

Understanding, Grasping, and Luck¹

Forthcoming in *Episteme*

Kareem Khalifa

Department of Philosophy

Middlebury College

Middlebury, VT 05753

kkhalifa@middlebury.edu

Abstract: Recently, it has been debated as to whether understanding is a species of explanatory knowledge. Those who deny this claim frequently argue that understanding, unlike knowledge, can be lucky. In this paper, I argue that current arguments do not support this alleged compatibility between understanding and epistemic luck. First, I argue that understanding requires reliable explanatory evaluation, yet the putative examples of lucky understanding underspecify the extent to which subjects possess this ability. In the course of defending this claim, I also provide a new account of the kind of “grasping” taken to be central to understanding. Second, I show that putative examples of lucky understanding unwittingly deploy a kind of luck that is compatible with knowledge. Finally, appealing to a number of works on explanation and its attendant epistemology, I argue that alleged instances of lucky understanding that overcome these two obstacles will invariably violate certain norms of explanatory inquiry—our paradigmatic understanding-oriented practice. By contrast, knowledge of the same information is immune to these criticisms. Consequently, if understanding is environmentally lucky, it is always inferior to the understanding that a corresponding case of knowledge would provide.

¹ Parts of this paper were presented at a conference on the Aims of Inquiry and Cognition Conference, May 2012 in Edinburgh. Various audience members, including Anne Baril, Stephen Grimm, Allan Hazlett, Andrew McGonigal, Paul Noordhof, Ted Sider, Nick Treanor, and Lani Watson provided useful feedback on these ideas at this time. Stephen Grimm provided detailed comments on an earlier draft of this paper. I thank all of them for their feedback.

Understanding why something is the case is clearly different than *knowing that* it is the case. But is such understanding any different than *knowing why* something is the case? In other words, is there any difference between understanding and explanatory knowledge? Those answering this question in the affirmative frequently cite understanding's unique compatibility with epistemic luck (Kvanvig 2003, 2009; Morris 2012; Pritchard 2008, 2009, 2010)². Roughly stated, advocates of lucky understanding hold that understanding, but not explanatory knowledge, can involve true beliefs that easily could have been false.

Defenders of lucky understanding also appeal to the ability to “grasp” how different propositions in a given domain hang together as a defining feature of understanding. However, beyond a few programmatic remarks, little else is said about the nature of grasping. In this paper, I shall present and defend a new account of grasping, and then argue that satisfying its conditions renders the prospects of lucky understanding quite bleak.

Section 1 presents the current arguments for lucky understanding. Section 2 then develops my account of grasping, and argues that alleged examples of lucky understanding feature subjects who lack this cognitive ability. Section 3 then highlights another defect with these examples, namely that they fail to provide clear cases of knowledge-undermining luck. Section 4 then presents possible ways of revising these examples so as to avoid these difficulties. I argue that even in these examples, lucky

² Others argue that understanding, unlike knowledge, is non-factive (Elgin 2007; Khalifa 2011b). I will not address this line of argument in this paper, though Khalifa maintains that non-factive understanding can still be a species of knowledge.

understanding is most charitably construed as always inferior to the understanding that a corresponding case of explanatory knowledge would provide.

1. Putative Cases of Lucky Understanding

Before proceeding, a clarification is in order. Here and throughout, I shall restrict my attention only to whether the subjects understand *why* something is the case, where understanding *why p* entails the true belief that *q explains p* for some *q*³. Except where needed, I will not broach the details wherein *q explains p*. Additionally, references to “understanding” hereafter refer to “exemplary understanding” unless otherwise noted. By “exemplary understanding,” I refer to the understanding characteristic of (though not limited to) well-established science. With this in mind, I shall now examine the argument upon which proponents of lucky understanding pin their hopes. In this section, I simply present these arguments. Subsequently, I highlight their paucities.

Epistemologists have grappled with several examples where luck undermines a putative instance of knowledge. While these are frequently all labeled as “Gettier cases,” I will follow Pritchard—a central advocate for lucky understanding—in distinguishing *Gettier-style* luck from *environmental* luck. According to Pritchard’s (2009, 21) dichotomy, in the former cases, bad luck “intervenes ‘betwixt belief and fact⁴,’” e.g.

Sheep Example

Abel is a shepherd who looks upon his pasture one day and sees something that looks like a sheep. He thus comes to believe that a sheep is in the pasture. As it

³ For some arguments to the contrary, see (Kvanvig 2009; Lipton 2009). For rebuttals, see (Khalifa forthcoming a, forthcoming b), respectively.

⁴ Pritchard credits Unger with coining this phrase.

turns out, Abel is looking at a shaggy dog that is perfectly occluding a sheep that is actually in the pasture. Had there been no sheep behind the dog, Abel would have falsely believed that a sheep was in the pasture.⁵

Abel's true belief is clearly lucky. To paraphrase Pritchard, the shaggy dog intervenes between the belief and the fact that the sheep is in the pasture. As such, beliefs that are susceptible to this kind of luck are not candidates for knowledge.

While some have argued that, unlike knowledge, understanding is compatible with Gettier-style luck (Morris 2012), I will not address these arguments here, and will instead focus on *environmental luck*. Unlike Gettier-style luck, environmental luck does not intervene between fact and belief. Goldman (1976), by way of Gettier, provides its classic incarnation:

Barn Façade Example

Bonnie is driving through a rural county, when she sees a barn-shaped object. Thinking this unexceptional, Bonnie forms the true belief that there is a barn in the field. However, unbeknownst to her, she is driving in Barn Façade County, where most objects that look like barns are in fact fakes. Moreover, had Bonnie looked at a fake barn, she would have believed it was a real barn.

Bonnie's true belief is clearly lucky, however, unlike Abel's shaggy dog, the piece of bad luck—that she is in Barn Façade County—does not intervene between the belief and the fact that a barn is in the field. Nevertheless, it is widely held that environmentally lucky beliefs of this sort do not amount to knowledge.

⁵ Paraphrased from (Chisholm 1966).

Using analogies with the Barn Façade Example, some have argued that, unlike knowledge, understanding is compatible with environmental luck⁶:

Nero Example

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 firefighters⁷.

Just as Bonnie could have very easily looked upon a barn façade and formed the false belief that a barn is in the field, Nero could have very easily consulted a partygoer and formed the false belief that, e.g. an arsonist caused the fire. Hence, Pritchard reasons that Nero's belief is environmentally lucky. Pritchard (2010, 79) then argues that "the agent concerned has all the true beliefs required for understanding why his house burned down, and also acquired this understanding in the right fashion. It is thus hard to see why the mere presence of environmental epistemic luck should deprive the agent of understanding."

Let us render Pritchard's reasoning explicit:

Environmentally Lucky Understanding Argument

⁶ The most prominent arguments to this effect come from (Kvanvig 2003; Pritchard 2010). Hereafter, I focus on Pritchard's argument, since it provides the most careful analysis of the relationship between understanding and luck. My treatment of Pritchard readily applies to Kvanvig.

⁷ Paraphrase from (Pritchard 2010, 79).

- (ELU1) If two beliefs are similar in all epistemically relevant respects, and if one of those beliefs does not amount to knowledge, then the other belief also does not amount to knowledge⁸.
- (ELU2) Bonnie's belief in the Barn Façade Example does not amount to knowledge.
- (ELU3) Bonnie's belief in the Barn Façade Example and Nero's belief in the Nero Example are similar in all epistemically relevant respects.
- (ELU4) Nero's belief in the Nero Example amounts to understanding.
- (ELU5) So Nero's belief in the Nero Example does not amount to knowledge, but it does amount to understanding.

I shall attack the third and fourth premises. Against (ELU4), Section 2 challenges Nero's claim to understanding, and presents the basic contours of my account of grasping.

Against (ELU3), Section 3 argues that the Nero Example differs in epistemically relevant ways from the Barn Façade Example. Finally, Section 4 shows that any new example that addresses these difficulties will result only in a pyrrhic victory: lucky understanding may be possible, but only as a way station to explanatory knowledge⁹.

⁸ Those who affirm pragmatic encroachment might take issue with this premise, e.g. (Fantl and McGrath 2009). This would provide yet another avenue for challenging this argument. I do not address it, because: (a) space prohibits the independent defense of pragmatic encroachment required to render this challenge cogent, and (2) it is unclear how to assess the pragmatic differences in the Barn Façade and Nero Examples.

⁹ My conclusions resemble Grimm's (2006), though our premises are quite different. Grimm's view of "grasping" differs from mine, and does not play a central role in his critiques of lucky understanding. Also, his arguments deploy neither the modal structure of luck nor the specific literatures on explanation that I draw upon.

2. Defective Understanding

My first challenge to arguments for lucky understanding is that Nero may lack understanding, casting doubts upon (ELU4). Consequently, genuine tests of lucky understanding must safeguard against this challenge in a way that current defenses of lucky understanding do not.

My argument against (ELU4) proceeds in two steps. First, I argue that understanding involves the ability to reliably evaluate explanations; I consider this ability to be a central but hitherto under-theorized aspect of “grasping” (2.1). Second, I argue that there is little evidence that Nero is a reliable explanatory evaluator, casting doubts on his understanding (2.2). In the process of defending this latter claim, I argue that the only other viable ability in the Nero Example—the ability to vet testimony—is less central to understanding than is explanatory evaluation. Consequently, the emphasis on testimony—shared by the vast majority of examples marshaled both for and against the claim that understanding is a species of knowledge—has added an unnecessary distraction to the discussions to date.

2.1. Reliable explanatory evaluation

First, regardless of one’s stance on understanding’s compatibility with luck, it is widely thought that understanding is constituted, in part, by the exercise or mastery of certain cognitive abilities (de Regt 2009a, 2009; de Regt and Dieks 2005; Grimm 2010; Kvanvig 2003, 2009; Pritchard 2008, 2009, 2010). Furthermore, if one holds that understanding involves such abilities, then presumably these abilities are reliable in producing the relevant beliefs. Thus, understanding is partly constituted by reliable cognitive abilities.

Importantly, all of those who currently hold that understanding both abhors Gettier luck and tolerates environmental luck endorse this idea¹⁰.

But which cognitive abilities figure in understanding? The ability to “grasp” explanatory and evidential connections within a body of information is frequently mentioned. The problem is that “grasping” has not been clearly articulated. I will go some ways towards filling this lacuna, by arguing that grasping entails reliable explanatory evaluation.

Before proceeding, I stress that this paper only requires that reliable explanatory evaluation is *necessary* for grasping—not that it is sufficient. Grasping frequently involves some kind of reflective awareness bearing close kinship to internal justification (Kvanvig 2003). The reader is welcome to append these internalist requirements to my reliability requirement, but they will not bear on what follows.

My argument for the indispensability of reliable explanatory evaluation is as follows:

Grasping Argument

- (G1) Understanding entails true beliefs of the form *q explains p*.
- (G2) Understanding entails that such beliefs must be the result of exercising reliable cognitive abilities.
- (G3) If understanding entails true beliefs of the form *q explains p*, and also entails that such beliefs must be the result of exercising reliable cognitive

¹⁰ This stems from certain commitments about virtue epistemology that need not concern us here. Those who hold the stronger claim that understanding tolerates both Gettier and environmental luck do not hold that understanding entails reliable cognitive ability. As mentioned earlier, I do not address these arguments here.

abilities, then these abilities involve evaluating (or discriminating between) explanations.

- (G4) So understanding entails that beliefs of the form *q explains p* are the result of exercising a reliable cognitive ability to evaluate explanations.

Thus, if “grasping” is the central ability for understanding, grasping involves reliable explanatory evaluation.

Since we have stipulated that Premise (G1) is a ground rule for the discussion, it will remain untouched. Moreover, to repeat, everybody who denies Gettier-lucky understanding but still accepts environmentally lucky understanding also accepts (G2). Thus, dialectically speaking, this premise is also on firm ground. But if both these premises are granted, then (G3) is highly plausible. If someone reliably forms beliefs about explanations, then they must not be susceptible to believing incorrect explanations. But this is just to say that they can discriminate between correct and incorrect explanations.

Now, admittedly, our ability to evaluate explanations is frequently quite modest, such that even in cases where it is negligible, we might be inclined to say that *some* understanding is nevertheless achieved. While this puts some pressure on (G3), it’s unproblematic for two reasons.

First, the quantity and quality of explanations we must evaluate in order to understand are naturally seen as context-sensitive. For instance, the explanations that students and laypeople must be able discriminate between in order to achieve understanding differ from those doing cutting-edge science. Thus, one may claim that since standards of reliable explanatory evaluation are context-sensitive, so too are

standards of understanding-attribution. Thus, so long as *some* reliable explanatory evaluation is required in any context, (G3) is defensible.

Second, even if there are some contexts in which no such evaluation is required, critics of (G3) seem destined to score only a pyrrhic victory. Observe that in cases where our ability to evaluate explanations is modest (or non-existent), our understanding is correspondingly modest. This is why it's natural to claim that scientists typically understand a given phenomenon *better* than students and laypeople. More broadly, in the alleged counterexamples to (G3), there is usually a hedge that we only have *some* (rather than exemplary) understanding. However, if defenses of lucky understanding hinge on rejecting (G3) along these lines, then it's hard to see why we should care very much about lucky understanding. On such a denial of (G3), only cases of relatively modest understanding could be lucky, meaning that in serious inquiry, lucky understanding would be little more than a second-rate surrogate for explanatory knowledge. Since that would rob lucky understanding of much of its interest, I assume its proponents would not make this concession, and that (G3) is plausible.

Alternatively stated, if exemplary understanding is a species of knowledge and understanding also admits of degrees, then a natural way of accounting for different degrees of understanding is to treat exemplary understanding as an ideal that lesser degrees of understanding only approximate¹¹. Thus, deniers of (G3) simply reinforce this picture if they appeal to cases of non-exemplary understanding in which the ability to evaluate explanations is muted. To make their case, they need a case of *exemplary*

¹¹ For an account of understanding that works along these lines, see (Khalifa forthcoming b).

understanding where the ability to vet explanations is nevertheless muted. Since no such examples have been offered, (G3) remains on firm ground.

Consequently, reliable explanatory evaluation is a necessary feature of understanding. Earlier, we acknowledged the characteristic ability of understanding is often called “grasping.” If we wish to maintain that convention, then reliable explanatory evaluation is an essential aspect of grasping. However, I stress that this is merely terminological, so if the reader feels strongly that grasping and reliable explanatory evaluation should be decoupled, so much the worse for the word “grasping” in this context. My argument only requires understanding to entail reliable explanatory evaluation.

Below, I discuss some further details about the nature of explanatory evaluation, but here is a quick sketch. Explanatory evaluators’ inputs are various potential explanations of a phenomenon plus a body of relevant evidence, and their outputs are beliefs about which of these potential explanations is an actual explanation of this phenomenon. Paradigmatically, the process of explanatory evaluation proceeds by comparing the evidence that would obtain if different potential explanations were true with actual evidence. Potential explanations are then eliminated when significant evidence that would have obtained had they been true fails to obtain. When this process of elimination runs smoothly, only correct explanations of the phenomenon remain¹². Explanatory evaluators are reliable just in case the resulting beliefs could not easily have been false.

¹² See, e.g. (Lipton 2004) for further details.

2.2. Does Nero Understand?

Thus far, I've argued that reliable explanatory evaluation is the characteristic ability of understanding. Thus, an unreliable explanatory evaluator does not understand. Returning to (ELU4), Nero's reliability as an explanatory evaluator is unclear, so the extent to which he understands is similarly unclear.

While Nero uses reliable methods to identify a firefighter (by looking at the person's uniform, etc.), it's far less clear that he's used reliable methods for evaluating explanations. Depending on the details of this testimonial exchange, Nero's ability to consider alternative explanations, consider evidence that would favor one of these explanations over its rivals, etc. might be very low. Indeed, the Nero Example is naturally interpreted as being a rather brief conversation between Nero and the firefighter in which none of these details are rehearsed. Consequently, it does not appear that Nero has engaged in reliable explanatory evaluation, and we should be very skeptical that he understands why his house caught fire.

However, the Nero Example involves *both* Nero's grasping *and* Nero's ability to evaluate testimonial evidence. Typically, a belief that can be produced by two independent abilities is reliably formed even if only one of these abilities is reliable. For instance, if someone with reliable vision wishfully thinks that a dollar is under the couch, her belief that a dollar is under the couch is nevertheless reliably produced upon looking under the couch and seeing a dollar. Analogously, one might argue that although Nero's grasping is unreliable, his ability to identify good informants is reliable. Hence, as long as testimonial evaluation can produce understanding in the same way as explanatory evaluation, Nero understands despite being an unreliable explanatory evaluator.

This argument is flawed. First, Nero is not—at least in this instance—a very reliable identifier of experts, for he could have easily consulted a partygoer. Second, I suspect that this reverses things: proponents of environmentally lucky understanding may well concede that Nero is lucky but unreliable in his testimonial evaluation, but his ability to “grasp” the explanatory information (the details of which they leave underspecified) then picks up the slack and provides him with understanding. However, since it’s very unclear why we should impute to Nero any robust ability to evaluate the explanation presented to him, it’s unclear why we should impute to him any robust ability to grasp that explanation.

However, even if proponents of lucky understanding can avoid these worries, it seems thoroughly implausible that testimonial evaluation can produce understanding in the same way as explanatory evaluation. To see this, consider two cases:

Case I: Tommy can discriminate real firefighters from fake firefighters, but cannot discriminate between correct and incorrect explanations of the fire. Thus, if a real firefighter offered him an incorrect explanation, he would believe it, and if a partygoer offered him a correct explanation, he would disbelieve it.

Case II: Mira cannot discriminate real firefighters from fake firefighters, but can discriminate between correct and incorrect explanations of the fire. Thus, Mira can identify correct and incorrect explanations of the fire regardless of who offers them.

It’s at least an open question whether Tommy really understands why his house caught fire in Case I. Furthermore, Nero does not even appear to have this much ability; he

would believe false explanations from real and fake firefighters. By contrast, Mira's understanding in Case II is on firm ground, and is all the more impressive as a cognitive achievement when she's fed a bogus explanation by a firefighter and still comes to believe the correct explanation through highly acute explanatory evaluation.

Hence, when it comes to understanding, explanatory evaluation is a more important ability than identifying experts. This suggests that the preoccupation with the role of testimony in producing understanding has been a rather significant distraction. Consequently, the best tests of lucky understanding will remove this element of testimony, and focus exclusively on explanatory evaluation.

In saying this, I do not deny that reliable testimony can provide understanding. Rather, my point is primarily about the *priority* of explanatory evaluation over testimonial evaluation: the extent to which an audience understands via testimony will always be parasitic on the extent to which a speaker's utterance enables that audience to evaluate explanations reliably.

3. Benign Epistemic Luck

In addition to difficulties with (ELU4), the third premise of the Environmentally Lucky Understanding Argument, (ELU3), faces difficulties. Specifically, it assumes a strong similarity between the classic Barn Façade Example and the Nero Example. I will argue that this is mistaken; the examples are in fact quite *different* in epistemically relevant ways, such that Nero's luck is compatible with his knowing.

To do this, I first render environmental luck's general structure explicit via the Barn Façade Example. I will then show that the Nero Example does not exhibit this same

structure, such that the sense in which Nero is lucky is “benign,” i.e. compatible with his *knowing* that the house caught fire because of a faulty breaker box. Consequently, the Nero Example is not a clear instance of understanding without explanatory knowledge.

To better clarify the structure of environmental luck, let’s examine the modal structure of the Barn Façade Example. Here I will take Pritchard’s (2009, 34) Safety Principle as my cue:

S’s belief is safe [i.e. not epistemically lucky] *iff* in most near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as the actual world, her belief continues to be true.

This notion of luck does not discriminate between Gettier and environmental luck. To motivate that distinction, I offer the following:

Schema for Environmental Luck

A subject *S* is environmentally lucky with respect to the belief that *p* if:

- (1) In the actual world w_0 ,
 - a. *S*’s belief that *p* is true, and
 - b. *S*’s belief that *p* is produced by one or more of *S*’s reliable cognitive abilities α and the fact that *p*; and
- (2) In a nearby world w_1 , everything is the same as w_0 except:
 - a. *S*’s belief that *p* is false, and

- b. S 's belief that p is produced by α and the fact that p -*façade* (rather than the fact that p .)¹³

In Gettier-style luck, condition (1.b) will differ, for S 's belief will not be produced by the fact that p , but by the fact that p -*façade*. For example, in the Sheep Example, the fact that a shaggy dog is in the pasture produces Abel's belief that the sheep is in the pasture.

More pressing for my purposes is the notion of *p-façades*. While this is an admittedly loose concept, the crucial bit for my argument is that *p-façades* play the same role in producing the subject's belief that p in w_1 as the truth-makers of p do in w_0 . For instance, the barn façade in the field plays the same role in producing Bonnie's belief in w_1 as the barn in the field plays in producing that same belief in w_0 . As I'll now argue, this is the crucial difference between the Barn Façade and Nero Examples, for if we were to follow the template for environmental luck above, the following would result:

- (1) In world w_0 ,
- a. Nero's belief that *the house's faulty breaker box explains why it caught fire* is true; and
 - b. Nero's belief that *the house's faulty breaker box explains why it caught fire* is produced by Nero's reliable ability to identify experts and the fact that the house's faulty breaker box explains why it caught fire.
- (2) In world w_1 , everything is the same as w_0 except:

¹³ Here, I profess agnosticism concerning any deep metaphysical issues about whether it is facts or something else (e.g. events) that produce beliefs, or whether production is reducible to a more rigorously analyzed concept (e.g. causation). The reader may fill in or substitute these variables as she sees fit. I also stipulate that α captures Pritchard's requirement that the belief be formed in the "same way" in both the actual and nearby possible world.

- a. Nero's belief that *the house's faulty breaker box explains why it caught fire* is false; and
- b. Nero's belief that *the house's faulty breaker box explains why it caught fire* is produced by Nero's reliable ability to identify experts and the fact that *something like* a faulty breaker box explains why the house caught fire¹⁴.

Call this the *Strict Model*. By contrast, the Nero Example is naturally interpreted as follows:

- (1) In world w_0 ,
 - a. Nero's belief that *the house's faulty breaker box explains why it caught fire* is true; and
 - b. Nero's belief that *the house's faulty breaker box explains why it caught fire* is produced by Nero's reliable ability to identify experts and the fact that **he consulted the real firefighter**.
- (2) In world w_1 , everything is the same as w_0 except:
 - a. Nero's belief that (e.g.) *arson explains why the house caught fire* is false; and
 - b. Nero's belief that *arson explains why the house caught fire* is produced by Nero's reliable ability to identify experts and the fact that **he consulted a partygoer**.

¹⁴ Here and throughout, I assume an "ontic" conception of explanation, i.e. that explanations are facts in the world. One could replace this with the other dominant view, in which explanations are sets of true propositions, with the only loss being in economy of prose. For more on these distinctions, see (Salmon 1989).

The Strict Model differs most starkly from the Nero Example in its “environmental factor.” The following summarizes the relevant similarities and differences:

	Belief	Fact	Relevant Contrast
Barn Façade	A barn is in the field.	A barn is in the field.	A barn (<i>rather than a barn façade</i>) is in the field.
Strict Model	A faulty breaker box explains why the house caught fire.	A faulty breaker box explains why the house caught fire.	A faulty breaker box (<i>rather than something like a faulty breaker box</i>) explains why the house caught fire.
Nero Example	A faulty breaker box explains why the house caught fire.	Not specified, though presumably the same as the Strict Model	A real firefighter (<i>rather than a partygoer</i>) provides the testimony.

Table 1. Differences between the Barn Façade, Strict, and Nero Examples.

Clearly, the Strict Model respects the structure of the Barn Façade Example better than the Nero Example. Even if we grant that the truth-maker in the Strict Model and The Nero Example (i.e. the second column) are the same, the latter still trades in very different counterfactual claims, as is evidenced by the contrast it invokes (third column): Nero could have very easily consulted a different person, and that person would have offered false testimony. By contrast, fidelity to the structure of the Barn Façade Example suggests that Nero would have been indifferent to subtle explanatory differences, e.g. he might have confused a faulty breaker box with another explanation of the fire.

Disanalogies *per se* are harmless. However, in this case, these differences undercut the claim that the Barn Façade and Nero Examples are similar in *epistemically relevant* ways, i.e. (ELU3). Specifically, the Nero Example seems to be flirting with an epistemically innocuous form of luck that arises when one could have easily received different evidence. For instance, consider the following:

Jesse James Example

Jesse James robs a bank and is just about to make his grand escape when his mask accidentally slips off. A bystander, Byron, gets a clear view of James's face at this moment, and owing to the latter's infamy, Byron immediately comes to believe that Jesse James robbed the bank. However, had Byron been standing at a slightly different location, he would not have formed a belief about the robber's identity¹⁵. Now, despite the fact that Byron could have easily been standing at an angle where he did not see Jesse James's face, Byron knows who robbed the bank. Pritchard (2005, 137) thus cites this example as a "benign" form of epistemic luck. More generally, he takes a whole class of examples of this sort to be compatible with knowledge:

Evidential epistemic luck

It is lucky that the agent acquires the evidence that she has in favor of her belief (Pritchard 2005, 136).

Returning to the Nero Example, a person's testimonial evidence is nothing more than the propositions she takes others to be asserting. Hence, in the Nero Example, the real firefighter offers different testimonial evidence than the partygoer that Nero could have easily asked. Thus, the Nero Example doesn't involve environmental luck; only evidential luck. Since evidential luck is compatible with knowledge, Nero may well *know* the explanation of the house fire. Thus, (ELU3) is highly suspect.

To summarize, two of the core assumptions of the arguments for lucky understanding—that the subjects are victims of "malignant" or knowledge-undermining luck (ELU3), and that they understand the phenomena confronting them (ELU4)—are on shaky ground. In highlighting these deficiencies, I do not claim to have established that

¹⁵ Paraphrased from (Nozick 1981).

understanding is a species of knowledge. Rather, I have tried to bring the contours of a proper instance of lucky understanding into relief: a genuine test of environmentally lucky understanding will involve something that respects the Strict Model's constraints and in which the subject exercises reliable explanatory evaluation. It is to these more telling examples that I now turn.

4. Environmentally Lucky Understanding?

The considerations about explanatory evaluation and evidential luck might inspire defenders of lucky understanding to revise their examples to control for these features. One tempting candidate is something along the following lines¹⁶:

Lucky Partygoer Example

O'Leary 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, and carefully relays the details as to how she arrived at this explanation. O'Leary eagerly and correctly absorbs all of these details. Unbeknownst to O'Leary, 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. O'Leary could have very easily asked these partygoers, and, had he done so, they would merely be guessing as to the fire's explanation, while failing to indicate that they were not real firefighters. However, as luck would have it, their guesses would have been *true* and identical to the real firefighter's in every detail.

¹⁶ I call it "tempting" because a version of this example is frequently raised when I have presented the arguments in Section 3.

Here, O’Leary’s testimonial evidence could not easily have favored another belief. Consequently, evidential luck *appears* dormant in this example. Moreover, we have stipulated that O’Leary absorbs detailed information about the firefighter’s explanatory evaluation, so he is positioned to be a reliable explanatory evaluator. Hence, it would appear that he also understands.

However, the Lucky Partygoer Example is *less* promising as an instance of environmental luck than the Nero Example. Indeed, the Lucky Partygoer Example is an even *more benign* example of evidential luck than the Nero Example. First, in the nearest possible world in which O’Leary’s belief is false, the partygoers offer different testimonial evidence than they offer in the actual world. Thus, initial appearances notwithstanding, this is yet *another* case of evidential luck. Second, suppose that both Nero and O’Leary had consulted the partygoers. *Ceteris paribus*, Nero’s belief would be false; O’Leary’s, true. Consequently, Nero’s belief more easily could have been false than O’Leary’s. Hence, if the Nero Example trades in evidential luck, then so does the Lucky Partygoer Example.

However, better examples of lucky understanding can be gleaned by paying closer attention to the aforementioned difficulties with the Nero Example. First, the Nero Example’s modal structure should be replaced with that of the Strict Model’s. This blocks the objection that the subjects only enjoy benign epistemic luck. Second, the subjects should acquire true beliefs through the exercise of reliable explanatory evaluation. This blocks the objection that the subjects do not genuinely understand. After presenting such an example (4.1), I will argue that even in these cases, understanding is still lacking (4.2).

4.1. A Proper Example

The preceding suggests a clear test for whether understanding is compatible with environmental luck. We will need an example where *S* understands why *p* and:

- (1) In world w_0 ,
 - a. *S*'s belief that *q* explains *p* is true
 - b. *S*'s belief that *q* explains *p* is produced by *S*'s reliable ability to evaluate explanations and the fact that *q* explains *p*.
- (2) In world w_1 , everything is the same as w_0 except:
 - a. *S*'s belief that *q* explains *p* is false
 - b. *S*'s belief that *q* explains *p* is produced by *S*'s reliable ability to evaluate explanations and the fact that *q-façade* (rather than *q*) explains *p*.

Call this the *Explanatory Model*. This respects the lessons of the Strict Model, but also makes *S*'s ability to evaluate explanations—rather than her ability to identify experts—the primary cognitive ability of interest. As we saw, this should assuage any worries that the subject does not really understand. Thus, if the Explanatory Model furnishes an unambiguous case of understanding, then lucky understanding is in good shape. The following respects the strictures of the Explanatory Model:

Fiona Example

Fiona is a firefighter who does an exemplary job in collecting all of the evidence that could be extracted from the embers of Nero's house. On the basis of this evidence, she comes to believe that a faulty breaker box is the root cause of the fire, and this is in fact true. However, at the moment of the fire, the breaker box

malfunctioned at the same time that the grounding wire shorted. The two events were causally independent, and the shorted grounding wire did not actually cause the fire. But, had the shorted grounding wire (rather than the faulty breaker box) caused the fire, Fiona would have discovered exactly the same evidence and still believed that the breaker box's malfunctioning explains why Nero's house caught fire.

Thus, in a nearby possible world w_I , Fiona continues believing the breaker box explanation when in fact the grounding wire explanation is correct.

The Fiona Example seems to be an ideal test case for the claim that understanding tolerates environmental luck. Fiona has exactly the same evidence in both the actual and relevant possible world, so there is no risk of confusing environmental luck with evidential luck. Her belief remains the same in both worlds, and it is simply a quirk of the environment—say that the shorted grounding wires and faulty breaker box leave very similar burn patterns in their aftermath—that Fiona has happened upon the correct explanation. So this is a compelling case of environmental luck. Furthermore, since explanatory evaluation is playing a leading role, the Fiona Example *shouldn't* be susceptible to the same worries as the Nero Example, namely that the subject doesn't really understand.

4.2. An Argument Against Lucky Understanding

However, as I'll now argue, those worries still linger. A very basic feature of reliable explanatory evaluators is that they can rule out confounds, i.e. other plausible explanations of the phenomenon of interest. But the Fiona Example assumes that Fiona is

deficient in precisely this aspect of evaluating explanations. Furthermore, this suggests an argument that generalizes to any example that fits the Explanatory Model:

The Rival Explanation Argument

- (RE1) If *S*'s understanding of why *p* is environmentally lucky, then *S*'s true belief that *q* explains *p* is produced by explanatory evaluation that fails to rule out rival explanations of why *p* that easily could have been true¹⁷.
- (RE2) If *S*'s belief that *q* explains *p* is produced by explanatory evaluation that fails to rule out rival explanations of why *p* that easily could have been true, then *S*'s ability to evaluate explanations is not reliable.
- (RE3) Hence, if *S*'s understanding of why *p* is environmentally lucky, then *S*'s true belief that *q* explains *p* is not produced by a reliable cognitive ability.
- (RE4) If *S* understands why *p*, then *S*'s true belief that *q* explains *p* is produced by a reliable cognitive ability.
- (RE5) So *S*'s understanding of why *p* is not environmentally lucky.

(RE1) is the idea that the true beliefs in understanding can be cognitive achievements regardless of the machinations of certain luck-inducing “façades” in the environment. Insofar as it is contentious, it is that rival explanations might strike some readers as odd things to include in an “environment.” However, we have already rehearsed the arguments as to why alternative explanations are the environmental wildcards that bear the strongest structural analogy to Goldman’s barn façades. Furthermore, as the Fiona Example illustrates, this might mean nothing more than that the explanation concerns a

¹⁷ I leave the notion of ‘ruling out’ unanalyzed, though the general idea is that a person would be able to readily identify some deficit in the explanation (say that it does not square with the evidence or is vastly inferior to another explanation.)

very delicate causal structure in the world, and such structures are far more common than counties replete with barn façades.

As mentioned above, all who hold that understanding tolerates environmental luck endorse (RE4). So, if there is any place to attack this argument, it is with (RE2). I will not provide a conclusive proof that it is true, but I will motivate it in three different ways.

First, it is important to stress that it is a fairly modest demand. There need not be underdetermination on a massive scale that every explanatory evaluator must overcome in order to be reliable. Rather, explanatory evaluators need only be able to rule out alternative explanations that *easily could have been true*. This accords with the idea that environmental luck entails a certain lack of epistemic safety, and that w_1 must be a *nearby* possible world.

Second, (RE2) accords with our intuitions. Imagine that Fiona's colleague, Keisha, is investigating the scene, and, recognizing that a shorted ground could have easily explained the fire, she believes that *either* a faulty breaker box *or* a shorted ground explains why the house caught fire. Keisha is clearly a more reliable explanatory evaluator than Fiona, and this appears to be entirely because she has considered alternative but plausible explanations.

Third, ruling out these kinds of alternative explanations is a pervasive feature in exemplary practices of explanatory evaluation. Indeed, our most reliable modes of explanatory evaluation—controlled experiments—are precisely designed to rule out alternative explanations. James Woodward's account of an ideal intervention helps to bring this point into relief. Let I , X , and Y be variables. Then an ideal intervention I on X

with respect to Y is a change in the value of X that changes Y , if at all, *only via* the change in X . More fancifully put:

(IV)

I is an *intervention variable* on X with respect to Y if and only if I meets the following conditions:

- I1. I causes X .
- I2. I acts as a switch for all the other variables [V] that cause X . That is certain values of I are such that when I attains those values, X ceases to depend the value of other variables that cause X and instead depends only on the value taken by I .
- I3. Any directed path from I to Y goes through X . That is, I does not directly cause Y and is not a cause of any causes [S] of Y that are distinct from X except, of course, for those causes of Y , if any, that are built into I - X - Y connection itself: that is, except for (a) any causes of Y that are effects of X (i.e. variables that are causally between X and Y) and (b) any causes of Y that are between I and have no effect on Y independently of X .
- I4. I is (statistically) independent of any variable Z that causes Y and that is on a directed path that does not go through X (Woodward 2003, 98).

Essentially, **(IV)** guarantees that our independent variable X is explanatorily relevant to Y , and conditions I2-I4 rule out V , I , S , and Z as other potential explanations of a change in Y . For instance, if I2 fails to obtain, then it is unclear whether a change in Y is explained by a change in X or a change in V . Parallel points apply to I3 (which rules out I and S as

alternative explanations), and I4 (Z). Thus, these conditions prohibit plausible alternative explanations from easily being true.




Figure 1: Conditions for interventions. The solid arrows indicate causal relations; the dotted line, a correlation; the black box, the requirement I1; and the grey boxes, the restrictions I2-I4 discussed above.

Hence, the epistemology motivating a perfectly controlled experiment—arguably our foremost means of evaluating explanations—entails that we must be able to rule out alternative explanations that easily could have been true. Of course, many of our explanations, such as Fiona’s, are not gleaned from controlled experiments, but, insofar as we can gather non-experimental evidence to evaluate these explanations, they generally attempt to approximate the evidence that would arise under the idealized conditions described by (I1)-(I4)¹⁸. As Woodward (2003, 114) puts it, these conditions provide a “regulative ideal.”

¹⁸ For a more thorough discussion of this, see (Shadish, Cook, and Campbell 2001).

Explanations that hold under a wide range of conditions are also preferred to those that do not. Depending on the author, this explanatory virtue is called scope (Lipton 2004), power (Lycan 2002), consilience (Thagard 1978), stability (Mitchell 1997, 2000), or resiliency (Skyrms 1980)¹⁹. This also is a sign that ruling out alternatives is an important feature of our practices of explanatory evaluation. Essentially, more powerful explanations of this sort gain additional justification by the evidence that they can explain and their competitors cannot. Hence, if the practice of preferring these kinds of explanations is reliable, ruling out rival explanations is important, as (RE2) states.

Woodward's view on this issue is that explanations are frequently judged superior when they involve generalizations that are *invariant under more interventions*²⁰. In other words, explanatory generalizations can be expressed as equations in which Y is a function of X , and these generalizations have greater explanatory power when they cannot easily be perturbed, i.e. they hold under a larger class of interventions. In this case, not only does each intervention provide the means of ruling out certain alternative explanations, but explanations that invoke generalizations that hold under a larger class of interventions are also deemed superior. Hence, the greater the invariance of an explanatory generalization, the greater resources it provides for ruling out rival explanations²¹.

¹⁹ This is not to say that these concepts are equivalent, only that they all seem to capture the idea that good explanatory generalizations hold under a wide variety of conditions.

²⁰ My use of Woodward is only meant to vividly illustrate that excluding rival explanations is an important aspect of explanatory evaluation, but any epistemology of explanation says this much. Hence, I submit that careful readings of, e.g. (Hempel 1966; Kitcher 1989; Lipton 2004; Popper 1963; Thagard 1978) will furnish very similar support for (RE2) of this argument. Indeed, insofar as explanatory evaluation is a species of hypothesis testing, virtually every theory of confirmation also requires that plausible alternatives be ruled out.

²¹ For more on invariance, see (Hitchcock and Woodward 2003; Woodward and Hitchcock 2003; Woodward 2003, Ch. 6).

Consequently, I take (RE2) of my argument to be on firm ground. With this, understanding's compatibility with environmental luck is unmotivated under the most telling of conditions, namely the Explanatory Model's.

At this point in the argument, one may simply resist the conclusion offered here by citing the bare intuition that Fiona understands why the house caught fire. However, this seems to raise my earlier worries that lucky understanding can only be defended if it is regarded as a second-rate surrogate of explanatory knowledge. Specifically, understanding admits of degrees, which in turn track with how well one's understanding approximates explanatory knowledge acquired via reliable explanatory evaluation. Such a position would entail Fiona has *some degree* of understanding, but that it is *lacking*. This preserves the intuition that motivates the objection, but it also explains a further intuition: namely that someone who could reliably rule out the grounding wire explanation would (*ceteris paribus*) better understand why the house caught fire than Fiona. Hence, even at the level of brute intuition, lucky understanding can at best be regarded as subordinate to explanatory knowledge.

5. Conclusion

I have argued that understanding's compatibility with environmental epistemic luck is on shaky ground. First, I have argued that reliable explanatory evaluation is the cognitive ability most central to understanding, yet the examples supporting lucky understanding do not specify this ability in any detail. In the process, I hope to have also shed new light on the kind of grasping involved in understanding. Second, by more closely scrutinizing the modal structure of the relevant examples, I have shown that analogies with

paradigmatic cases of environmental luck are strained. Third, I have argued that even when these objections are addressed, lucky understanding emerges as an inferior substitute for the understanding provided by explanatory knowledge. Consequently, understanding's status as a species of knowledge is unthreatened by its relationship to environmental luck.

- Chisholm, R. (1966), *Theory of knowledge, (Foundations of Philosophy Series)*. Englewood Cliffs: Prentice-Hall.
- de Regt, H. W. (2009), "Understanding and scientific explanation", in H.W. de Regt, S. Leonelli, and K. Eigner. (ed.), *Scientific Understanding*, Pittsburgh: University of Pittsburgh Press, 21-42.
- (2009a), "The epistemic value of understanding", *Philosophy of Science* 76 (5):585-597.
- de Regt, H. W., and D. Dieks (2005), "A contextual approach to scientific understanding", *Synthese* 144 (1):137-170.
- Elgin, C. (2007), "Understanding and the facts", *Philosophical Studies* 132 (1):33-42.
- Fantl, J., and M. McGrath (2009), *Knowledge in an uncertain world*. Oxford: Oxford University Press Oxford.
- Goldman, A. I. (1976), "Discrimination and Perceptual Knowledge", *Journal of philosophy* 73 (20):771-791.
- Grimm, S. R. (2006), "Is understanding a species of knowledge?", *British journal for the philosophy of science* 57 (3):515-535.
- (2010), "The goal of understanding", *Studies in the history and philosophy of science* 41 (4):337-344.
- Hempel, C.G. (1966), *Philosophy of natural science, Prentice-Hall Foundations of Philosophy Series*. Englewood Cliffs: Prentice-Hall.
- Hitchcock, C.R., and J. Woodward (2003), "Explanatory generalizations, part II: Plumbing explanatory depth", *No ̄s* 37 (2):181-199.
- Khalifa, K. (2011b), "Understanding, Knowledge, and Scientific Antirealism", *Grazer Philosophische Studien* 83 (1):93-112.
- (forthcoming a), "Is understanding explanatory or objectual?", *Synthese* (doi: 10.1007/s11229-011-9886-8).
- (forthcoming b), "The role of explanation in understanding", *British Journal for the philosophy of science* (doi: 10.1093/bjps/axr057).
- Kitcher, P. (1989), "Explanatory unification and the causal structure of the world", in Philip Kitcher and Wesley C. Salmon (eds.), *Scientific Explanation*, Minneapolis: University of Minnesota Press, 410-506.
- Kvanvig, J. L. (2003), *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.

- (2009), "The value of understanding", in A. Haddock, A. Millar and D. Pritchard (eds.), *Epistemic value*, Oxford: Oxford University Press, 95-111.
- Lipton, P. (2004), *Inference to the best explanation*. 2nd ed. New York: Routledge. Original edition, 1991.
- (2009), "Understanding without explanation", in H.W. de Regt, S. Leonelli, and K. Eigner. (ed.), *Scientific Understanding*, Pittsburgh: University of Pittsburgh Press, 43-63.
- Lycan, W. G. (2002), "Explanation and epistemology", in Paul K. Moser (ed.), *The Oxford Handbook of Epistemology*, Oxford: Oxford University Press, 408-433.
- Mitchell, S. D. (1997), "Pragmatic laws", *Philosophy of Science* 64:468-479.
- (2000), "Dimensions of scientific law", *Philosophy of Science* 67 (2):242-265.
- Morris, Kevin (2012), "A Defense of Lucky Understanding", *The British Journal for the philosophy of science* 63 (2):357-371.
- Nozick, Robert (1981), *Philosophical explanations*. Cambridge, Mass.: Harvard University Press.
- Popper, K. R. (1963), *Conjectures and refutations : the growth of scientific knowledge*. London: Routledge & K. Paul.
- Pritchard, D. (2005), *Epistemic luck*. New York: Oxford University Press.
- (2008), "Knowing the answer, understanding, and epistemic value", *Grazer Philosophische Studien* 77:325-339.
- (2009), "Knowledge, Understanding and Epistemic Value", *Royal Institute of Philosophy Supplements* 64:19-43.
- (2009), "Safety-Based Epistemology: Whither Now?", *Journal of Philosophical Research* 34:33-45.
- (2010), "Knowledge and understanding", in D. Pritchard, A. Millar and A. Haddock (eds.), *The nature and value of knowledge: three investigations*, Oxford: Oxford University Press, 3-90.
- Salmon, W. (1989), "Four decades of scientific explanation", in Philip Kitcher and Wesley Salmon (eds.), *Scientific Explanation*, Minneapolis: University of Minnesota Press, 3-219.
- Shadish, William R., Thomas D. Cook, and Donald Thomas Campbell (2001), *Experimental and quasi-experimental designs for generalized causal inference*. Boston: Houghton Mifflin.
- Skyrms, Brian (1980), *Causal necessity: a pragmatic investigation of the necessity of laws*. New Haven: Yale university press.
- Thagard, P. (1978), "The best explanation: criteria for theory choice", *Journal of philosophy* 75:76-92.
- Woodward, J. (2003), *Making things happen: a theory of causal explanation*. New York: Oxford University Press.
- Woodward, J., and C. Hitchcock (2003), "Explanatory generalizations, Part I: A counterfactual account", *Nous* 37 (1):1-24.