an explanation that having an explanation improves understanding.

seems to me, however, that absent a compelling reason to think otherwise, we should *not* take this intuitive connection between explanatory understanding and knowledge to involve a robust concept of *knowledge*—the concept that gets debated in epistemology, that requires more than justified true belief, and so on. The discussion above in terms of the distinction between *knowledge that* and *knowledge why*, and how we can insist on a distinctive epistemic value associated with *knowledge why* even if it is explicated in terms of *knowledge that*, does not seem to presuppose a substantive conception of knowledge and can be easily recast in terms of less-demanding positive epistemic states (like true belief).²

More generally, it *is* plausible that explanatory understanding requires true belief: on a causal view of explanation, to understand why an event occurred will involve having true beliefs about the causal history of that event. In this sense, explanatory understanding will imply understanding_{TB}. But it is much less clear that it should require *knowledge*. It is not clear, for example, that explanatory understanding of an event should require *knowledge* of the causal history of that event, particularly once we note that whether a true belief counts as *knowledge* may turn on factors (say, the etiology of the belief) that do not seem to determine whether one has an *explanation* of an event (a point that will be of some importance in §5). In this case, explanatory understanding will not imply understanding_K. It is, of course, compatible with this that knowledge-based understanding can go together with explanatory understanding; we do not think that in coming to *know* that a certain event has a causal history, rather than merely truly believing that it has that causal history, we thereby lose our explanatory understanding of that event. The present claim, rather, is just the subject may have explanatory understanding of the event prior to *knowing* that it has a certain causal history, so long as the subject *truly believes* that it has that causal history. If this is right, while we may continue to speak of the epistemic value of having an explanation in terms of “explanatory knowledge”, we should insist that this may involve a positive epistemic state, such as true belief, that falls short of the demanding conditions needed for *knowledge*.

² Similarly, when Kim speaks of explanatory knowledge, it is doubtful that he takes this to involve a substantive, robust conception of knowledge.

3. Barnes on the Explainability of Brute Facts

It is an interesting question what facts about the world are brute and, more generally, what sort of facts *can* be brute.³ Reflection on the concept of a brute fact may also be capable of providing insight into explanation and understanding. Nonetheless, there has not been much sustained discussion of the notion of a brute fact. And I shall now argue that extant discussions of the notion of a brute fact have neglected the aforementioned distinction between understanding_K and explanatory understanding, and that this neglect threatens the conclusions that have been drawn from reflection on the concept of a brute fact.

According to the first line of thought that I shall consider, brute facts can, in fact, be explained. Indeed, Eric Barnes argues that several prominent theories of explanation can be faulted precisely on the grounds that they imply that brute facts are “unexplainable” and thus represent “a lack of scientific understanding” (1994, 61). For Barnes, a brute fact is one with no *explanatory basis*; he contends that it is compatible with this that such facts can be *explained* (and thus that *having an explanation* does not require *having an explanatory basis*). It is true, Barnes concedes, that if a fact F is brute, but is not *known* that F is brute, then we may lack some understanding with respect to F. However, he writes,

[I]f F is known to be brute, then it seems to me that F represents no mystery whatsoever—F is simply partly constitutive of the way the world is... [O]ur understanding suffers a gap just in case there is some hidden explanatory basis for a fact that we hold true—where there is no hidden explanatory basis, and we know this, there is nothing lacking in our understanding—for there is no explanation that we fail to have.

He continues,

I see no reason not to go further and claim that a correct theory of explanation ought to entail that brute facts are perfectly explainable: their explanation consists of the stipulation that such facts are brute (64 – 65).⁴

This is startling view; it implies that all facts can be explained, since all facts will either have an “explanatory basis” or will be explainable by noting that they lack an explanatory basis.⁵ While I am not

³ Chalmers 1996, for example, argues against the suggestion that there can be brute “inter-world facts”, facts about the space of possible worlds; see also XX.

⁴ I have made some changes to the notation for consistency. I have also omitted Barnes’ use of the word “ontological” to qualify the uses of “brute” in these passages; Barnes uses this qualification to distinguish brute facts that are known to be brute from those that are “epistemically brute”, where an epistemically brute fact is one that may have an explanatory basis but we do not know what it is.

sure that it makes good sense to say that an explanation for a fact can consist in noting that that fact is brute—or even whether there is a clear distinction between *having an explanation* and *having an explanatory basis*—we can set aside these worries, since it seems that Barnes illicitly conflates understanding_K with explanatory understanding. In particular, he seems to reason as follows:

- EB1. There is a sense in which brute facts do not “represent a gap in our understanding”.
- EB2. If brute facts do not represent a gap in our understanding, then they can be explained.
- EB3. So, brute facts can be explained—in particular, they can be explained by noting that they are brute, and known to be so.

If a fact F is brute, and it is known that F is brute, there is a sense in which F’s lacking an “explanatory basis” does not threaten our understanding of the world: there is a sense of understanding under which it is not correct to say that once we know that F is brute, our understanding of the world is impugned by F’s being a brute fact. The problem, however, is that we can account for this by appealing to the knowledge-based sense of understanding described in §2. Since understanding_K does not require the having of an explanation or explanatory understanding, the sense in which brute facts “do not represent a gap in our understanding” will not support such facts being explainable.

This interpretation of the sense in which brute facts may not “represent a gap in our understanding” is supported by the explicit remarks that Barnes provides in favor of EB1. Thus consider again the claim that “where there is no hidden explanatory basis, and we know this, there is nothing lacking in our understanding—for there is no explanation that we fail to have”. (1994, 66). The plausibility of this claim is readily accountable in terms of understanding_K. When we learn that a fact is brute, we learn something about our world, and we come to know something about how our world that we previously did not know. And once we know that a fact is brute, there is a sense in which it will not “represent a scientific mystery” (ibid). But there is little reason to think that the sense of understanding at work here involves anything more than understanding_K.⁶

⁵ As noted in Fahrback 2005.

⁶ Notably, once we invoke the notion of understanding_K, it is not quite right to say, as Barnes does, that once we learn that a fact F is brute, “there is nothing lacking in our understanding”. It is true that if we know that F is brute, whereas we previously mistakenly took F to have an explanation, we improve our understanding_K of the world. But we will nonetheless lack explanatory understanding of F: while F will not represent a “scientific mystery” in the sense of understanding_K, it is entirely compatible with

4. Fahrbach on Explanation and Understanding

In his insightful paper “Understanding Brute Facts”, Ludwig Fahrbach agrees with Barnes that brute facts may not represent a gap in our understanding, but denies that such facts can be explained. Fahrbach thus contends that the notion of a brute fact can allow us to distinguish understanding from explanation, a sense in which understanding does not require explanation. In response, I will first suggest that if we do *not* invoke a distinction between explanatory understanding and understanding_K, Fahrbach’s position may in fact threaten the idea that there is a distinctive epistemic value associated with having an explanation. Second, I will argue that once we make a distinction between explanatory understanding and understanding_K, we should not follow Fahrbach (2005, 460) in holding that the concept of being a brute fact provides a basis for an interesting distinguishing between explanation and understanding.

Regarding the first point, suppose we follow Fahrbach in holding that a brute fact F may not represent a “scientific mystery”, even if F cannot be explained. Now, Fahrbach writes that his position supports the claim that “the epistemic gain imparted by an explanation is different from the epistemic gain imparted by the statement that a fact is brute” (ibid). Yet he assumes that the epistemic gain associated with learning that a fact is brute can be described as a kind of understanding. In this case, however, absent some reason to think otherwise we could just as well reason that since brute facts cannot be explained (as Fahrbach holds), but do not threaten our understanding of the world, we should conclude that the epistemic gain associated with explanation does *not* require explanation. To put things a bit differently, we may just as well suppose that there is a single kind of epistemic value, understanding, that while conferred by the having of an explanation can *also* be possessed without possessing an explanation, since it can be possessed with respect to facts that cannot be explained. In this case, we will not be able to maintain that there is a *distinctive* epistemic value associated with having an explanation.

Such a worry, however, is immediately dispelled once we invoke the distinction between understanding_K and explanatory understanding. In particular, we should hold that insofar as we are

this that F cannot be understood in an explanatory sense. This is precisely what we should expect if we hold, as is plausible, that a brute fact is simply one that cannot be explained.

inclined to spell out the epistemic gain associated with learning that a fact is brute in terms of understanding, this should be restricted to a claim about understanding_K. But since understanding_K is not explanatory understanding, there will be no temptation to think that the sense of understanding associated with learning that a fact is brute threatens the idea that there is a unique epistemic value associated with having an explanation, and that this value can be described as a kind of understanding.

Similar remarks apply to Fahrbach's reasons for thinking that reflection on the concept of a brute fact can provide a basis for distinguishing explanation from understanding. Thus he writes (459 – 60):

...the epistemic process that occurs in [a subject] Smith when he learns that F is brute results in an improvement in his understanding of the world. When Smith requests an explanation of F, it is because his understanding of the world suffers from a gap. It was incomplete with respect to F. When he gains the information that F has no explanation, the gap in his understanding is closed. [...] Now, let us assume that Smith not only learns that F is brute, but comes to know everything else about F and its place in the order of nature. If this is the case and if his epistemic state is optimal with respect to F, then it is plausible to assert that his understanding of F is complete.⁷

Fahrbach concludes from this passage that “the concept of being a brute fact provides a basis for distinguishing between the concept of explanation and the concept of understanding”, since it shows that brute facts can be fully understood without being explainable. But the key aspects of this passage are naturally explicated in terms of understanding_K, and in this case will only support the unsurprising conclusion that understanding_K can be possessed without explanation. Upon learning that a fact F is brute, it is natural to say that the epistemic “gap” that gets closed is precisely a gap in knowledge: in requesting an explanation for F, Smith falsely believed that F is not brute; upon learning that F is brute, Smith no longer falsely believes that F is brute, and perhaps may be said to know that F is brute. Smith's false belief is thus replaced with knowledge, and this his understanding_K of the world improves. But this improvement in understanding_K need have nothing to do with the sort of understanding we might think goes along with the having of an explanation. Contra Fahrbach, then, considerations relating to brute facts do not support a significant distinction between explanation and understanding; the distinction here only amounts to the unsurprising possibility of understanding_K without explanation.

⁷ I have made some changes to the notation for consistency.

In advancing these critical remarks both here and in §3, I do not intend to suggest that there is nothing that can be learned about explanation and understanding by appealing to the notion of a brute fact. And I believe that there are interesting questions about our endorsement of the “bruteness” of a fact, and when such an endorsement is warranted.⁸ But if my conclusions are on track, extant work on the notion of a brute fact has not succeeded in producing substantive conclusions about explanation and understanding.

5. Knowledge, Understanding_k, and Understanding

If the considerations advanced in §2 are on track, we can consistently maintain that explanatory understanding *is* a kind of knowledge *and* that there is a distinctive epistemic value associated with explanatory understanding.⁹ Nonetheless, it is worth considering how the notion of understanding_k works into the extant debate about the extent to which understanding is a kind of knowledge. I shall claim that one reason why we might think that understanding is *not* a species of knowledge can be defended, in part, precisely by invoking the distinctions sketched in §2.

According to Jonathan Kvanvig, understanding can be distinguished from knowledge on the grounds that while knowledge does not tolerate *accident*, understanding does. In defending this view, Kvanvig develops a case in which a subject S learns, through reading a textbook, how the Comanche came to dominate the southern plains of the United States. However, suppose that most textbooks about Native Americans are full of misinformation, and that one was quite lucky in picking up a textbook with true information about the Comanche. On one hand, this “epistemic luck” seems to threaten S’s knowledge about the Comanche. But, Kvanvig suggests, it does not threaten S’s understanding; we should just say that S is lucky to understand how the Comanche came to dominate the southern plains (2003, 198 – 199). So, understanding can be distinguished from knowledge on the grounds that it is insensitive to its etiology precisely where knowledge *is* sensitive.

⁸ See XX.

⁹ As noted in §2, this claim can be made with epistemic states less demanding than knowledge and is compatible with the plausible suggestion, which I shall now defend in more detail, that explanation does not require knowledge, but rather some less demanding positive epistemic state like true belief.

In a recent critical discussion, however, Stephen Grimm argues that understanding *is* sensitive to such etiological considerations. Grimm's arguments, however, are naturally understood as concerning understanding_K, in which case it should hardly be surprising that understanding can be defeated in the same manner as knowledge. This, of course, is entirely compatible with other sorts of understanding, including explanatory understanding, being immune to defeat in such a manner.

Grimm first presents the following hypothetical case in favor of his position, which he says raises the "veridical hallucination problem" for Kvanvig's position (2006, 520):

... suppose that the CIA slips a hallucinogen into Albert's coffee, and that as a result he "sees" his dog bump into the table, causing a vase to crash to the floor. Putting things together, he takes himself to understand why the vase fell from the table and crashed to the floor: because the table was bumped by the dog. As it happens, moreover, this hallucination exactly matches the events that are actually unfolding in front of him.

However, while Albert truly believes that the vase fell because the table was bumped by the dog, "we're still tempted to say", Grimm writes, that "Albert doesn't really understand why the vase fell—specifically, he doesn't understand that the vase fell because the table was bumped by the dog." And the reason, he contends, is because Albert too easily might have misidentified the cause of the fall (2006, 520 – 521).

It should be noted that while Grimm does not draw such a connection, if we endorse his position on understanding, the relationship between having an explanation and having understanding will become somewhat obscure. In particular, even if we follow Grimm in holding that Albert does not understand why the vase fell, it seems right to say that Albert can explain why the vase fell (under a broadly causal view of explanation). It would thus seem that if we follow Grimm, having an explanation will not itself confer understanding with respect to the explanandum. Somewhat differently, having a causal explanation does not seem to require *knowing* that the given causal relation obtains; it rather simply seems to require truly believing that the given causal relation obtains. But if understanding requires knowledge, explanation will not itself be capable of conferring understanding.

I personally do not find Grimm's assessment of the "Comanche case" described to be *especially* compelling; it is not compelling in the same sense in which we are compelled, in Gettier's examples, to conclude that the subject does not possess the relevant piece of knowledge (Gettier 1963). But we have

the resources to explain Grimm's assessment: Albert does not understand_K why the vase fell, since Albert does not *know* that the vase fell because the table was bumped by the dog. Indeed, we can interpret Grimm as showing that we can distinguish understanding_K from understanding_{TB}. In the case that he describes, Albert has explanatory understanding and any understanding that merely requires true belief. But he does not have, and cannot have, understanding_K, since he does not possess the relevant knowledge. This provides a coherent analysis of Grimm's intuition without forcing us to draw a general conclusion about understanding or even a conclusion about explanatory understanding; we get the unsurprising conclusion that understanding_K can be defeated by the sort of considerations (in particular, those relating to luck and accident) that can defeat knowledge. Essentially the same considerations can be applied to the other cases that Grimm advances.¹⁰ We can thus insist on a close connection between explanatory understanding and having an explanation: as the quality of an explanation does not seem to turn on its etiology, so also can we maintain that explanatory understanding, while requiring true belief, does not require knowledge, and so is not sensitive to its etiology.

6. Conclusion

Understanding is an elusive concept and the elusiveness of understanding may to some extent explain well-known difficulties in providing a counterexample-free analysis of explanation. The suggestion implicit in my discussion is that we can make some progress here by distinguishing different senses of understanding and should set aside a knowledge-based sense of understanding in our discussions of the epistemic value of explanation—that is, explanatory understanding. Drawing such a distinction, I argued, allows us to reassess the significance of the notion of a brute fact for our position on explanation and understanding. Moreover, it allows us to defend an attractive way of distinguishing understanding from

¹⁰ For example, Grimm describes the following “lucky environment” case in which a subject Becky observes a blacksmith who enjoys testing his timing by hammering chestnuts at the precise moment they explode from the heat of an anvil. On this occasion, however, “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” (521). So, Becky believes that the chestnut exploded because of the hammer blow and her belief is true. Intuitively, however, Grimm contends that she does not “genuinely understand” why the chestnut shattered, and that the problem has to do with the luck involved in her having a true belief about the cause of the chestnut's shattering. But the same “divide and conquer” strategy can be advanced here: Becky does not understand_K why the chestnut shattered; but it is compatible with this that she has explanatory understanding of why the chestnut shattered, since explanatory understanding does not require understanding_K.

knowledge by allowing us to concede that there is a sense of understanding that is associated with knowledge, but that it is compatible with this that other senses of understanding, including explanatory understanding, may not be perspicuously taken to be a kind of knowledge.

References

- Chalmers, D. 1996. *The Conscious Mind*. New York: Oxford University Press.
- Barnes, E. 1994. Explaining Brute Facts. *PSA: Proceedings of the Philosophy of Science Association 1994* 1: 61 – 68.
- Fahrbach, L. 2005. Understanding Brute Facts. *Synthese* 156: 449 – 466.
- Gettier, E. 1963. Is Justified True Belief Knowledge? *Analysis* 23: 121 – 123.
- Grimm, S. 2006. Is Understanding a Species of Knowledge? *The British Journal for the Philosophy of Science* 57: 515 – 535.
- Kim, J. 1994. Explanatory Knowledge and Metaphysical Dependence. *Philosophical Issues 5: Truth and Rationality*: 51 – 69.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.