

Is Understanding A Species Of Knowledge?

Stephen R. Grimm

ABSTRACT

Among philosophers of science there seems to be a general consensus that understanding represents a species of knowledge, but virtually every major epistemologist who has thought seriously about understanding has come to deny this claim. Against this prevailing tide in epistemology, I argue that understanding is, in fact, a species of knowledge: just like knowledge, for example, understanding is not transparent and can be Gettiered. I then consider how the psychological act of “grasping” that seems to be characteristic of understanding differs from the sort of psychological act that often characterizes knowledge.

- 1 *Zagzebski's account*
 - 2 *Kvanvig's account*
 - 3 *Two problems*
 - 4 *Comanche cases*
 - 5 *Unreliable sources of information*
 - 6 *The upper-right quadrant*
 - 7 *So is understanding a species of knowledge?*
 - 8 *A false choice*
-

Among philosophers of science there seems to be a general consensus that understanding represents a *species of knowledge*. Thus, for instance, we find Peter Lipton arguing that a good theory of understanding should characterize its subject in a way that is “unmysterious and objective,” and that the first step toward reaching this goal is to see that, “Understanding is not some sort of super-knowledge, but simply more knowledge: knowledge of causes” (Lipton [2004], p. 30). Peter Achinstein likewise writes that, “Explaining q has been defined as uttering something with the intention of rendering q understandable... Such understanding I take to be a

form of knowledge” (Achinstein [1983], p. 23).¹ Of course, as these quotes effectively remind us, philosophers of science have hardly agreed on just which *sort* of knowledge it is that characterizes understanding: whether understanding why P comes from knowing the cause of P, or from knowing that P can be subsumed under laws, or what have you, is still widely debated. But that understanding is a kind of knowledge is something that—so far at least—has rarely been doubted by those working in the field.²

That makes it all the more striking that virtually every major epistemologist who has thought seriously about the nature of understanding—including Catherine Elgin ([1996], [2004]), Linda Zagzebski ([2001]), and Jonathan Kvanvig ([2003])—has come to the conclusion that understanding is *not* a species of knowledge.³ They variously claim, for example, that understanding is immune to *Gettier problems* whereas knowledge is not, that understanding is *transparent* whereas knowledge is not, and that understanding is possible even in the absence of *truth*, whereas this is an impossibility for knowledge. All in all, they argue, understanding and knowledge are simply pulling in too many different directions for the former to be thought of as a species of the latter.

In this paper I will evaluate the case against taking understanding to be a species of knowledge by focusing specifically on the accounts offered by Zagzebski and Kvanvig, and I will first argue against this growing consensus in epistemology that—just like knowledge—understanding requires truth, is not transparent, and can be *Gettiered*.⁴ In short, that understanding certainly seems to behave a lot like a species of knowledge. Later in the paper I will then ask whether there are independent reasons for supposing that understanding is not a species of knowledge, and I will end by considering what, on balance, we should say about this subject.

1 Zagzebski’s account

According to Zagzebski ([2001]), one of the most striking differences between knowledge and understanding is that while understanding is *transparent*,

¹ Other philosophers of science who seem to assume, without need for argument, that understanding is a species of knowledge include Wesley Salmon ([1989], pp. 134–5), James Woodward ([2003], p. 179), and Philip Kitcher ([2002], *passim*).

² This is to set to one side issues raised by those who, like J. D. Trout ([2002]), claim that understanding is not really an epistemic good at all but rather merely a “subjective feeling.” I will not be dealing with that sort of understanding skepticism in this article.

³ And, for that matter, not just epistemologists. According to John Haugeland: “Understanding is always “of” something—objects, in a broad sense—but this of-ness is not the same as that of beliefs or desires. Thus, understanding is not the same as knowledge, a special kind of knowledge, or even a complex structure or totality of knowledge” (Haugeland [1998], p. 1).

⁴ Addressing Elgin’s main concern—namely, that not only is it the case that understanding can be had in the absence of truth but that in some instances understanding results precisely *from* idealizations that depart from the truth—will take us too far afield. Strevens [forthcoming] proposes a way of responding to this concern that is highly congenial to the account I develop here.

knowledge is not. On her view, understanding is fundamentally a matter of grasping how various pieces of information relate to one another; it is a matter of making connections among them, of seeing how they hang together. But if that is the case, she argues, then it is natural to think that the *object* of understanding—apparently, the connection that we “see”—must be internal. After all, how can one “grasp” or “see” something that isn’t right there, open to our mental view?

In this as well as other respects, however, Zagzebski claims that understanding differs from knowledge. As she writes:

Understanding has internalist conditions for success, whereas knowledge does not. Even when knowledge is defined as justified true belief and justification is construed internalistically, the truth condition for knowledge makes it fundamentally a concept whose application cannot be demonstrated from the inside. Understanding, in contrast, not only has internally accessible criteria, but it is a state that is constituted by a type of conscious transparency. It may be possible to know without knowing that one knows, but it is impossible to understand without understanding that one understands. ... [U]nderstanding is a state in which I am directly aware of the object of my understanding, and conscious transparency is a criterion for understanding. Those beleaguered by skeptical doubts therefore can be more confident of the trustworthiness of putative understanding states than virtually any other state. (Zagzebski [2001], pp. 246–7)

It is not hard to see what Zagzebski is driving at. For one thing, we can grant that there is some sort of connection that is immediately open to view when we take ourselves to understand; it seems to be a hallmark of having understanding that we can typically articulate (or explain) what it is in virtue of which we take ourselves to understand, for instance, and that kind of articulacy might be thought to require a direct apprehension of the thing understood. For another, it is plausible that there are some types of understanding that *do* allow for the kind of transparency Zagzebski has in mind. Our understanding of concepts might fall into this category.

But to grant that some components of understanding are transparent in this way, or even that some types of understanding might allow for this kind of transparency, is a long way from granting that understanding in general allows for the kind of transparency described by Zagzebski. In particular, it is a long way from granting that our understanding of natural phenomena—arguably the paradigm case of understanding—allows for this kind of complete transparency.

A simple example helps to illustrate the point. Suppose that you open the refrigerator and notice that the light has gone off, so you put your hand inside and feel around, finding it warm to the touch. While looking around for an explanation eventually you notice something unusual: the cord has been

unplugged. You now take yourself to understand why the refrigerator stopped working, but it should be obvious that you might be mistaken. It might have broken down as the result of a short circuit, and the cord might have been unplugged for fear of fire. In that case, despite your sense that you understand why the refrigerator stopped working, in fact you don't understand. You've mischaracterized how things stand in the world.

The basic problem with Zagzebski's account is therefore that our understanding of natural phenomena seems conspicuously *factive*—what we are trying to grasp is how things actually stand in the world, and there is no reason to think that how things stand in the world is consciously transparent to us. To this it might be objected that even when we miss out on understanding in this way due to some failure of fit with the world, there is nonetheless a desirable form of understanding that we *can* nonetheless lay claim to. In the case of the refrigerator, for instance, even though you fail to understand how things stand with respect to this particular situation, you still seem to understand something else of value: roughly, that refrigerators depend on electricity, and in the absence of a source for electricity, they stop working. But generalizations of this sort are themselves substantive claims about how things stand in the world; they claim to describe real physical dependencies. And, again, there is no reason to think that we have any kind of transparent access to whether or not such dependencies actually obtain.⁵

2 Kvanvig's account

Kvanvig's ([2003]) account of understanding is therefore an improvement on Zagzebski's, in this respect at least, because he insists on the factive character of understanding. For Kvanvig, what this means is that in order for a mental state to count as a state of understanding, the beliefs that constitute the state must be true. As he sees things, however, what is important and distinctive about understanding is entirely a matter of what one makes *of* one's true beliefs, once one has them. Specifically, Kvanvig insists along with Zagzebski that understanding is fundamentally a matter of grasping or seeing connections among one's beliefs. But on his view what this entails is that, so long as the beliefs happen to be true, facts about how one's beliefs were acquired—in Kvanvig's terms, external facts about the *etiology* of one's

⁵ Zagzebski has more to say here, about how understanding is internal to (what she calls) practices, and about how practices allow us to identify and correct mistakes along these lines. I will not address these further proposals here, beyond noting that there are some practices—astrology, voodoo, etc.—that would confirm very strange dependencies indeed, and that from the inside these might well have all of the same conscious marks as more conventional practices (such as, to cut things thickly, contemporary science).

beliefs—are irrelevant when it comes to assessing whether or not one understands.

Kvanvig argues for this claim by means of the following example (Kvanvig [2003], pp. 197–8). Suppose you pick up a textbook on Native American History and read through a chapter documenting the Comanche dominance of the southern plains, until eventually you seem genuinely to understand *why* the Comanches dominated the southern plains. But suppose as well that while the book you happened to pick up is accurate, most other books on this topic are full of errors. If you had picked up one of these other books instead (and we can imagine that they are all within easy reach!), your beliefs about the Comanches would have been almost entirely false.

What then should we say about the status of your Comanche beliefs? Do they amount to knowledge? Many would say that they fail to amount to knowledge because the fact that your beliefs about the Comanches hit the truth seems to be a mere accident. Unfortunately for you, you have stumbled into a Gettier environment, where a bit of good epistemic luck (happening by chance upon accurate information) has canceled out a bit of bad epistemic luck (being in an environment full of misinformation).

But now, as Kvanvig notes, we still have a further crucial question to consider: Once we see how easily your beliefs might have been false, should we take the same negative view toward your (alleged) understanding? Should we say that it fails to amount to a case of *genuine* understanding? According to Kvanvig, the answer is No. The specific story about how one's beliefs came to be true—whether as a result of accident or luck or some other means—has no impact on their ability to contribute to understanding. Once the truth condition has been met, understanding is a purely internal matter. He sums up these claims as follows:

Understanding does not advert to the etiological aspects which can be crucial for knowledge. What is distinctive about understanding, once we have satisfied the truth requirement, is internal to cognition. It is the internal seeing or appreciating of explanatory relationships in a body of information which is crucial to understanding. When we think about knowledge, however, our focus turns elsewhere immediately, if we have learned our lessons from the Gettier literature: we think about the possibility of fortuitousness, of accidentality, of being right only by chance. We focus, that is, on what kind of further external connections there are between mind and world, beyond the fit required for the belief to be true.

The basic idea here is that, though knowledge is incompatible with a certain kind of epistemic luck, understanding is not. Upon learning of the disturbed etiology of beliefs about the Comanches, as in the case imagined here, we might say that the person has true beliefs or even true justified beliefs, but no knowledge, if we have heeded our lessons from Gettier. We would not, at least we should not, say that because of

these factors, she is lucky to have the knowledge she has, for knowledge rules out this kind of luck. But we needn't say the same thing about the claim of understanding. If the etiology were as imagined, one would be lucky to have any understanding at all of the Comanche dominance of the southern plains. So such understanding would count as understanding not undermined by the kind of luck in question. (Kvanvig [2003], pp. 198–9)

Understanding is therefore a paradigm instance of a luck-proof cognitive state. And since knowledge is essentially vulnerable to luck, we can conclude that understanding is not a species of knowledge.

In the remainder of this chapter I will try to show that this general way of thinking about understanding—what we might call “understanding internalism”—is misguided. Not only is understanding susceptible to luck, I will argue, but etiology clearly matters when it comes to understanding. This will then leave us with the following question: What should we make of Kvanvig's Comanche case? Specifically, has he really shown that there can be understanding in the absence of knowledge? I will argue that the Comanche case shows none of these things. As we will see, it matters how we unpack the details of this kind of case. But on any way of filling out the details knowledge and understanding seem to sway together.

3 Two problems

Beginning with the claim about luck, there are at least two different genres of counterexamples to understanding internalism: we might call the first *veridical hallucination problems* and the second *bad environment problems*. Further ways in which understanding is vulnerable to luck will emerge in later sections, but these two types of cases should be enough to introduce some basic difficulties with the view.

To illustrate the first sort of problem, 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 then 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, his hallucination exactly matches the events that are actually unfolding in front of him.

When we add the pieces up, we seem to have all the elements for understanding that Kvanvig requires. There is the true belief that the vase fell from the table, as well as the true belief that the stand was bumped by the dog. There is, in addition, the genuine dependence that (presumably) holds between the two: that the vase fell *because* the table was bumped by the dog. But still, we're tempted to say, 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 seems to be straightforward (not to mention, for those acquainted with the ongoing debate in epistemology about the nature of knowledge, familiar): namely, that too easily might he have misidentified the cause of the fall. If the drug had caused him to hallucinate something else, had it portrayed his cat (or a squirrel, or a Frisbee, and so on) bumping into the table instead of his dog, then he would have believed that the vase fell because the table was struck by the cat (or any of the other things we care to invent).⁶

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.⁷ 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. To clarify why, suppose that Becky lingers for a while and watches as the blacksmith performs his chestnut trick in the usual way. In each case, the blacksmith times his swing

⁶ If one is worried that Albert's belief about the dog might not count as true because in virtue of the hallucination it isn't *de re* enough—in other words, not really *about* the dog—we can easily change the example. Thus suppose that Albert's roommate Stan mistakenly thinks that he is responsible for the broken vase (it was a rough night last night, and the details are vague in Stan's mind). In order to avoid responsibility, he makes up an alternative story about the dog—a story I then accept on his testimony—which turns out, by merest chance, to be exactly right. (Thanks to John Turri for this version of the case.)

⁷ The case is adapted from one of J. L. Mackie's well-known examples (Mackie [1980], pp. 43–4).

perfectly so that the hammer impacts the chestnut just as the nut is about to explode. The first time this happens, we can imagine Becky thinking to herself: “That chestnut just shattered because of the blow from the hammer.” And so on for the second time, and the third time, etc. Indeed, we can imagine that all the blacksmiths in the area perform the same trick, so that time after time the chestnut explodes precisely upon impact by the hammer.

If we are reluctant to ascribe understanding to Becky on the lone occasion she gets it right, however, I suggest that this is for essentially the same reason that we are reluctant to ascribe knowledge to someone driving through fake barn country on the lone occasion when they happen to spot a real barn.⁸ Just as someone seems to lack barn knowledge on the lone occasion when (while driving for several hours through fake barn country and repeatedly pausing to remark “That’s a mighty fine barn!”) he comes across the only real barn in the area, so too Becky seems not to grasp why the chestnut exploded when she might so easily have been mistaken. Knowledge, as Nozick once observed ([1981], p. 213), should be made of sterner stuff than that—not so easily should the knower be capable of mistake. But now, in light of the previous example, the same seems to hold true of understanding. If Becky would have failed to identify the genuine cause of the shattering 99 out of a 100 times in the blacksmith shop, the natural temptation is to say that, on the lone occasion when she gets things right, she fails to understand. Like knowledge, understanding too should be made of sterner stuff.

4 Comanche cases

I think these examples show how etiology can matter a great deal to understanding, and thus that there is good reason to suspect that it is not a purely internal affair. It is important to see, however, that Kvanvig can accept all this and still have his fundamental claim go through: namely, that understanding is not a species of knowledge. All he needs to show in order to establish *that*, after all, is that there are cases where understanding and knowledge can come apart—specifically, that we can have cases of understanding that are not cases of knowledge. And Kvanvig claims to have produced a case of this sort: namely, where someone understands all about the Comanches even though, because of the chancy way they came to believe the relevant information, none of the relevant beliefs amounts to knowledge.

Nothing special hangs on the fact that the example concerns the Comanches *per se*, of course, but for convenience we can hereafter refer to

⁸ Where by “fake barn country” I mean the kind of place famously described by Goldman in his ([1976]): namely, a place where the majority of barn-like structures in the area are in fact barn *facades*, set up to make travelers think that the area is more prosperous than it actually is.

Table 1. Kvanvig's Chart

Information Source	Information environment	
	Good	Bad
Good	Knowledge Understanding	No knowledge Understanding
Bad	No knowledge Understanding	No knowledge Understanding

the broader type to which this case belongs as a “Comanche-style case” and characterize it as follows. A Comanche-style case is one in which we form true beliefs on the basis of trusting some source, and either (a) the source is unreliable, or (b) the source is reliable, but in the current environment one might easily have chosen an unreliable source. Cases along the lines of (a) are easy enough to imagine: they are cases in which we trust a bad source (a crystal ball, a pathological liar, etc.) that just happens to get it right. Cases along the lines of (b) are also fairly straightforward: suppose by luck you happen across the only honest man in a room of pathological liars, and the man explains Einstein's theory of general relativity to you. If all goes well you will then understand the theory, but you will (so Kvanvig suggests) fail to know it: too easily might you have trusted some other source of information.

Since Kvanvig is not entirely clear about whether his Comanche example is supposed to be an instance of (a) or (b), to see how knowledge and understanding sway together throughout Comanche-style cases we will need to lay out the various possibilities a bit more carefully and consider them in turn.

The following chart is one way to make sense of the landscape. I will say more about what I mean by “sources of information” and “information environments” in a moment.

On this way of looking at things, a *source of information* is the primary notion, and we can distinguish between a good source and a bad source as follows. A good source of information with respect to a certain domain is a source that is more likely than not to provide accurate information about that domain.⁹ A bad source is one that is more *unlikely* than not to provide accurate information.¹⁰ For example, a good source of information about the past will be someone with a reliable memory, and a bad source of information about the past will be someone who speculates about the

⁹ Or, at any rate, that places above whatever the reliability threshold happens to be.

¹⁰ Or that places *below* whatever the reliability threshold happens to be.

past by using an Ouija board.¹¹ A good source of information about the (well-lit, normal) immediate environment will be someone with normal perceptual abilities; a bad source will be someone who has just been taking hallucinogenic drugs.¹²

The notion of an information environment, as I am understanding it, builds on the notion of a source of information. A good environment, information-wise, is an environment where one's sources of information—the sources of information within easy epistemic reach, as it were—are by and large good (i.e. reliable). A bad environment, information-wise, is an environment where the sources of information within one's easy epistemic reach are *not* by and large good. Suppose you are at an intersection with a dozen people and that you need to ask for directions. A bad information environment will be one in which only one person is reliable about the directions in question, while the other eleven will send you off on a wild-goose chase.¹³ A good information environment would be one in which virtually all of the people at the intersection are good sources of information, in the sense that virtually all will tend to get things right with respect to the domain in question.

With this in mind, Kvanvig's claim is that knowledge and understanding come apart in three of the four quadrants. On his view, whenever you have an element of luck involved, either from a bad information source or from a bad information environment, although knowing on the basis of this information is ruled out, understanding is still entirely possible.

As this last comment suggests, when we say that for Kvanvig knowledge and understanding come apart in these areas, we need to be careful about interpreting this in the right way. What he means, presumably, is something along the following lines: that the potential for knowledge or that the possibility for knowledge can come apart from the potential or possibility for understanding. For example, just because one bases one's (true) belief on a good information source in a good information environment, it does not follow—on anyone's view—that the belief amounts to knowledge. You might have defeaters for the belief that you are overlooking, for instance, or (a misleading) reason to think that the source, in fact, is not reliable. We should therefore read Kvanvig as claiming that believing on the basis of an unreliable source, or on the basis of a reliable source but in a bad information environment, permits or allows for understanding, but does not permit or allow for knowledge, in three of the four quadrants described above.

¹¹ I am cutting things quite thickly here: someone could have reliable beliefs about their college years, for example, and unreliable beliefs about their childhood.

¹² On this view, things like books will count as derivatively reliable.

¹³ There are complications here about people's *willingness* to share their accurate information. If this is your worry, imagine that you are choosing among a pile of 12 maps instead, only one of which contains accurate information.

5 Unreliable sources of information

In this Section I will argue that Kvanvig is mistaken about the bottom half of the chart: just as basing one's beliefs on a bad source does not allow for knowledge, so too it does not allow for understanding. In other words, the two states in fact do not come apart. Although I will try to show that this holds for any Comanche-style case, for concreteness it will help to work with a specific example, and Kvanvig's own (Comanche!) case will do as well as any.

The first thing to bear in mind when evaluating these possibilities, especially when understanding is ultimately at issue, is that there are two different ways to think about the nature of the information—in this case, the information contained in the Comanche book—we are considering. On the one hand, we can imagine that the book consists, at least in part, of *explanations*—or better (in order to emphasize their potential fallibility) of *explanatory stories*. On the other hand, we can imagine that the book does not contain explanatory stories at all. Instead it is the kind of book that just offers descriptive information and is silent about how this information should be interpreted or about how one bit of information illuminates the rest.

Case I: Explanatory Story from a Bad Source. Taking the first way first, suppose that Kvanvig's Comanche book offers explicit explanatory stories. Perhaps, for example, the book claims (and here, owing to ignorance, I am just making things up) that the Comanches dominated the southern plains of North America because their economy was based on a barter system, and that economies based on barter systems were more likely to produce members who were competitive and cunning, traits which in turn were enormously conducive to success in battle.

Suppose that all this is true (i.e. that the alleged dependencies are genuine), that I accept this explanatory story, and that I take myself to understand why the Comanches dominated the southern plains on the basis of “seeing” or “grasping” how the various elements described in the story depend on one another. In keeping with the bottom half of the chart, what we now need to *add* is that the author of the textbook came to accept this explanatory story on the basis of extremely shoddy research. For instance, we can imagine that the author concluded that economies based on barter systems were more likely to be competitive and cunning on the basis of only one sample, and without controlling for the influence of other factors. In these circumstances, then even if it turns out that the story is entirely correct, and that there is a genuine dependence between barter systems and competitiveness on the one hand and between competitiveness and success in battle on the other, then presumably the author of the theory doesn't genuinely understand why

the Comanches dominated the southern plains. And if the author doesn't understand this, it is very hard to see how the *reader* of the textbook could understand in turn.¹⁴

Case II: Descriptive Information from a Bad Source. For completeness we also need to consider a Comanche textbook that *doesn't* traffic in explanations. In other words, to consider a book that is entirely descriptive, one that lays out the facts and leaves the insightful reader free to construct explanations on his or her own. Since the relevant accomplishment—the piecing together, as it were—would then appear to be entirely internal, perhaps the facts about how the descriptive information was acquired would then become irrelevant to the possibility of understanding.

But that does not follow. Suppose that the entire textbook is the product of baseless reporting, written by a 19th century version of the *New York Times's* Jayson Blair—a negligent writer who could not be bothered to do the research and simply invented his claims about the Comanches out of the blue. Even if all the stars line up in his favor and his claims turn out to be entirely correct, it is hardly the case that the textbook reader who pieces things together and develops an account of Comanche dominance would really come to understand *why* the Comanches dominated the Southern Plains.

Of course, in saying this we do not have to deny that there is a *kind* of internal accomplishment involved in piecing together the facts and constructing an explanatory story in this way. All we have to deny is that the accomplishment is really one of understanding why the Comanches dominated the Southern Plains. For that accomplishment, apparently, a stronger (alternatively: less accidental) connection to the Comanches is required.¹⁵

¹⁴ Indeed, if understanding could be acquired in such a slipshod way (by consulting only one sample, without controlling for the influence of other factors, etc.), then it is hard to see why we would have the need for the sociologists, economists, and historians—with their established methodologies and established canons for constructing and evaluating explanations—that we evidently have.

¹⁵ There is perhaps an even more straightforward way to illustrate the problem with thinking that understanding can flow from poor sources of this kind, one that can be made to fit either Case I or Case II examples. Thus consider a *swampbook* version of Kvanvig's Comanche book, that is, an exact replica of a normal book brought into being by a stray lightning bolt and a patch of swampy gunk. We can imagine either that the swampbook contains explicit explanatory stories as in Case I, or we can imagine a swampbook that just presented the facts, as in Case II (all entirely accurate, of course). On either way of thinking about the swampbook, however, we would be hard pressed to say that the readers of swampbook understood anything about the actual Comanches. It is much more plausible to say instead that what they understood was something about the logic of the story, or perhaps about the patterns of dependency that held among some possible tribe. To understand the actual Comanches, again, we seem to need a firmer connection to the Comanches themselves. (If one thinks the failure of understanding here depends on failure of reference, then it is easy enough to imagine that one has originally learned about the Comanches in a more orthodox way.)

6 The upper-right quadrant

Assuming that it is likewise impossible to acquire knowledge from such poor sources, so far we have the following results. First, that the possibility of knowledge and the possibility of understanding sway together on the upper left quadrant, where a good source meets a good environment. Second, where we have a bad—i.e. unreliable—source (a terrible methodology, a Jayson Blair, a swampbook, etc.), just as knowledge on the basis of such information seems ruled out, so too does understanding. So on all of the ways of filling in the backstory that we have considered so far, Kvanvig's point fails to hold: knowledge and understanding do not come apart.

Arguably, however, the best case for Kvanvig remains to be considered: namely, the upper right quadrant, or the case where our beliefs are based on accurate information from a good source, but in the midst of a bad information environment. Here again, we can fill in the details of this quadrant in one of the two ways: we can either suppose that the information from this good source traffics in explicit explanatory stories, or we can suppose that it only provides descriptive information. It is simpler to evaluate Kvanvig's claim, however, if we suppose that the information is purely descriptive.¹⁶

Suppose then that the accurate information contained in the textbook comes from a reliable source: someone who is not only a Comanche expert, but whose methods are impeccable, etc. Moreover, all of the other information sources about the Comanches within easy reach are unreliable: if they hit upon the truth, it is only by accident. For the sake of argument, suppose we agree that coming to understand on the basis of such information (i.e. the information from the good source) is possible. What I want to argue now is that, pace Kvanvig, there is a perfectly legitimate sense in which such a person can come to *know* on the basis of this information as well. And if that can be established (and a few more details sorted out), then Kvanvig will not have shown that understanding and knowledge can come apart after all.

The key point to appreciate here is that Kvanvig overstates his case in claiming, without appropriate qualification, that knowledge is incompatible with luck. In reality, knowledge is compatible with some forms of luck and incompatible with others (See e.g. Pritchard [2005], ch. 5). For example, although believing on the basis of information from a bad source never seems to allow for knowledge, even when by luck the bad source happens to get things right, believing on the basis of a good source in a bad information environment *does* seem to allow for knowledge. In suggesting in a blanket way that knowledge is incompatible with luck Kvanvig therefore paves over significant differences.

¹⁶ Moreover, although it is hard to say exactly, this seems to be the model he has in mind.

John Hawthorne has recently proposed a case that illustrates the compatibility nicely. He writes,

I give six children six books and ask them each to pick one of the books at random. All but one contains misinformation about the capital of Austria. I ask the children to look up what the capital of Austria is and commit the answer to memory. One child learns ‘Belgrade,’ another ‘Lisbon,’ another ‘Vienna,’ and so on. I ask an onlooker who has witnessed the whole sequence of events (or someone to whom the sequence of events has been described) ‘Which one of the schoolchildren knows what the capital of Austria is?’ or ‘How many of the children know what the capital of Austria is?’ It is my experience that those presented with this kind of case will answer, not by saying ‘None of them,’ but by selecting the child whose book read ‘Vienna’—even though that child was only given the correct answer by luck. (Note in this connection that if I make a five dollar bet on a certain child knowing that Vienna is the capital of Austria, you will pay up as soon as you are convinced that the child believes the capital to be Vienna. You will not inquire further about how the child came by that information—whether by dumb luck or from an informant that normally lies—even if you have reason to suspect such an unreliable source.) (Hawthorne [2003], pp. 68–9)

Hawthorne’s judgment about the case seems very plausible; it seems right to credit the child with knowledge. Hawthorne does not go on to discuss *why* this seems right, but a few things can be said to try to clarify the case.

First, when (as third party evaluators) we have reason to believe that the source of the information is good—here, that the textbook came from a reliable author, etc.—we tend to focus on the sense in which the belief is not lucky: it was no matter of luck, we think, that the textbook author identified Vienna as the capital of Austria, even if it was a matter of luck that this particular textbook ended up in the student’s hands.

But now suppose we learn that the author’s identification of Vienna was *itself* the product of luck. Perhaps, for example, the information in all of the textbooks, including the one the lucky boy received, was produced entirely with the help of a Ouija board. As a result, one book contained the information that the capital of Austria was “Oog,” another “Tuyy,” and another “Vienna” (or we could substitute the names of real capitals if you like; the point remains the same). If we fill in the details of Hawthorne’s case in this way, our inclination to judge that the boy *knows* the capital of Austria drops considerably. He has a true belief about the capital of Austria, naturally, but we draw the line: you simply can’t gain knowledge from an Ouija board, regardless of whether the board happens to get things right.

My intention in raising Hawthorne’s case is not, needless to say, to enter into an extended account of the various ways in which knowledge is and is not compatible with luck. The intention instead is simply to point out that, if

we are tempted to think that someone can come to *understand* various things about the Comanches on the basis of good information source in a bad information environment, then in cases of this sort there is also a strong tendency to say that the person *knows* these things about the Comanches as well.

Naturally, there are still other ways to imagine the case—cases where we make the luckiness of the belief seem more salient—in which we would be inclined to deny that basing one's belief on a good source might allow for knowledge. If we learn that before picking up the good-source Comanche book you picked up 20 other books on different subjects (dinosaurs, space exploration, etc.), all of which were tissues of lies, and that you accepted everything these books had to say, then when you come across the good-source Comanche book, I take it that our inclination to say that the beliefs you form about the Comanches amount to knowledge again drops significantly.¹⁷

That point is no solace to Kvanvig, however, just so long as our judgments about the possibility of understanding *sway along with* our judgments about the possibility of knowledge. And they certainly seem to. If I pick up the same 20 books and I take myself to understand various things about these subjects on the basis of this information (in the first book about dinosaurs, in the second book about space exploration, etc.), then when I come across the lone accurate book about the Comanches, the temptation to say that I genuinely understand these things about the Comanches seems to drop again, even though all of the connections that I draw might be perfectly correct.

We can sum up these results with the following revised chart.

Table 2. Revised Chart

Information source	Information environment	
	Good	Bad
Good	Knowledge Understanding	Knowledge* Understanding*
Bad	No knowledge No understanding	No knowledge No understanding

¹⁷ And the same thing for Hawthorne's capitals case. If we feed the child 20 books filled with false information about world capitals, and he blithely accepts all of this, then there is a strong tendency to think that when he comes across the one book that says "Vienna is the capital of Austria" he fails to know it.

I have placed a star (“*”) next to the contents of the upper right quadrant because these circumstances have clearly emerged as a special case. According to the ground rules I laid out in Section 4—where a quadrant counted as knowledge-friendly or understanding-friendly just in case it was *possible* for knowledge or understanding to emerge under those conditions—in one sense we should simply remove the stars: if we accept the Hawthorne-style cases, then it *is* possible for knowledge as well as understanding to emerge under those conditions.

As we just saw, however, cases involving a good information source in a bad information environment quickly become complicated, and it would be misleading to suggest that (defeaters aside) knowledge and understanding are the norm here. Although this starred status makes the chart less theoretically elegant, in terms of the larger project the lack of elegance is in a way beside the point: for us, all that matters is that when knowledge and understanding *do* shift with particular circumstances, they shift together. And this last fact has been repeatedly established: where the connection with the truth has been strong enough to classify a mental state as a state of understanding, the beliefs that constitute the state have also been strong enough to count as knowledge. When, on the other hand, the connection has been too weak for understanding, it has likewise been too weak for knowledge.

7 So is understanding a species of knowledge?

So, is understanding a species of knowledge, then?¹⁸ From what we’ve seen so far there seems to be a *truth component* to understanding; if you identify the cause of the refrigerator’s shutting down as the pulled plug when in fact it was a short-circuit, then no matter how vivid or transparent the pulled-plug story might seem to you, you will fail to understand why the refrigerator shut down. There also seems to be an *anti-luck component*: hitting upon a true explanatory story by accident—as through a Jayson-Blair-style source, or as in a blacksmith’s shop in the midst of eccentric nut-smashing activity—rules out understanding. In all of these cases, a stronger connection to the truth is needed.

Although the parallels between knowledge and understanding therefore run quite deep, to this point we have failed to consider one further, and presumably crucial, component of knowledge that threatens to upset the comparison: namely, the *psychological component* of knowledge.

The psychological component of knowledge is generally taken to be *belief*, and by way of elaboration it is usually claimed that when we believe

¹⁸ I am particularly indebted to the anonymous reviewer for helpful feedback on the remainder of the paper, which led to several clarifications and (I hope) improvements.

something we take it to be the case, we hold it to be true.¹⁹ A slightly different way to characterize belief is in terms of the act of assent—to believe that something is so is to *assent* to the claim that things are so. On this view, just as denying some claim is like saying No to it, so too believing some claim is like saying Yes to it.²⁰

But now notice, if we think that knowledge is a species of belief, and that understanding is a species of knowledge, then understanding too will turn out to be a species of belief. And given what we've just claimed about belief, it would then follow that the psychological component of understanding should likewise be thought of along the lines of an act of assent (a "saying Yes" of some kind). It is at this point, however, that the parallels between knowledge and understanding start to look questionable.

To see the problem, suppose for a moment that the psychological component of understanding *could* be characterized as an act of assent. What would be the thing to which we were assenting? The most plausible answer here is a claim of the form *A because of B*. For example, if you are curious about why the tides exist, then (given the appropriate background beliefs) your curiosity will be satisfied when you learn that the tides exist *because of* the moon's gravitational pull. From a psychological point of view, the thing you will be assenting to—saying Yes to—will therefore presumably be a claim to the effect that the tides exist because of the moon's gravitational pull.

But there seems to be a significant problem here: specifically, an act of assent of this sort (or even, arguably, *several* acts of assent of this sort) seems too thin to capture what is going on when we take ourselves to understand. Suppose that one day I get into my 1991 Volkswagen and none of the gauges—the speedometer, tachometer, etc.—come to life; they're all dead. I drive the car to the garage and my mechanic tells me that the reason why the gauges are dead is because I have a bad ignition switch. I then seem to have excellent reason to assent to the claim that my gauges are dead because of a bad ignition switch, even though I fail to grasp how a bad ignition switch might lead to this result. In other words, even though I currently lack the epistemic wherewithal to grasp A's dependence on B, I nonetheless now have an excellent basis for assenting to the fact (or proposition) that A is because of B.²¹

¹⁹ See, for example, Eric Schwitzgebel: belief is "the attitude we have, roughly, whenever we take something to be the case or regard it as true" (Schwitzgebel [forthcoming], p. 1).

²⁰ Not that these attitudes are exhaustive—there are still withholdings to consider.

²¹ It might be objected, as the anonymous reviewer objects, that even though I do not understand why a bad ignition switch leads to dead gauges, I still understand something significant: namely, that if the switch had not been bad, the gauges would not have been bad. The difference between the two mental states would then end up being a difference in degree (or depth of understanding) rather than a difference between non-understanding and understanding. I am sympathetic

Examples from mathematics also bear this out. If you are my teacher and you tell me that a certain proof explains a theorem, then on the basis of your testimony I will assent to the fact that the theorem can be explained because of, or in virtue of, the things you cite in your proof. But it seems obvious that I may still fail to understand how the one explains the other. Grasping the way in which the theorem depends on the elements you cite in your proof is different from assenting to the claim *that* the dependency holds.

As this discussion suggests, when trying to offer an account of understanding the notion of *grasping* arises almost irresistibly. Moreover, when we grasp some claim we are apparently doing something significantly different from merely saying Yes to it. In reply, it might be said that this attitude of grasping is nothing significantly different from an act of assent: it just involves *more* assent. But it seems clear that one can pile up assents as high as you like without getting a grasping. In the mathematics case, for example, I can be told that the key to the proof lies in the fact that the first element is divisible by the second, and I can assent to this (again, based on the testimony of someone I trust). But I still might not understand the proof; I might fail to see or grasp how the truth of the theorem depends on the fact that the first element is divisible by the second. And it seems like this game could go on indefinitely. You might continue to spell out how the various dependencies are supposed to work. Moreover, based on your testimony I might assent to these claims at every step of the way. But none of these assents, in and of themselves, adds up to a genuine grasping on my part. A new kind of cognitive achievement seems to be needed.

Following James Woodward ([2003]), one promising way to spell out the requisite notion of the *kind* of achievement that is needed—hence the kind of grasping that is involved in understanding—is in terms of having an ability to answer “what-if-things-had-been-different?” questions.²² To have an ability to answer questions of this sort, Woodward argues, is to be able to anticipate the sort of change that would result in the thing we want to explain (the explanandum) if the factors cited as explanatory (the explanans) were different in various ways. Brian Skyrms ([1980], p. 11) expresses this same point in a more metaphorical way. According to Skyrms, grasping involves having an ability to anticipate that “wiggling” one variable will

to this objection, and I think my final proposal in Section 8 helps to illuminate what’s at issue. In essence, in order to assent to a dependency claim along the lines of A because of B in a way that really amounts to assent (as opposed to just a mouthing of the words “A because of B”), one must possess a distinctive ability or aptitude that goes beyond merely saying Yes to the claim. In other words, to assent meaningfully to a dependency claim of this kind arguably requires the very sort of counterfactual manipulative ability that I stress in Section 7, a claim that would further support the suggestion that some assents are richer than others.

²² Henk de Regt ([2004]) likewise suggests that understanding should be conceived of as a kind of ability (see also de Regt and Dieks [2005]).

characteristically lead to a “wiggling” of another variable. Moreover, if you lack the ability to see how things are connected in this way, however—specifically, if you lack the ability to see how a change in the value of one variable will lead to a change in the value of another variable—then although you might justifiably assent to the fact (based on testimony, etc.) that the two variables depend on one another, there remains an important sense in which you will not have grasped how they depend on one another.

There is one further reason to think of the psychological component of understanding in terms of the richer notion of grasping, rather than the thinner notion of assent: namely, such a shift would help to shed light on why the epistemic gain we experience when we understand is so universally valued. For notice, if understanding involves the ability to answer “what-if-things-had-been-different?” questions of the sort described by Woodward, then this ability will doubtless bring with it the promise of tremendous control over one’s environment. If there are various ways in which, from your point of view, the world might be, and if you are able to identify the features that determine (in some sense) the various possible states of the thing you want to explain, then you will presumably be in a very enviable position with respect to influencing the future course of the world. Specifically, you will be in an excellent position to bring about the things you desire and to avoid the things you care to avoid. That such an ability would be highly valued and sought after hardly needs argument.

8 A false choice

Despite the strong case for thinking of the psychological component of understanding along the “rich” lines sketched above, what I want to suggest in closing is that the choice between thinking of understanding on the model of belief/assent and thinking of it on the model of grasping is not a forced one. More specifically, I want to suggest that, while belief *can* be quite a thin psychological state, some forms of belief are thicker than others. On this way of looking at things, belief emerges as a kind of umbrella category for “ways of saying Yes” to a claim, and within belief there are simple acts of assent and then there are also acts of assent that, as it were, combine an element of grasping with the act of assent.

Perhaps the most powerful reason to accept that some beliefs are thicker than others—and in particular that it makes sense to think of acts of grasping that combine an element of assent—is that we already *have* a nice model of how this works: namely, *a priori* knowledge. *A priori* knowledge, after all, is not distinguished merely by its content but also—perhaps primarily—by *one’s attitude toward* this content. My knowledge that $2 + 2 = 4$, for example, does not qualify as an instance of *a priori* knowledge simply in virtue of the fact that $2 + 2 = 4$ is a necessary truth. Thus, if I were to assent

to this claim simply on the basis of accepting the testimony of a reliable authority, I would thereby know it, but I would not know it *a priori*. What *makes* it an instance of *a priori* knowledge for me (assuming it is) instead depends on the presence of some other psychological attitude that I bear toward the content in addition to (or in some way mingled with) the attitude of assent. What exactly this other attitude amounts to is a topic I will not explore here (thankfully!), but in light of what we have said so far it worth pointing out that in this respect as well talk of “grasping” the truth or “seeing” that things could not have been otherwise seems to emerge almost unavoidably.²³

If we are willing to count *a priori* knowledge as a species of knowledge, at any rate, as we manifestly are, then the reservations we raised earlier about the inherent “thinness” of the psychological state of belief seem misguided. Although belief can be quite thin—can simply amount to a simple saying Yes to a claim—it can also have additional layers of richness. Moreover, so long as we are willing to accept that one of these layers contains an element of grasping, then on balance there seems to be no compelling reason to deny that understanding is a species of knowledge and, in light of our findings about how etiology matters to understanding, plenty of good reason to accept that it is. More directly, on balance we have good reason to think that understanding is a species of knowledge.

Acknowledgements

In addition to the very helpful anonymous reviewer for the journal, I would like to thank Michael DePaul, Fritz Warfield, Robert Audi, Catherine Elgin, Jaegwon Kim, Mark Silcox, Ernest Sosa, and John Turri for helpful comments and feedback on earlier versions of this paper.

Department of Philosophy
University of Montana
Missoula, MT 59812
USA
stephen.grimm@umontana.edu

²³ Roderick Chisholm, for example, begins one of his discussions of the *a priori* with the following passage: “There are propositions that are necessarily true and such that, once one understands them, one *sees* that they are true. Such propositions have traditionally been called *a priori*” (Chisholm [1989], p. 26).

References

- Achinstein, P. [1983]: *The Nature of Explanation*, New York: Oxford University Press.
- Chisholm, R. [1989]: *Theory of Knowledge, Third edition*, Englewood Cliffs: Prentice Hall.
- De Regt, H. [2004]: 'Discussion Note: Making Sense of Understanding', *Philosophy of Science*, **71**, pp. 98–109.
- De Regt, H. and Dennis D. [2005]: 'A Contextual Approach to Scientific Understanding', *Synthese*, **144**, pp. 137–70.
- Elgin, C. [1996]: *Considered Judgment*, Princeton: Princeton University Press.
- Elgin, C. [2004]: 'True Enough', *Philosophical Issues*, **14**, pp. 113–31.
- Goldman, A. [1976]: 'Discrimination and Perceptual Knowledge', *Journal of Philosophy*, **73**, pp. 771–91.
- Haugeland, J. [1998]: *Having Thought*, Cambridge: Harvard University Press.
- Hawthorne, J. [2003]: *Knowledge and Lotteries*, New York: Oxford University Press.
- Kitcher, P. [2002]: 'Scientific Knowledge', in P. Moser (ed.), *The Oxford Handbook of Epistemology*, New York: Oxford University Press.
- Kvanvig, J. [2003]: *The Value of Knowledge and the Pursuit of Understanding*, New York: Cambridge University Press.
- Lipton, P. [2004]: *Inference to the Best Explanation, Second edition*, New York: Routledge.
- Mackie, J. L. [1980]: *The Cement of the Universe*, Oxford: Clarendon Press.
- Newton-Smith, W. H. [2000]: 'Explanation', in W. H. Newton-Smith (ed.), *A Companion to the Philosophy of Science*, Oxford: Blackwell.
- Nozick, R. [1981]: *Philosophical Explanations*, Cambridge: Harvard University Press.
- Pritchard, D. [2005]: *Epistemic Luck*, New York: Oxford University Press.
- Salmon, W. [1989]: 'Four Decades of Scientific Explanation', in P. Kitcher and W. Salmon (eds.), *Minnesota Studies in the Philosophy of Science*, vol. 13, Minneapolis: University of Minnesota Press.
- Schwitzgebel, E. [forthcoming]: 'Belief', in E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, <plato.stanford.edu>.
- Skyrms, B. [1980]: *Causal Necessity*, New Haven: Yale University Press.
- Strevens, M. [forthcoming]: 'Why Explanations Lie: Idealization in Explanation.'
- Trout, J. D. [2002]: 'Scientific Explanation and the Sense of Understanding', *Philosophy of Science*, **69**, pp. 212–33.
- Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*, New York: Oxford University Press.
- Zagzebski, L. [2001]: 'Recovering Understanding,' in M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, New York: Oxford University Press.