

Environmental Luck and the Structure of Understanding

Forthcoming in *Episteme*

1. Environmental Luck

Conventional wisdom holds that there is no lucky knowledge: if it is a matter of luck, in some relevant sense, that one's belief that *p* is true, then one does not know that *p*. One sense of luck that is generally recognized to be incompatible with knowledge is *environmental luck*¹. This kind of luck is exemplified by Carl Ginet's "barn façade" case (see Goldman, 1976), in which one forms a belief in a way that is typically knowledge-conducive, but where a significant part of the explanation as to why one's belief is true in a particular instance is due to the nature of one's environment, such that one could have easily ended up with a false belief if things had been slightly different. Discussions of environmental luck tell us something important about knowledge, namely that knowledge requires a certain robustness of belief: if one forms a belief in an environment in which one could have easily gone wrong, then one does not have knowledge.

While knowledge has traditionally been the primary interest of epistemologists, *understanding* has recently been receiving significant attention. While there is as of yet little consensus regarding a theory of understanding, one way we can work towards developing such a theory is by considering whether the kinds of factors that are important for determining whether one knows are also important for determining whether one understands. The question I want to

¹ Sosa (see Sosa, 2001) is a notable exception here, demarcating two different kinds of knowledge, one which is compatible with environmental luck and one which is not.

address here is the following: is environmental luck incompatible with understanding in the way that is for knowledge? I am going to answer that it is.

Epistemologists seem to be largely split on this issue: for example, Kvanvig (2003), Pritchard et al., (2010), Morris (2012), and Hills (2015), amongst others, argue that environmental luck is compatible with understanding, while Sliwa (2015), Riaz (2015), and Kelp (2014), amongst others, argue that it is incompatible. The split is not unprincipled: the majority of those who hold that one cannot have lucky understanding tend to adhere to a form of *reductionism*, which states that understanding is a species of knowledge, whereas many of those who hold that one can have lucky understanding adhere to *nonreductionism*, which states that it is not. We can see why views about the compatibility of luck and understanding might follow commitments to reductionism: if understanding is a species of knowledge, and environmental luck prevents one from knowing, then it stands to reason that it will also prevent one from understanding. If one is a nonreductionist, however, the relationship between luck and understanding is less clear, and it becomes an open question as to whether part of what makes understanding a unique epistemic state is that it is compatible with at least some forms of luck.

Here I will remain agnostic on the question of reductionism: I will argue that, regardless of whether one holds that understanding is a species of knowledge, understanding is incompatible with environmental luck. My argument has three parts. First, we need to determine how we evaluate whether one has understanding, which requires determining what I will call understanding's *evaluative object*. I argue that as the evaluative object of (at least a traditional conception of) knowledge is a belief in a proposition², the evaluative object of understanding is a

² Throughout I will have a "traditional," i.e. justified-true-belief plus luck-immunization-factor conception of knowledge in mind. It is this conception, I take it, which is at the heart of most discussions of epistemic luck. I will not here have anything to say about "knowledge-first" conceptions, like those endorsed by Williamson (2001).

mental representation of a relational structure. Next, I show that arguments that environmental luck is incompatible with understanding miss the mark by considering cases in which one has a belief in a proposition is lucky to be true, instead of ones in which one's mental representation of a relational structure is lucky to obtain. I agree, then, with those who argue that one can have understanding when one's beliefs are environmentally lucky to be true, but that this compatibility is not relevant when considering the question of whether one can have environmentally lucky understanding. I then present what I take to be a properly constructed case which shows the incompatibility of environmental luck with understanding.

2. The Evaluative Object of Understanding

To determine whether understanding is compatible with environmental luck we should first consider the way in which it is incompatible with knowledge. The paradigmatic case that demonstrates the effects of environmental luck on knowledge is the aforementioned barn façade case: while driving through the countryside Henry sees what appears to be a barn and forms the true belief that there is a barn. In normal circumstances – that is to say, circumstances in which one is driving through a normal countryside populated by normal farmers – Henry's belief would qualify as knowledge. However, Henry is, unbeknownst to him, driving through barn façade country, a place that is populated by set designers and epistemologists who have erected a large number of barn façades, ones that are indistinguishable from the real thing (at least from the road). Henry, then, happens to form a true belief about the one real barn for miles around, although he could have easily formed a false belief about any one of the numerous façades instead. In this situation, according to the orthodoxy, Henry fails to know that there is a barn: after all, Henry just lucked out by forming a true belief about the one real barn he happened to

lay his eyes on. One lesson that we can draw from this case is that in determining whether one has knowledge that p , one factor that we need to take into account is the relationship between one's true belief that p and the environment in which p obtains, namely whether one's environment makes it such that one's belief could have easily been false.

As the barn façade case shows, in determining the effects of environmental luck on knowledge we are concerned with the conditions under which one's belief in a proposition is made true. This is because, at least according to a traditional conception of knowledge, a belief in a proposition is the *evaluative object* of knowledge: knowledge is a special kind of belief, and when we are determining whether one knows we evaluate one's belief to see if it meets the right kinds of conditions for it to qualify as being known³. We can now generalize from this discussion to consider the effects of environmental luck on other epistemic states. In general terms, luck in epistemology is a phenomenon concerning one's mental representation of the world, the way that the world actually is, and the way things could have been. To evaluate the effects of luck on any other kind of epistemic state, then, we need to consider the relationship between the evaluative object of that state and the conditions under which the contents of that state obtains and could have failed to obtain⁴. In order to determine whether luck is compatible with understanding, then, we first need to ask: what is the evaluative object of understanding?

³ One way of thinking about why one cannot have environmentally lucky understanding is that environmental lucky beliefs are not *safe*, and that knowledge requires safe belief. Roughly, a belief is safe just in case in most nearby possible worlds (relevantly construed) in which one forms one's belief in the same way as in the actual world, one's belief is true. When one's belief is lucky, though, it is true in the actual world, but generally fails to be true in the relevant nearby possible worlds, and thus fails to be safe. There is more to be said about a precise construal of a safety condition on knowledge, but I take this general conception to be sufficient for my purposes here.

⁴ One way of asking this question is again to generalize from questions concerning safety and knowledge to safety as it might pertain to other epistemic states, where safety is determined in general in terms of whether the evaluative object of the epistemic state obtains in the actual world as well as in nearby possible worlds (relevantly construed) in which the relevant representation was formed in the same way.

To answer this question I should first say something about the nature of understanding. Understanding is sometimes taken to be a more demanding epistemic relationship than knowledge: in general it seems like acquiring understanding is more difficult than acquiring knowledge, and while I may know a lot, I understand many fewer things than I know. It has been argued that understanding is epistemically valuable (see Pritchard, Millar, and Haddock (2010)), and that it is important in determining expertise (see Wilkenfeld et al. (2016)). Understanding can also take a number of different forms: we can, for example, *understand-why* something is the case⁵, *understand-how* to do something⁶, *understand-that* something is the case⁷, or *understand some object or phenomenon*⁸.

Despite the differences in views on and types of understanding, it is widely accepted that understanding requires *grasping*. What “grasping” consists in, exactly, is again a matter of debate. Some have argued that to grasp one must possess some cognitive abilities such that one has “cognitive control” over some relevant information (Hills, 2015); others argue that grasping involves being able to provide the right kinds of explanations of a given phenomenon in different circumstances (Khalifa, 2013); while others still prefer to use the term in an intuitive and largely metaphorical sense, such that one “sees” or “draws connections” between various bits of information. In general, when I understand something I grasp how things fit together, how causes are related to effects, how reasons support truths, and so forth.

These discussions refer to grasping as a *process* or set of processes. What, though, do we have a grasp *of* when we say that we grasp something? We can again look to recent discussions on understanding for possible answers: Kvanvig (2003: 3) argues that understanding is

⁵ See Pritchard, Millar, and Haddock (2010), Hills (2015), Khalifa (2013), and Wilkenfeld et al. (2016).

⁶ See Zagzebski (2008).

⁷ See Bourget (forthcoming).

⁸ See Kvanvig (2003), Wilkenfeld (2013), and Kelp (2015).

concerned with “the ways in which pieces of information are connected with each other”; Wilkenfeld (2013: 1000) argues that understanding is a mental phenomenon concerned with one’s representation of the phenomenon that one understands; Grimm (2006:12) argues that understanding involves grasping a “structure”; and Hills (2015: 3) argues that understanding is ultimately concerned with relationships between propositions. From these views we can see that, unlike knowledge, understanding is not primarily concerned with a belief in an individual proposition, but instead what I will call a *mental representation of a relational structure*. While in the case of knowledge we talk of a belief in a proposition p , henceforth when we talk of understanding I will talk of a mental representation of relational structure ϕ .

What, exactly, is the nature of a relational structure? I will not here defend a particular view on this matter. Again, since I am agnostic on the question of reductionism, I leave it open as to whether a relational structure can be expressed in terms of propositions. As I noted above, understanding comes in many different forms, and it is up for debate whether any or all of these forms of understanding can be expressed ultimately in terms of propositions one represents. Consider, for example, one form of understanding that readily suggests such an expression, namely understanding-why. Some reductionists have argued that understanding why p is equivalent to knowing that p because q , where q is a reason or set of reasons. In this case the object of understanding would be a relational structure that consisted of p and the reasons one holds for p . Even if one analyzes understanding-why in this way, though, we would still evaluate whether one understands why p with regards to a different evaluative object than if we were to evaluate whether one knows that p . Consider the way we evaluate whether one knows that p : we check to see if one’s representation that p (in this case, a belief that p) meets the criteria we have set out for knowing. Consider a case of understanding why p : in this case we do not check to see

if one's representation *that p* meets a relevant set of criteria, but whether one's representation *that p and the reasons why p* meet those criteria. The representation of *p* and the reasons why *p*, then, would constitute the evaluative object for understanding-why. Again, one could express *p* and the reasons why *p* as a conjunctive proposition, and would most likely do so if one adheres to reductionism. The bottom line, though, is that although the surface grammar of attributions of knowledge and understanding why both involve atomic propositions – one knows that *p* and understands why *p* – it remains that knowledge that *p* is evaluated with respect to one's representation of *p*, whereas understanding why *p* is evaluated with respect not to the proposition *p* by itself. This difference, I will argue in what follows, is crucial in how we portray and address the question of whether one can have environmentally lucky understanding.

We can now say something about the relationship between grasping as a process and what is grasped: in general terms, the process of grasping is a mental function that takes as inputs any kind of doxastic or epistemic state, evidence, or information, and outputs a mental representation of a relational structure. Additionally, grasping seems to require more than *just* mentally representing a relational structure, since there can be cases in which one takes there to be relations between bits of information for very bad reasons. In such cases we would say that while one mentally represents a relational structure, one does not grasp it. Consider, for example, that after reading my horoscope I come to form a number of beliefs: that today I will receive a monetary windfall, that as a Scorpio my temperament is a double-edged sword, and that Saturn is at its perihelion. I then take these pieces of information to be importantly related, such that Saturn's distance from the sun has a calming effect on my fiery Scorpio temperament, which in turn allows me to make more level-headed financial decisions. While I thus possess a mental representation of a relational structure between beliefs and bits of information, I do not *grasp* it: I

did not, after all, employ any of the abovementioned grasping processes, as I do not plausibly have cognitive control over any of the relevant information, nor am I able to provide good explanations of any related phenomena.

Grasping as a process, then, is a species of a more general kind of representational process, namely one that outputs a representation of a relational structure. In these terms representational processes are analogous to belief-forming processes: while one can form beliefs in any number of different ways, only some of these ways are good in that they produce beliefs that are candidates for knowledge. Similarly, we can identify grasping processes as those good representational processes which produce candidates for understanding. I will not have anything specific to say about these processes here, however, as my claim is simply that the evaluative object of understanding is a mental representation of a relational structure, and that when we are determining whether one has understanding we look to such a structure to see whether it meets the right conditions.

In addition to the grasping condition, another widely accepted requirement for understanding is that understanding requires *truth*: just as we cannot know something that is false, we cannot understand something that is not the case⁹. We need to be careful, however, about how to conceive of this requirement. Since the evaluative object of understanding is a mental representation of a relational structure, we need to speak more generally in terms of representations of relational structures that either obtain or fail to obtain, namely whether the structure I represent is one that actually holds in the world. For example, in the above horoscope example another reason why I lack understanding is that the relational structure I represent fails

⁹ Although I take this to be the majority view, it is not the unanimous view. See, for example, Zagzebski (2008) and Elgin (2007) for dissent.

to obtain, as it is not in fact the case that there is any relationship between the position of Saturn and my financial circumstances.

To illustrate the claims I have put forth so far, consider as an example some that I *do* understand, namely why that wine glass, the one that my cat knocked onto the floor, broke. In determining whether I have this understanding we can look to the relational structure that I represent, namely the relations between my beliefs that wine glasses are fragile, that the tile floor is hard, that when fragile things fall from a significant height onto a hard floor they tend to shatter, that my cat is responsible for the glass falling in the first place, etc. If this relational structure is one that I represent as the result of a grasping process – say, I employed some relevant set of good cognitive abilities that tends to produce representations of structural relations which obtain – and obtains – because the glass really did break for the reasons that I represent – then the relational structure I represent is a plausible candidate for understanding. On the other hand, if the relational structure I represent does not obtain – that is, if it was not the cat but my accidentally jostling the table that knocked over the glass, or if the glass broke not because of impacting the floor but because an unknown defect – or if the process I used to represent the relational structure was not a good one – say, I drew the connections that I did out of wishful thinking because I did not want to be held responsible for the glass breaking – then I plausibly do not possess understanding.

We now have enough of a picture of the structure of understanding to consider whether it is compatible with environmental luck. A few remarks before proceeding. First, despite the comparisons I have made between understanding and the traditional conception of knowledge, I am not here advocating any particular list of necessary and sufficient conditions for understanding. What is important for my purposes is that understanding requires, at a minimum,

a grasp of a relational structure which obtains. I take it that these requirements are common to all forms of understanding; as such, I do not in what follows have a specific form of understanding in mind¹⁰. Second, I take everything that I have said so far to be compatible with either reductionism or nonreductionism. This is because I leave it open as to how, exactly, one mentally represents a relational structure. Again, I am remaining agnostic concerning the reductionism versus nonreductionism debate: it is my view that both sides have largely misconstrued the relationship between environmental luck and understanding.

With all this in mind, we are now in a position to formulate the question of how environmental luck affects understanding.

3. Formulating the Question

The difference of evaluative objects between knowledge and understanding requires that we pose the question of whether environmental luck is compatible with understanding differently. We have seen that the question as to whether one can have environmentally lucky knowledge is something like the following:

Environmental Luck and Knowledge (ELK): Can S know that p if the environment in which S forms her true belief that p makes it such that S could have easily been wrong about the fact that p , were S to have formed her belief in the same way?

¹⁰ The relevant differences between different forms of understanding, I take it, comes down to the components of the relevant relational structure and the grasping processes that are employed. For example, the components of a relational structure involved in understanding-why might be different from those in a relational structure involved in understanding-how. What these forms have in common, though, is that there is a relational structure that has to be grasped that also obtains. It is this structure that all forms of understanding have in common that make them incompatible with environmental luck.

The answer, we have also seen, is generally taken to be “no.” How should we formulate the question of whether understanding is compatible with environmental luck? One way would be to formulate a question similarly to ELK, substituting “understand” for “know”. Here, then, is what might result:

Environmentally Lucky Belief and Understanding (ELBU): Can S understand (why/how/that) p if the environment in which S forms her true belief that p makes it such that S could have easily been wrong about the fact that p , were S to have formed her belief in the same way?

But this is the wrong question to ask, since, as I have argued, we do not evaluate one’s understanding with respect to one’s belief *that* p . ELBU, then, assumes that the evaluative object of understanding is the same as that of knowledge, but this is a mistake. Instead we need to ask:

Environmentally Lucky Representations and Understanding (ELRU): Can S’s mental representation of ϕ qualify as understanding if the environment in which S represents ϕ makes it such that ϕ could have easily failed to obtain, were S to have represented ϕ in the same way?

Again, that luck in general affects the relationship between one’s mental representation and the way that the world could have been means that to determine whether environmental luck is compatible with understanding we need to consider situations in which one’s representation of ϕ obtains, but could have easily failed to. ELRU, then, is the question that needs to be addressed.

Cases that attempt to show that environmental luck is compatible with understanding, however, typically address ELBU instead. Consider, for example, a case from Duncan Pritchard (Pritchard et al., 2010)¹¹:

Nero and the Firefighters: Nero comes home to find his house in flames. When he asks a firefighter what caused the fire, she gives him the correct answer that it was a faulty breaker box. Unbeknownst to Nero, the person he asked is one of the few real firefighters on the scene, as many nearby people are dressed as firefighters en route to a costume party. Nero could have very easily asked these partygoers, and, had he done so, they would have given him a false answer while failing to indicate that they were not real firefighters. (78)

Nero, Pritchard claims, understands why his house burned down, despite his being lucky that he has a true belief about the cause of the fire. There is reason to think that Nero does indeed have this understanding: even though Nero could have easily received false information, and thus formed a false belief about the cause of the fire, this does not seem to matter in determining whether he understands. Again, consider the relational structure that Nero represents: first, it obtains (the faulty breaker box *did*, after all, cause the fire); and second, although it is not specified explicitly in the case, it is plausible that Nero not only represents that structure but also grasps it (we can assume that Nero employed a good grasping process). Nero's case is meant to elicit the following kind of intuition: as long as one is able to grasp the right kinds of relationships between true beliefs, in determining whether one has understanding it does not

¹¹ I am here using Khalifa's (2013) presentation of the case.

matter whether those beliefs are true in a way that they could have easily been false. Since this *does* matter when it comes to knowledge we are to conclude that environmental luck does not affect understanding in the way that it affects knowledge.

By reasoning in this way, though, we are answering the wrong question: the *Nero* case, at best, gives us reason to think that the answer to ELBU is “yes,” as it shows that one can have understanding in the face of a belief that is environmentally lucky to be true. What does Nero’s case say about ELRU? In order to answer this question we need to ask whether the *relational structure that Nero represents* could have easily failed to obtain due to relevant facts about his environment. I argue that this is not the case: given the details of Nero’s case, there is no reason to think that Nero’s representation of the relevant relational structure could have easily failed to obtain. As a result, his is not a genuine case of environmental luck for understanding.

It is true, of course, that Nero *could* have easily represented a relational structure which failed to obtain: if he had asked any of the costumed partygoers instead of the real firefighter he would have formed different beliefs – e.g. that the cause of the fire was an errant cigar tossed onto a pile of oily rags – and thus would have represented a relational structure that did not obtain – e.g. one in which the house burned down because of the combustibility of rags, the heat of smoldering cigars, etc. However, in those cases in which Nero asks one of the costumed partygoers he is not forming his mental representation *in the same way* as when he asks the real firefighter. Although there is luck involved in Nero’s case, then, it is merely *circumstantial* luck.

Call a *belief* circumstantially lucky when one’s belief is lucky to be true because one is lucky to have been in the circumstances that made it possible to form the belief in the first place. A belief that a car accident has occurred, for example, might be circumstantially lucky, given that it was unlikely that one was in a position to witness it; this, however, does not prevent one

from knowing that a car accident occurred. Or consider Peter Unger's (1968) classic example of a man who forms a belief that there is a turtle, but is lucky to have formed that belief since he only narrowly avoided being assassinated moments before; again, this kind of circumstantial luck does not prevent the turtle-watcher from knowing that there is a turtle. Circumstantial luck, then, is compatible with knowledge.

We can similarly define a conception of circumstantial luck for understanding. Say that S's mental representation of relational structure ϕ is circumstantially lucky to obtain when S is lucky to have been in the circumstances that made it possible to form a representation of ϕ in the first place. For example, the car accident witness can understand why the accident occurred, despite it being the case that she was unlikely to have witnessed the accident in the first place, and Unger's turtle-watcher could become the world's greatest herpetologist, acquiring a vast amount of understanding of turtles, despite it being continually unlikely that she would be in the position to observe any turtles as she spends her life narrowly avoiding incompetent assassins. When it comes to determining the way in which one represents a structure of relations between beliefs, though, part of what constitutes the relevant circumstances in which one forms the relevant representation is the fact that one has those beliefs. After all, just as one would not form one's belief in the same way were one in totally different circumstances (as illustrated in the car accident and turtle-watcher cases), one would not form a representation of a relational structure in the same way were one to have different beliefs. That one's beliefs constitute part of one's *circumstances* when it comes to evaluating understanding means that cases like Pritchard's do not address the question of whether understanding is susceptible to *environmental* luck.

Indeed, this is, I think, the general problem with cases that attempt to show that environmental luck is compatible with understanding, namely that they are ones in which it is

only one's beliefs that constitute the relational structure one represents that are lucky to be true¹².

Here is another example which exemplifies the same problem, an oft-discussed case from

Kvanvig (2003):

Comanche Dominance: Consider, say, someone's historical understanding of the Comanche dominance of the southern plains of North America from the late seventeenth until the late nineteenth centuries. Suppose that if you asked this person any question about this matter, she would answer correctly. Assume further that the person is answering from stored information; she is not guessing or making up answers, but is honestly averring what she confidently believes the truth to be. Such an ability is surely constitutive of understanding, and the experience of query and answer, if sustained for a long enough period of time, would generate convincing evidence that the person in question understood the phenomenon of Comanche dominance of the southern plains. But does she have knowledge? Ordinarily, yes; but it is not required. For, on the usual theories of knowledge, all those answers could be given from information possessed and still fail to be known to be true, because the answers might only be accidentally true. For example, most history books might have been mistaken, with only the correct ones being the sources of the understanding in question and with no basis in the subject for preferring the sources consulted over those ignored. (197-8)

In Kvanvig's case, while one's environment makes it such that one could have easily formed false beliefs by picking the wrong book (and thus at least potentially makes it such that one fails

¹² Again, if we consider the type of mental representation involved in understanding to be a belief, we then need to distinguish between a belief in ϕ and the beliefs that comprise ϕ . What a case like Pritchard's shows, then, is the latter, when the question of the compatibility of understanding and environmental luck needs to address the former.

to know), it does not seem to prevent one from *understanding* why the Comanche dominated the southern plains. Indeed, it again seems plausible that one *does* acquire understanding in this case, since one represents a relational structure which obtains, and which is (plausibly) the result of a good grasping process.

Comanche Dominance is similar to *Nero and the Firefighters*, as in both cases one is lucky to form true beliefs given one's environment, and in both cases one intuitively understands regardless. But again we need to ask: although the book-reader's representation of ϕ – say, a structure that involves one's various beliefs and information about the Comanche dominance of the southern plains – obtains, is it the case that, in relevantly similar circumstances in which S represents ϕ in the same way, ϕ also obtains? Since, as I argued above, S's beliefs that form part of the basis of her representation of ϕ are part of her circumstances, the answer to this question appears to be “yes”: there is no reason to think that S would have represented any structure other than ϕ given that she had the same beliefs and represented ϕ in the same way (that is to say, using the same grasping process). Kvanvig's case, then, like Pritchard's, does not show a genuine case of environmentally lucky understanding.

Others have also argued that cases like Pritchard's and Kvanvig's do not show what they aim to show. Generally, such arguments are part of a defense of a reductionist picture which attempts to show that attributions of knowledge and understanding always stand or fall together. Cases like Pritchard's and Kvanvig's threaten the reductionist thesis specifically by attempting to show that, due to the luck involved in the cases, one can have understanding without knowledge. Responses from the reductionist camp attempt to show that in such cases one either possesses neither knowledge nor understanding, or that one possesses both knowledge and understanding. Wilkenfeld et al. (forthcoming), for example, present an empirical study which purports to show

that the folk deny both knowledge and understanding in a similar case that is meant to show the compatibility of luck with understanding¹³, while Sliwa (2015), Riaz (2015), Kelp (2014), and Khalifa (2013), attempt to show that the luck involved in cases like Pritchard's and Kvanvig's is not of a type that is incompatible with knowledge, since these cases are not sufficiently analogous to quintessential cases demonstrating environmental luck (i.e. barn façade cases).

As I am agnostic about the reductionist thesis I can similarly remain agnostic about whether cases like Pritchard's and Kvanvig's are, in fact, ones in which one lacks knowledge due to the effects of environmental luck, as well as whether these cases are structurally similar enough to the barn façade case¹⁴. My argument is that the kind of luck that is involved in these cases is compatible with understanding not necessarily because of the structure of the cases, but because the cases target the wrong evaluative object.

So far I have argued that prominent cases that purport to show the incompatibility of environmental luck and understanding miss the mark because they address the wrong question. Let's consider, then, a case which I think addresses the right question.

4. A Case of Incompatible Environmental Luck for Understanding

¹³ This is Grimm's (2006) "exploding chestnut case." The case is too long and detailed to outline here. However, because it is so complex, I am generally skeptical of the value of folk judgments of the case.

¹⁴ I think there is a relevant difference between the *Comanche* and *Nero* cases and standard cases of circumstantial luck that make it plausible that they are genuine cases of environmentally lucky beliefs. The difference is between being lucky to have acquired a true belief and being lucky to have acquired a true belief when one could have just as easily acquired a false belief. Sliwa (2015: 63), for instance, compares the *Comanche* case to a case involving a detective who, while investigating a murder scene, happens to spot a clue that leads to the arrest of the murderer, a clue that could have easily been missed. She then claims that the *Comanche* case is more similar to the detective case than the barn façade case. But I think this is mistaken: in Sliwa's case the detective was lucky to form a true belief regarding a vital clue, but he was not in any danger of believing anything false *about that clue*. Compare this to a case in which I am in some sense lucky to have a true belief that a car accident occurred: in this situation I was not in any danger of forming the false belief that the car accident did not occur. In the *Nero* and *Comanche* cases, though, I take it that there is a danger in believing a relevant falsehood in place of a truth. I take there to be good reason, then, to think that the *Comanche* and *Nero* cases involve a kind of luck that is generally incompatible with knowledge.

Consider the following variant on *Nero and the Firefighters*:

Claudius and the Electrician: On his way home, Claudius finds that all the houses in his neighbourhood in Rome are on fire. Wondering what caused the fires, he stops by one of the houses, which happens to be Nero's. When Claudius asks a nearby electrician what caused the fire, she gives him the correct answer that it was a faulty breaker box. The electrician is trustworthy and knowledgeable, and there are no masquerading electricians nearby. On the basis of his true belief that the breaker box was faulty, plus background beliefs regarding electricity, the nature of breaker boxes, heat, and fires, Claudius forms a mental representation of the relationship between Nero's breaker box and the fire: without a functioning breaker box to control the flow of electricity, wires in Nero's house were heated to the point that the surrounding wood combusted. And, in fact, Nero's house did burn down as a result of the faulty breaker box in the way that Claudius represents. Unbeknownst to Claudius, however, due to a recent outbreak of house fires, Roman officials decided that all houses would be retrofitted with lower-voltage breaker boxes. As a result, every house in the city has been retrofitted, except for Nero's, which still uses the old breaker box technology. Despite this change, rampant house fires persist as the result of faulty breaker boxes. When a faulty breaker box causes a fire in a retrofitted house, however, it is not for the same reasons that it caused Nero's house fire: instead, the new breaker boxes give off sparks which ignite flammable materials that Romans like to store in their basements.

Claudius, I submit, does not understand why Nero's house burned down. This, despite the facts that he forms true beliefs on the basis of the information given to him by the electrician, ones

that are not themselves true as a matter of environmental luck, and that Claudius represents a relational structure which obtains. Nevertheless, had Claudius happened to have stopped by any of the other house fires and formed the *same belief* – specifically, the belief that the house fire is caused by a faulty breaker box – he would not represent a relational structure that obtained, since given his background beliefs concerning how breaker boxes typically cause fires and his being unaware of the recent breaker-box upgrades, he would similarly take the breaker box malfunction and fire to be related, but in the wrong way. Claudius, then, got lucky in that the relational structure he mentally represents obtains in his particularly situation: had he stopped by any other house he would have represented the same relational structure in the same way, and that relational structure would not have obtained. As in cases of environmental luck that prevent one from knowing, a key part of the reason why the relational structure that Claudius represents obtains is due to facts about his environment of which he is unaware.

To illustrate the point further, consider what would happen if Claudius had stopped by Livia's house instead. Livia's house, unlike Nero's, is one of the many houses that has been retrofitted with new breaker boxes, and is also on fire. When Claudius asks a nearby electrician what caused the fire, she responds that it was a faulty breaker box. In this case Claudius again forms a mental representation of the relationship between Livia's breaker box and the fire, one pertaining to overheating wires and combustible wood. Unlike Nero's house, however, Claudius' representation is not reflective of reality: Livia's faulty breaker box did not cause her house to combust in the way that Claudius represents. In this case it seems clear that Claudius does not understand why Livia's house burned down. But for Claudius, there is no relevant difference between Livia's house and Nero's house: he would form the same mental representation in the

same way when told that the breaker box was the cause of the fire, regardless of whose house he happened to visit.

Here, then, is the parallel between Claudius' case and the barn façade case: in the barn façade case an important part of the explanation as to why Henry's belief is true is because of the nature of his environment, details of which he is unaware: typically, when you see what looks like a barn on the side of the road, it is a barn, except in barn façade country, when it typically isn't. In Claudius' case the same thing happens, but with a representation of a relational structure: typically, when you are told that a breaker box caused a house fire it is because of overheating wires, except when you are in Rome, when that typically isn't the case. In both Henry's and Claudius' cases they come across the one exception to a rule that they do not realize holds of their environment, and in both Henry's and Claudius's cases they would form the same mental representation regardless of whether it obtained. Thus, just as Henry fails to know, Claudius fails to understand.

A motivation for holding the view that environmental luck is incompatible with understanding is the intuition that I am attempting to elicit from you in the Claudius and the Electrician case. One might think that there could be a different, better explanation for this intuition, one that does not require us to hold the view that Claudius does not have understanding. Morris (2012) makes such an argument: he argues that at least some forms of understanding are, in fact, compatible with epistemic luck (what he calls "explanatory understanding"), and that cases that attempt to elicit intuitions to the contrary can be explained away by the fact that we are not always careful in ordinary discourse about distinguishing ascriptions of propositional knowledge from ascriptions of understanding. As a result, we will mistakenly deny understanding in cases in which we correctly deny knowledge due to the effects

of luck. What our intuitions might be tracking in a case like Claudius and the Electrician, then, is not a lack of understanding, but a lack of *knowledge*.

There is something odd about this response, though: to say that it is sometimes a mistake to deny understanding because one has denied knowledge already assumes that understanding is not a form of knowledge. As I am here a reductionism agnostic, I cannot assume that it is ever a mistake to deny understanding because one has denied knowledge. Indeed, Grimm (2006) argues that because relevant knowledge and understanding attributions tend to stand or fall together, this constitutes evidence for reductionism, not a phenomenon to be explained away. That is not to say that these intuitions could still not be explained away. However, under the assumption that reductionism is false, it is not clear *why* we would we mistakenly deny understanding when we deny a corresponding instance of knowledge. Indeed, in the *Claudius and the Electrician* case I argue that Claudius lacks understanding in a way that is independent of his possession of knowledge, namely that he does not represent a relational structure in the right way. Failing to represent a relational structure does not directly impact whether one has knowledge, though, since relational structures are not the evaluative objects of knowledge. We cannot, then, explain away the intuition in this case in the way that Morris might want.

There are a few reasons beyond the intuitive why we should think that Claudius does not understand. First, I have already argued that once we get a fix on the right evaluative object for understanding we can see how *Claudius and the Electrician* is relevantly analogous to the barn façade case. Second, Claudius' case is of a general type, namely cases in which what one attempts to understand is an *exception to a rule*. When there is a general rule that governs phenomena of a certain type, but for which there are exceptions, it seems that one cannot understand that exception without also recognizing the fact that it is exceptional. These rules and

exceptions can come in different forms. In chess, for example, pawns can only capture a piece that is on the diagonal, except in very specific circumstances when they can capture *en passant*. One can correctly represent the relations that exist between the conditions under which the rule is applied and the state of the chessboard, but one cannot understand the *en passant* rule without furthering recognizing that it is an exception that only applies in certain circumstances. Again, we have seen reason why this is the case: one hallmark of understanding is that it allows one to explain similar phenomena in a range of different circumstances. By attempting to understand an exception without recognizing that it is exceptional, though, one will be unable to recognize the range of circumstances in which the rule can be applied, and will thus be lead to provide incorrect explanations in different circumstances.

Claudius' inability to recognize that Nero's house constitutes an exception provides the basis for the second reason why we should think that he does not have understanding: he is unable to provide the right explanation of why the other houses in his neighbourhood burned down (again, consider the explanation that he would provide for Livia's house burning down). Indeed, that he is able to provide the right explanation of Nero's house fire is a matter of luck, as he happened to pick out the one house for which the processes he employs results in a mental representation of a relational structure that obtains. Again, note that in neither *Comanche Dominance* nor *Nero and the Firefighters* is the phenomena that is understood *exceptional*, and in both cases the subject is able to provide robust explanations: given the same beliefs in similar circumstances, the respective subjects would continue to represent the correct relational structures, and would, in turn, be able to provide correct explanations of similar phenomena. So there is again reason to think that in these cases one has understanding, but that environmental luck does not play a relevant role.

Perhaps, though, you are still unmoved: perhaps you think that Claudius does, in fact, understand why *Nero's house in particular* burned down, despite him not understanding why, for example, Livia's house burned down¹⁵. There is, after all, an unsurprising explanation as to why Claudius does not understand why Livia's house burned down: Livia's house did not burn down for the reasons that Claudius thinks it did. It would be perfectly consistent, one might say, to claim that Claudius understands in one case but not the other. The crucial point, though, is that if a hallmark of understanding is that it requires that one be able to not only provide a correct explanation in the relevant case, but in all cases that one takes to be relevantly similar, then Nero does not possess the relevant understanding in the individual case. Consider again the barn façade case: reasoning in the same way we might say that Henry knows that *that* is a barn, even if he does not know that any of the barn façades are barns. I do not think, though, that it is tempting to reason in this way, and similarly I do not think we should be tempted to reason in the same way about Claudius' case. We can again consider the *en passant* case. Let's say that I see you use your pawn to capture a piece and ask you why you were able to do that. You tell me that it was because pawns can capture *en passant*. I can form a relevant true belief, but without having the additional information that pawns can *only* capture in *this specific way in these specific circumstances*, it is hard to see how I could be said to understand either the *en passant* rule or why you were able to capture the relevant piece. This is because I would not be able to distinguish the relevant circumstances in which the exceptional circumstances obtained. Once I possess this additional information I can then recognize that the rule is an exception to a pawn's normal capturing behavior, and be said to understand, but I cannot be said to understand without this additional information.

¹⁵ Thanks to an anonymous editor's suggestion of this interpretation of the case.

5. Conclusion

My argument that environmental luck is incompatible with understanding can be summarized as follows. Luck in epistemology in general concerns the way in which one's mental representation is related to the way the world is and the way it could have been, and environmental luck concerns the way in which facts about one's environment influence the way in which said evaluative object obtained and could have obtained. Since the evaluative object of understanding is the mental representation of a relational structure, to evaluate the compatibility of environmental luck and understanding we need to look at cases in which said representation obtains in one's specific circumstance, but not in a wide range of related circumstances. I have argued that cases that attempt to show the compatibility of luck with understanding miss the mark because they are evaluating instances of understanding when considering the wrong evaluative object, thereby conflating circumstantial luck and environmental luck. I have also presented a case which I take to exemplify the way in which environmental luck is incompatible with understanding, namely one in which one's mental representation of a relational structure obtains in one's specific circumstances, but would fail to obtain in a wide range of similar circumstances when formed in the same way. As the fact that environmental luck is incompatible with knowledge says something important about the concept of knowledge, so too does the fact that environmental luck is incompatible with understanding say something important about the concept of understanding, namely that understanding requires that one's mental representation not only obtains, but obtains in a robust way.

Works Cited

- Bourget, David. Forthcoming. "The Role of Consciousness in Grasping and Understanding." *Philosophy and Phenomenological Research*: 1-59.
- Elgin, Catherine. 2007. "Understanding and the facts." *Philosophical Studies* 132.1:33 - 42
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23.6: 121-3.
- Goldman, Alvin I. 1976. "Discrimination and perceptual knowledge." *The Journal of Philosophy* 73.20: 771-791.
- Grimm, Stephen. 2006. "Is understanding a species of knowledge?" *British Journal for the Philosophy of Science* 57.3:515-535.
- Hills, Alison. 2015. "Understanding Why." *Noûs* 50.2: 1-28.
- Kelp, Christoph. 2015. "Understanding phenomena." *Synthese* 192.12:3799-3816.
- . "Knowledge, Understanding and Virtue." In A. Fairweather (ed.), *Virtue Scientia. Virtue Epistemology and Philosophy of Science*. Springer.
- Khalifa, Kareem. 2013. "Understanding, Grasping and Luck." *Episteme* 10.1: 1-17.
- Kvanvig, Jonathan L. 2003. *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
- Morris, Kevin. 2012. "A Defense of Lucky Understanding." *British Journal for the Philosophy of Science* 63.2: 357-71.
- Pritchard, Duncan, Alan Millar, and Adrian Haddock. 2010. *The nature and value of knowledge: Three investigations*. Oxford University Press.
- Riaz, Amber. 2015. Moral understanding and knowledge. *Philosophical Studies* 172.1:113-128.

- Sliwa, Paulina. 2015. "Understanding and Knowing." *Proceedings of the Aristotelian Society* 115.1:57-74.
- Sosa, Ernest. 2001. "Human knowledge, animal and reflective." *Philosophical Studies* 106.3:193 - 196.
- Unger, Peter. 1968. "An Analysis of Factual Knowledge." *The Journal of Philosophy* 65.6: 157-70.
- Williamson, Timothy. 2001. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wilkenfeld, Daniel. 2013. "Understanding as Representation Manipulability." *Synthese* 190.6: 997-1016.
- Wilkenfeld, Daniel A.; Plunkett, Dillon & Lombrozo, Tania. 2016. "Depth and deference: When and why we attribute understanding." *Philosophical Studies* 173.2:373-393.
- . (forthcoming). Folk attributions of understanding: Is there a role for epistemic luck? *Episteme*: 1-26.
- Zagzebski, Linda. 2008. *On Epistemology*. Belmont, CA: Wadsworth.