



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

Ampliative Understanding

Alex Whalen

PhD Philosophy

The University of Edinburgh

2017

Abstract

Virtue-theoretic accounts of knowledge start by capturing the value of knowledge as an achievement and work from there to develop a full theory of knowledge. But environmental luck, which is compatible with achievements but typically defeats knowledge, introduces some unique challenges for these accounts to overcome. While far from devastating for the virtue-theoretic project, several authors have viewed these worries as an opportunity to shift their focus towards understanding.

In the past, understanding has been mostly ignored by epistemologists who considered it to be a psychological state rather than something worth further inquiry. Over the past decade, this view has changed and understanding is quickly becoming a topic of great interest and lively debate. Among the key questions in this debate is the relationship between knowledge and understanding, the role of epistemic luck, and whether understanding has final value as a cognitive achievement. However, the debate is taking place in the absence of a useful theory of understanding that can provide a principled means of addressing these topics.

This project aims to help remedy the situation by identifying a kind of understanding, which I call ampliative understanding, that can provide a framework in which the current debate can take place. In staying true to the virtue-theoretic approach, this account of understanding starts by focusing on its value as a cognitive achievement and working from there. On this view, an agent with ampliative understanding will be able to acquire true beliefs in a way that manifests her cognitive abilities. While there are certainly other kinds of understanding that may be of epistemological import, ampliative understanding is able to accommodate our intuitions about the value of understanding and can capture most of the necessary features for understanding that we find in the literature. My hope is that, with the framework of ampliative understanding in place, we can have a debate that is both rigorous and productive.

Lay Summary

Suppose you want to learn about something new like astronomy or the French Revolution. The goal of this kind of inquiry is not to just acquire knowledge about these topics; instead, we tend to aim for understanding. In other words, it's not enough to know information on a topic—we want to see how that information is connected as a whole. This intuition has helped motivate a shift in focus in epistemology away from classic questions about knowledge and towards understanding. In addition, there are some seemingly intractable problems for knowledge that understanding may be able to address.

While the literature on understanding, despite its relative infancy, has presented us with some unique insights, we have lacked a framework from which we can provide principled and well-informed responses to the central questions in the discussion of understanding. This project aims to provide such a framework—that is, one that can help to guide future research on understanding. We will also look at a full account of understanding in order to get a feel for what this framework does for developing these accounts.

The main feature of this approach, which I call ampliative understanding, is a person's ability to acquire true beliefs because of her understanding. In other words, someone who understands a topic, explanation, or subject matter will be able to use her understanding to navigate the information in that topic and even figure out other, new, true beliefs that she can include in her understanding. Thus we are able to hang on to the intuitions that understanding is valuable and is a suitable goal for inquiry.

Table of Contents

Introduction.....	1
From Knowledge to Understanding.....	7
Two Criteria for Achievement.....	22
The Value of Understanding Thesis.....	41
Ampliative Understanding.....	61
Further Research.....	74
References.....	77

1

Introduction

Understanding has historically not received much attention within epistemology. The past decade, however, has seen a surge in both the number of epistemologists thinking about understanding and the amount of literature on the topic. What was once considered nothing more than a psychological state, the notion of understanding is quickly becoming a hot topic within epistemology. Much of this surge can be traced back to 2003 with the publication of Jonathan Kvanvig's seminal book *The Value of Knowledge and the Pursuit of Understanding*. In this book, Kvanvig looks at a number of problems for analyses of knowledge and its value. Ultimately, he uses these worries to shift the discussion to understanding as it pertains to epistemology.

Despite the surge in the literature on understanding, epistemologists still lack a robust account of understanding. To borrow from Jaegwon Kim, "Producing a usable account of understanding is the hard part; one scarcely knows where to begin. What we lack is an understanding of understanding that is clear and rich enough for useful theorizing" (1994: 69). The goal of this work is to meet Kim's challenge—to develop an account of understanding that can help guide the debate. There are currently a number of open questions regarding both the nature and value of understanding. Is understanding a suitable goal for inquiry? How are understanding and knowledge related? Is understanding valuable for its own sake? There have been some wonderful insights into these questions in the literature; but without an account of understanding in the background to help inform the debate, it's difficult to see where these insights are taking us.

This project aims to establish a framework for developing an account of understanding. We will use this framework to put together an analysis of a kind of understanding, what I call *ampliative understanding*. The key feature of ampliative understanding is that an agent with this kind of understanding has the ability to acquire true beliefs based on her cognitive engagement with a body of information. In other words, an agent with ampliative understanding can apply her understanding to new scenarios or problems rather than simply knowing the answer through, say, rote

memorization. But before we forge ahead, some introductory remarks are in order.

1.1 The Road Ahead

As Kim suggested above, the hard part with developing an account of understanding is figuring out where to start. The literature on understanding is still in its infancy—at least when compared to more traditional topics within epistemology. Despite this, there have already been a number of important insights with which we can engage. But our starting point is going to be two guiding intuitions that will frame the dialectic of this work. The first intuition is that understanding is an epistemic state that is over and above knowledge. This is not to suggest that understanding is *more* valuable than knowledge or that understanding is distinct from knowledge. It is compatible with this intuition that understanding is a kind of ‘super knowledge’. The key idea is that a typical case of understanding involves more cognitive engagement from the agent than a typical case of knowledge. Chapter 2 will pick up on this intuition to motivate the overall project: a shift in epistemological focus away from knowledge and towards understanding.

The second intuition is the claim that understanding is valuable for its own sake. In other words, understanding has *final value*. In Chapter 3 we will begin to develop this claim by looking at achievements. This project relies on a virtue-theoretic approach which captures the final value of understanding as a cognitive achievement. We will take the lessons learned from our discussion of achievements and, in Chapter 4, develop this into what I call the *value of understanding thesis*, which captures understanding’s final value.

Chapter 5 begins with an analysis of ampliative understanding, which we develop from the framework established over the previous chapters. We will also look at two influential discussions on understanding—one from Jonathan Kvanvig and one from Stephen Grimm—to see how our framework can accommodate the main claims on these accounts. Throughout the discussion, we will consider possible amendments and revision to the analysis on offer. What is important in this project in the framework we develop; the analysis serves more as a concrete example built from this framework with which we can engage.

We will return to some of these key questions on understanding in Chapter 6 to see what principled answers ampliative understanding can provide. We will also consider other projects can be developed from the work we have done here. Before

getting started, we should first be clear about the kind of understanding we are discussing and give a picture of what a successful account of understanding should aim to do.

1.2 The Uses of Understanding

Before a conversation about understanding can even get off the ground, we need to make it very clear what we are talking about. There are a number of ways that understanding is used in ordinary language; some of these uses are of epistemological import while others aren't. Broadly speaking, we are looking for uses of understanding or understands that pick out a particular epistemic state. For example, if someone says "Bob is an understanding person," this is describing Bob's psychological or emotional state. Something like this might be said in response to accidentally breaking Bob's window—he's not going to flip out because he's an understanding person. Similarly, when someone says "God understands", this is trying to convey that God has empathy for the human condition. Neither of these uses tells us about Bob's or God's epistemic states—their cognitive relationship to some proposition or body of information.

We can also identify hedging uses of understands. If I suspect (but am not completely convinced) that my wife is angry with me for forgetting to take out the trash, I might say "Look, I understand that you're mad at me, but I had a lot on my mind." This usage is a little less clear-cut since one might reasonably conclude that I believe that my wife is angry with me. But ultimately this usage is similar to what we saw in "God understands." I am trying to empathize with my wife by acknowledging the state of affairs and hopefully diffuse the situation. Even though we might try to draw some epistemic conclusions from this usage of understands, it really comes down to a psychological state rather than an epistemic one.

So where can we draw the line? For Kvanvig, what distinguishes the uses of understanding that are of epistemological import is whether or not they are factive. Kvanvig identifies two uses of understanding that we should focus on; he calls these *propositional* and *objectual* understanding. Propositional understanding "occurs when we attribute understanding in the form of a propositional operator" (Kvanvig (2003): 191). In other words, when 'understands' takes a proposition as its grammatical object, we have propositional understanding. A straightforward example of this would be "Sam understands that gravitation is one of the four fundamental forces". Kvanvig also intends

for *wh*- clauses to be captured by propositional understanding as well. This is because *wh*- clauses can be put in terms of understanding *that* something is the case. Understanding why *X* is the case, for example, can be explicated in terms of understanding that *X* is the case because of *Y*. Objectual understanding occurs when an agent understands some particular object such as a subject matter, event, or field of study. Examples of this include “Becky understands astrophysics” or “Jeff understands the French Revolution.” Notice that objectual understanding isn't as straightforwardly factive as propositional understanding since astrophysics and the French Revolution are not the sorts of things that can be true or false. But while propositional understanding wears its factivity on its sleeve, objectual understanding still has factivity “in the background” (*ibid.*).

It isn't clear whether Kvanvig intends for propositional and objectual understanding to be two distinct *kinds* of understanding or whether he is simply identifying two different uses that are of the same kind. Although Kvanvig never refers to these as kinds of understanding, other authors¹—whether intentionally or not—do just this. Ampliative understanding can handle both uses, so we don't need to come down on one side of this debate. But for ease of discussion I will occasionally refer to these two uses (or kinds) of understanding.

Kvanvig's focus on factive uses of understanding seems to do a nice job picking out only those uses that are of epistemological import. But there is a problem, which we will explore in Chapter 4, with his factivity requirement for understanding. At the end of the day, this won't be a problem because the account of understanding on offer will be able to pick out those important uses of understanding. But in the meantime, I want to move away from factivity and try to draw the distinction we need along different lines. One thought might be to look at only those uses of understanding that are propositional. This will easily capture propositional understanding; objectual understanding—while not straightforwardly propositional—does seem to have propositions in the background, much like it does with factivity. But this move will have the undesirable effect of including hedging uses of understanding such as “I understand that you're angry with me”, since ‘You are angry with me’ is a proposition.

Instead, our focus will be on *doxastic* uses of understanding—those which (at least implicitly) attribute certain beliefs to the agent. Whether an agent understands the French Revolution or understands that some logical systems are necessarily incomplete,

¹c.f. Elgin (2009)

there is a straightforward sense in which an agent will have beliefs about the French revolution or logical systems. Hedging uses, on the other hand, don't ascribe any particular belief state to the agent. In the example we saw, I suspect that my wife is angry with me, but fall short of believing this claim. It is this lack of belief that makes this a hedging use of understanding. Notice that this doxastic focus is in the same spirit as Kvanvig's distinction based on factive uses, just without the general problems of factivity that we will see later on.

We will also need to discuss what desiderata a successful account of understanding should aim to meet. But first, a quick caveat is in order. Both the factive and doxastic distinctions end up capturing a particular kind of understanding: *linguistic* understanding, or what we might call comprehension. These are uses such as "I understand what you are saying" or "Jim understands German." I want to avoid talking about linguistic understanding for two important reasons. The first is that there is already some rich and interesting literature on linguistic understanding that is by and large tangential to the project at hand. Second, this sort of understanding seems to be a prerequisite for knowledge, or, at least, for knowledge that is more than just a mere 'parroting back' of information. In order to know, for example, that Peano Arithmetic is incomplete, I need to be able to comprehend this statement. But the kind of understanding I'm interested in here is something over and above knowledge and not a prerequisite for it. If it turns out that ampliative understanding contributes to the conversation on this front, then so much the better. But this is not a goal I have in mind. The same goes for what we might call *conceptual* understanding, such as a basic understanding of causation or of fragility (i.e. that a glass will break if it falls). With these clarificatory remarks out of the way, let's look at what a successful account of understanding should accomplish.

1.3 The Desiderata for an Account of Understanding

The desiderata for an account of understanding aren't all that different from an account of knowledge; we want to get the cases right and account for the value of understanding. There are, however, some important differences when it comes to how we meet these desiderata. With an account of knowledge, getting the cases right comes down to ascribing knowledge to an agent when it's appropriate and withholding such an ascription otherwise. With understanding, getting the cases right is more complicated

than simply ascribing understanding to an agent.

For starters, understanding comes in degrees. We can make sense of the claim that Sally understands quantum mechanics better than Sam, so our account of understanding should be able to explain (or at least accommodate) this difference in degree. There are, of course, cases in which an agent simply lacks understanding; but we also have cases of *misunderstanding* that our account should also be able to capture. Very little has been said in the literature on this notion of misunderstanding, but it does seem to be an important class of cases in which an agent attempts to understand some topic or state of affairs but, for whatever reason, fails to do so. In a straightforward sense, these sorts of cases fall under the category of a lack of understanding. After all, an agent who tries and fails to understand lacks understanding. But an account that could explain what's going on in cases of misunderstanding would have more going for it than one that was unable to capture this distinction. So instead of having two kinds of cases to deal with—S knows and S doesn't know—we have several. There are varying degrees of understanding to account for, a lack of understanding and misunderstanding.

As for the desideratum of accounting for the value of understanding, there are several approaches we can take. The strongest approach would work to develop the value of understanding over knowledge. Given the current state of the literature, this is an unrealistic goal. Instead, we will focus solely on the final value of understanding as a cognitive achievement. Virtue-theoretic accounts of knowledge begin with this claim but, as we will see in Chapter 2, even this claim is problematic. Thus, we cannot compare the value of understanding to the value of knowledge when the latter has not been definitively established. We will show (in Chapter 4) that the final value of understanding as an achievement can be defended from the current objections in the literature. But we would also need to show that knowledge lacks this value in order to adequately demonstrate the truth of the strongest approach (i.e. that understanding is more valuable than knowledge).

Since the virtue-theoretic approach starts by looking at the value question, we will begin our project there. But we will start by looking at some of the problems with accounting for the value of knowledge in order to motivate the project at hand.

2

From Knowledge to Understanding

The goal of this chapter is to provide some motivation for making a shift in focus away from knowledge and towards understanding. While the past decade has seen a surge in the literature on understanding, this only indicates a growing interest on the topic. It would be too quick to say that, because there is interest in a topic, that it is something worth thinking about. We can think of this transition to a focus on understanding as part of a broader narrative that starts with some general problems for providing a satisfying account of knowledge. Our first task will be to discuss a general difficulty for knowledge: making the adjustments needed for an account of knowledge to be compelling in the post-Gettier era of epistemology while preserving the intuition that knowledge is a valuable epistemic state. We will frame this discussion in terms of these two goals to show the difficulty of meeting them both; it seems the closer we get to one goal, the farther away we get from the other.

In §2.2, we will discuss the conditions under which an account of understanding can meet these knowledge goals. It may not be obvious that an account of understanding should seek to give a satisfying answer to goals that are relevant to a theory of knowledge; but if we take a step back and view the shift to understanding as a broad narrative, it will be clear that, in order to advance the dialectic of the narrative, we should seek out epistemic states, such as understanding, that can meet these goals.

We will then explore, in §2.3, how understanding is in a better position than knowledge to provide compelling answers to these two epistemic goals as well as some of the problems we should expect to face along the way.

2.1 The accuracy and value goals

I suspect that most people reading this are familiar with the challenges that a satisfying theory of knowledge must overcome. But to properly motivate the project on offer, it's important to see where understanding fits in the landscape of the traditional goals of

epistemology. At the very least, we want an account of knowledge that can (1) get the cases right and (2) account for value of knowledge. Call these the *accuracy goal* and the *value goal*, respectively. Epistemologists rightly want to achieve both goals in providing an account of knowledge. What we tend to find in the literature, however, is that certain accounts will use one or the other as a primary focus or starting point. This is similar (at least, in spirit) to the classic *problem of the criterion*: do we start with putative cases of knowledge to see what they have in common or do we work to develop a theory of knowledge that we can apply on a case-by-case basis. I want to use this accuracy vs. value distinction to highlight the difficulty for developing a theory of knowledge that meets both goals. This, in turn, will provide some motivation for shifting our attention to understanding. If we run into problems when trying to meet both goals for an account of knowledge, then it would be worthwhile to consider how other epistemic states such as understanding fare on this front.

In the post-Gettier era of epistemology, getting the cases right (i.e. ascribing knowledge to an agent or withholding such ascriptions as appropriate) has become a significant challenge. Early approaches sought to debunk Gettier's counterexamples to the JTB account of knowledge. One such example was the "no false lemmas" approach. In Gettier's original paper, the cases involved agents who would reach a true, justified belief—but only by including a false belief or assumption in their reasoning. One case that Gettier (1963) considers is a man, Smith, who is applying for a job. Smith believes that another man, Jones, will get the job and also believes that Jones has 10 coins in his pocket. Putting these two beliefs together, Smith believes that the man who will get the job has 10 coins in his pocket. It turns out that Smith himself gets the job, at which point he realizes that his belief was true—Smith has 10 coins in his pocket. The "no false lemmas" move notes that Smith used a false belief in his chain of reasoning; he believed that Jones was going to get the job. Thus, we wouldn't ascribe knowledge to Smith on this account because of a false step in his inference.

As the literature on Gettier cases has grown, so too have the counterexamples to this move—and others—that attempt to deal with Gettier cases. Take the case of the sheep-shaped rock. In this scenario, a farmer, let's call him Rocky, sees what appears to be a sheep standing at the top of the hill. Rocky forms the subsequent belief that there is a sheep on the hill. What Rocky is actually looking at is a rock that, due to erosion, looks very much like a sheep. It would seem, then, that Rocky has a false belief, except for one detail: behind that rock, just out of Rocky's view, is a sheep. The end result is that Rocky

has a true belief that's seemingly justified (after all, he's looking at what he thinks is a sheep), but that falls short of knowledge.

The standard contemporary diagnosis of these kinds of cases is that they introduce *epistemic luck* into the mix. For theories of knowledge focused on the accuracy goal, the primary move is to make an adjustment to the traditional justified-true-belief (JTB) account of knowledge to handle cases like this. The kinds of adjustments we find in the literature are (unsurprisingly) quite varied, but the core idea is basically the same. To meet the accuracy goal, we need to introduce an additional constraint on knowledge, which is typically referred to as an anti-Gettier, or anti-luck, condition. We tend to find this extra condition either tacked on to the traditional tripartite account of knowledge or we will find an attempt to handle epistemic luck by amending some aspect of the JTB account (often it's justification that is amended). In Chapter 3, we will spend more time cashing out this notion of epistemic luck; what's important to understand in this discussion is that, while an anti-Gettier condition may end up helping a theory of knowledge meet the accuracy goal, it will have a tough time accounting for the value of knowledge. In other words, it's a difficult balancing act for a theory of knowledge to deal with epistemic luck and to preserve the long-held intuition that knowledge has value. To see why this is, let's take a closer look at the value goal.

When we talk about the value of knowledge, there are several ways to frame the question. We might ask what value knowledge has over mere true belief, in which case justification (or, at least, some accounts of justification) might be able to explain the difference in value. Plato uses this approach in his *Meno*, arguing that knowledge is "tethered" in a way that mere opinion is not. Fast forward a few millennia, and this tethering comes down to some form of justification. Pritchard (2010) calls this the *primary* value problem, and it is the most basic version of the value problem for knowledge. But even answering the primary value problem can be difficult on some accounts.

Take reliabilism (broadly construed) as an example. The overall strategy here is to develop an account of justification that can adequately handle the kind of epistemic luck present in Gettier (and Gettier-like) cases. At the core of the reliabilist approach is a focus on whether the agent in question has used some reliable process or method in forming her belief. This seems like a very attractive move—especially when our primary focus is on the accuracy goal. Many agents who (through a variety of thought experiments) find themselves arriving a true, but epistemically lucky, belief do so because they haven't

stuck with a reliable epistemic process. A reliable method or process, in this context, is one that will be objectively likely to lead an agent to true beliefs.

We can set aside the potential worries for the reliabilist approach because I merely want to emphasize the tension between our two goals (accuracy and value). Bringing in reliabilism introduces what's called the *swamping problem*: the fact that a belief is true *swamps* any value that reliabilism might bring to the table. As Richard Swinburne (2009: 58) puts it:

So long as the belief is true, the fact that the process which produced it usually produces true beliefs does not seem to make that belief any more worth having.

Linda Zagzebski (2004) explains the problem by having us consider two equally good cups of espresso – one made from a top-of-the-line espresso machine and the other made by a cheap knock-off that just happened to produce an amazing cup of espresso. If all you care about is drinking a good cup of espresso, then it doesn't matter which cup you choose. In other words, there's no additional value in the cup that was made by a machine that is objectively more likely to produce good espresso. In the same way, if all we care about is true belief, then a true belief arrived at through luck seems no more value than one that is arrived at through a meticulous, reliable process or method. And this isn't just a problem for reliabilism—any account of justification that is cashed out in terms of the agent belief's being more likely to be true will encounter this problem. The main worry here is that the value of the true belief *swamps* any value that objective likelihood of truth might bring to the table.

Think about it like this: you have the choice between two bags of money: bag A and bag B. Both contain £1 million in cash—this much you're sure of. But, by the way, bag B is objectively more likely to contain £1 million. Since you already know that each bag has £1 million in it, you have no reason to prefer bag B to bag A; it's not more valuable in virtue of being objectively likely to be valuable. If all you care about is money, then these bags have the same value; just grab one and get out of there before the person giving you this choice changes their mind!

The above examples are not meant to suggest that we don't (or shouldn't) care about the etiology of our beliefs—that is, the method or process through which our beliefs are formed. Instead, they illustrate the tension we mentioned earlier between accuracy and value when it comes to developing a theory of knowledge. The key feature

of these accuracy approaches is that they are looking to add to or alter some aspect of the traditional JTB account of knowledge to handle cases involving epistemic luck. Whatever this move turns out to be, it is likely not going to have the resources to properly address the value problem for knowledge. And this is because there is no reason to think this extra condition—which distinguishes unreliable true belief from knowledge—is going to add anything of value beyond what is already captured in the notion of a true belief.

The flipside to the accuracy approach, which we'll call the *value approach*, focuses on capturing the value of knowledge and working from there to develop a full theory of knowledge. This so-called *value turn* in epistemology is the hallmark of *virtue epistemology*. And if my earlier claim holds true, then what we'll see here is a difficulty in meeting the accuracy condition—that is, in getting the cases right. While it's difficult to lump the variety of virtue epistemic approaches together, they do have one key feature in common: rather than dealing with the primary value question for knowledge (i.e. is knowledge more valuable than true belief?), virtue epistemology starts with the intuition that knowledge is valuable. Given the amount of focus that knowledge has received, this intuition seems to have a strong pull. Put more properly, the virtue-epistemic approach starts with the idea that knowledge is valuable for its own sake—that is, knowledge has *final value*.

A typical move in the virtue-epistemic approach is to locate the value of knowledge as a kind of achievement (in particular, a cognitive achievement). We will spend more time talking about the nature and value of achievements in the next chapter; for now, the basic idea behind virtue epistemology is that an agent who knows that P has arrived at a true belief because of her cognitive abilities.² Virtue epistemology provides a simple yet elegant approach by relying on our intuitions about achievements. The argument in support of virtue epistemology is (on the face of it) pretty straightforward. Achievements are *finally valuable*—that is, valuable for their own sake. Hitting a home run in baseball, scoring a goal in the World cup, or even just winning a poker game amongst friends: all of these are successes that are because of one's abilities. Regardless of what other sources of value attach to these successes (e.g. monetary awards or fame), the intuition is that these achievements are valuable for their own sake. By focusing on knowledge as a kind of achievement, we get the conclusion that knowledge has final value

² The 'because-of' relation here can be spelled out in a variety of ways. Since our discussion here is to simply get an idea of the issues surrounding the value problem for knowledge, nothing is going to hang on how we spell out this relation. We'll save unpacking this concept for the [next] chapter.

qua achievement. For the sake of clarity, here is the argument in premise-conclusion form.

An argument for the final value of knowledge

P1: Achievements have final value.

P2: Knowledge is a kind of achievement.

C: Therefore, knowledge has final value *qua* achievement

Assuming the argument goes through (a claim we will thoroughly examine later on), we now have a clear answer to the value goal for knowledge. Even better, this approach might be just what we need to avoid the above problems of epistemic luck. In classic Gettier-style cases, an agent does arrive at a true belief, but it's not because of her cognitive abilities, it's because of epistemic luck.

Let's go back to the case of Rocky, who has a justified true belief that there is a sheep on the hill. Since Rocky is actually looking at a rock, it's safe to say that his true belief is due to luck rather than any abilities he possesses (cognitive or otherwise). Since knowledge requires the right sort of relation between an agent's cognitive abilities and her true belief, Rocky isn't in the market for knowledge. So far, the virtue-epistemic approach seems to be in a stronger position than approaches that are focused on accuracy; virtue epistemology captures the final value of knowledge and ends up getting a few check marks on the accuracy score card for free. We don't need to add any extra conditions to the virtue-epistemic account to handle these kinds of cases since the agent's belief is not a cognitive achievement. But the value approach runs to problems when we consider different kinds of cases involving what's called *environmental luck*.

In Gettier cases, the epistemic luck is often called *intervening luck* because it, according to Pritchard (2009) "intervenes betwixt belief and fact". This is why the virtue-epistemic approach handles intervening luck so well—had the luck not been present (i.e. had the sheep not been on the hill or had Smith not had 10 coins in his pocket) then the belief would have turned out to be false. The luck in these cases intervenes in such a way that the agent arrives at a true belief, rather than a false one. But environmental luck is a different creature entirely; in cases involving environmental luck, the agent is in an unfriendly environment (i.e. one in which she could have easily had a false belief, but didn't). A classic example of this comes from Carl Ginet and further developed by Alvin Goldman (1976). The short version is that you're driving along a rural road and on either side of the road are lovely-looking barns. You decide to pull over to admire one of these barns. You now have a belief that you're looking at a barn, and you're right—you are

looking at a barn. But what you don't realize is that you're driving through barn façade county. In order to make the region look more prosperous than it actually is, the locals have erected a ton of fake barns that, from the side of road, look completely convincing. You've managed to stop in front of one the few actual barns in this county, so your belief is true; but you're in what we would call an *epistemically unfriendly environment*. That is, your belief that there's a barn in front of you could have very easily been false.

The problem now for the virtue-epistemic approach is that your true belief is the result of your cognitive abilities; but because of the luck involved here, we would hesitate to ascribe knowledge in this case. If the virtue-epistemic approach is to handle cases like this, it will need to add an anti-luck condition, much like the accuracy approach we looked at earlier. And, like earlier, the addition of this anti-luck condition doesn't add anything of value of the account of knowledge. Even though we started off by focusing on the value of knowledge, the need to handle cases of environmental luck forced us to add an anti-luck condition. Thus we cannot capture the value of knowledge, which was the primary motivation of the value approach!

There is an additional worry for the value approach that is even more troubling. Recall that the key move behind the virtue-epistemic approach is to capture the final value of knowledge as an achievement. But what about ordinary cases of knowledge that don't seem to require any real effort or ability. For example, I believe that the lights are on in my office. While my belief is the result of my abilities, it hardly seems like the sort of thing we would classify as an achievement. We will discuss this worry of so-called *easy knowledge* in the next chapter; what's important is again to notice the tension between the goals of accuracy and value. In labeling knowledge as a kind of finally valuable achievement, we might have excluded a wide class of cases of putative knowledge since they don't seem to intuitively be achievements.

It's important to remember that the above discussion is meant simply to motivate a shift in focus to understanding; it would be too quick to throw up our hands and call the above worries intractable. There are genuine worries in the literature that any anti-luck or anti-Gettier condition for knowledge would have to be so complex that the resulting analysis would fail to capture the intuitive value of knowledge. As Timothy Williamson puts it: "This importance [of knowledge] would be hard to understand if the concept *knows* were the more or less ad hoc sprawl that analyses have had to become;

why should we care so much about *that?*”³ Williamson may have a point, but it’s far from obvious that the *complexity* of an account of knowledge is inversely proportional to its capacity to account for the *value* of knowledge.

Our takeaway from this is subtler. Knowledge does seem like the sort of thing that has value; but perhaps this is because knowledge often goes hand-in-hand with understanding. Pritchard (2010) picks up on this idea when discussing the final value of knowledge, noting “... there is a kind of epistemic standing—a species of understanding—which is finally valuable and which also tends to coincide with knowledge. This fact can thus help to explain why when we initially think about knowledge we tend to suppose it is finally valuable” (p. 63). Kvanving (2003)—which serves as an important motivation for this project—makes this move (shifting focus to understanding) because of the problems we’ve discussed for knowledge. What’s important for our purposes is that we can use these problems to motivate a discussion of understanding, which was precisely my goal in this section. We will ultimately want to see why understanding is in a better position to address these worries than knowledge. But first, we need to see how the accuracy and value goals are applicable to an account of understanding. Translating these goals over to an entirely different epistemic state turns out to be a bit tricky.

2.2 Translating these goals to understanding

So far, we’ve seen that providing an account of knowledge that gets the cases right and can account for its value is a difficult balancing act. But to complete the move to understanding, we need to show that it is in a better position than knowledge to meet the worries from the previous section. This is trickier than it sounds because the twin goals of accuracy and value are a bit more complicated when it comes to giving an account of understanding. We will need to spend some time thinking about what it means for an account of understanding to meet these goals.

Let’s start with the goal of value. As a starting point, we can think back to what this goal looks like for an account of knowledge. We mentioned the *primary* value problem for knowledge—that is, whether knowledge is valuable than mere true belief. Reliabilist approaches to knowledge ended up having some difficulty answering this

³ Williamson (2001) pp. 30-31 (*italics original*).

question, but this is only one way of inquiring about the value of knowledge. We could also ask whether knowledge is more valuable than that which falls short of knowledge (the *secondary value problem*). An example here is a justified true belief that is epistemically lucky, and so fails to qualify as knowledge. We saw in the previous section that the addition of an anti-luck condition is needed to keep cases involving epistemic luck from being in the market for knowledge; but the anti-luck condition fails to provide any additional value over that of a justified true belief that is lucky.

Trying to translate these value problems into the realm of understanding, however, brings in some potential hazards. Consider the primary value problem. One might be tempted to say that understanding is more valuable than mere true belief because understanding involves a network, or body, of beliefs. In other words, understanding has more value because it involves more true beliefs—an agent with understanding will, in virtue of her understanding, have more than one true belief. But this answer completely misses the point of the primary value problem.

When we are comparing mere true belief to knowledge, the assumption is that we're comparing the same belief. In the classic example from the *Meno*, the belief in question involves the correct way to Larissa. We're then asked to compare someone who *knows* the way to Larissa to someone who only has a true belief about the correct way. The key question is what separates these two agents, given that they (effectively) have the same true belief. A satisfying answer to this question will refer to some feature that confers value on knowledge, which is lacking in the mere true belief. Simply stating that an agent who understands has more true beliefs misses this important aspect of the primary value problem. We've said nothing about the value-conferring features of a belief that is involved in understanding—only that there are more true beliefs.

A better approach may be to talk about the connectedness of the target belief within the overall network of beliefs. Using this approach, we can make a more legitimate comparison between a mere true belief and that same belief integrated into an agent's understanding. Here, we can bring in epistemically valuable features that the belief will have: it might justify (or be justified by) other beliefs or stand in some coherence relation with related beliefs that are involved in an agent's understanding. These seem like value-conferring features that would make a belief that is part of an agent's understanding more valuable than a mere true belief. But we still have the worry that, by talking about the relational properties of a particular belief, we are making the same mistake as before: we are bringing in additional true beliefs when accounting for the value of understanding,

albeit through the back door.

Answering the secondary value problem as it relates to understanding brings its own set of challenges. Here, we're asking whether understanding is more valuable than other epistemic states that fall short of understanding. In the previous chapter, we stipulated that the kind of understanding we are focused on here—ampliative understanding—is something over-and-above knowledge. But we left hanging the question of exactly what this amounts to because of the myriad of unanswered questions that we would need to know before properly cashing out this idea. Suppose, for example, that understanding is *not* a species of knowledge. Our answer to the secondary value may involve value that is different in kind from knowledge; in other words, the value of knowledge may not be commensurable to that of understanding. We might also find that understanding does not break down as neatly into constituent parts as knowledge does.

Our only other option is answering the *tertiary value problem*—accounting for the final value of understanding. And this is something we will be well-equipped to answer. We noted in the introduction that the approach on offer will be motivated and informed by the virtue-epistemic approach to knowledge. This means we will focus on the final value of understanding as a cognitive achievement, a claim that will of course need to be defended. The reader may wonder here why we are relying on the virtue-epistemic approach, given the problems noted in the previous section. It turns out that the problems we saw for knowledge won't be nearly as problematic when it comes to understanding. The next section will give a preliminary explanation for why this is.

With the value goal a little clearer, we can turn our attention to the accuracy goal. Remember that, for knowledge, this amounted to getting the cases right—ascribing or withholding knowledge as appropriate. So far, we can translate the accuracy goal to understanding in a straightforward way; it's just a matter of ascribing understanding when and only when it is appropriate to do so. But there are other features of understanding that an adequate theory will need to address—in particular, the fact that understanding comes in degrees. We'll need to be able to make sense of the idea that an agent can have a better (or worse) understanding than another agent. In other words, our account needs to be able to handle claims like the following: Sarah understands astronomy better than Jim. So we'll need to ascribe understanding when appropriate, but also say something about the *degree* of the agent's understanding.

It's also worth noting here that understanding is not *binary* like knowledge is—that is, you either know that P or you don't. An account of knowledge need not distinguish

between, say, being misinformed about P (i.e. being told that P is false when it is, in fact, true) or believing that P, but failing to know it (whether because one's belief is unjustified, lucky, etc.). In both cases, we simply withhold a knowledge ascription and move on. When it comes to understanding, however, there are a few different things that can go wrong that a useful account should be able to distinguish. In particular, an accurate account of understanding should be able to distinguish between (1) an agent who has tried, but failed, to understand something, and (2) an agent who has gotten the wrong sort of understanding. In other words, we want to distinguish between a general *failure* to understand and a *misunderstanding*.

If we think about understanding as making appropriate connections between pieces of information, we can get an idea of what this answer will look like. A failure to understand will be a case in which an agent is unable to make a sufficient number of these connections to be in the market for understanding. For an agent to misunderstand some concept or topic, she will have made connections between pieces of information, but not in the right way. As an example, consider two astronomy students learning about the process of hydrogen fusion that fuels stars, including our Sun. There is a lot of information here to connect, including the actual fusion process and the process by which energy is produced during fusion. We can imagine a student who simply cannot put these connections together at all—he is completely lost when it comes to this topic. Compare this to a student who recognizes that energy and fusion are involved and incorrectly believes that the energy from a star pushes hydrogen atoms together to form helium atoms. She has the pieces there, but has things backwards in a sense—the fusion leads to a release of energy, not the other way around. In this case, she has *misunderstood* the process; she has made connections between the pieces of information, but not in the right way. To fully meet the goal of accuracy, our account of understanding should have a principled means of distinguishing between these two cases.

In summary, the account of understanding we are developing should meeting the following criteria:

1. Account for the final value of understanding as an achievement
2. Ascribe understanding to an agent when (and only when) such ascriptions are appropriate
3. Account for the degrees of understanding
4. Distinguish between a general *failure to understand* and a *misunderstanding*

Notice we still have our two broad goals of value and accuracy, with criteria 2-4 indicating the different levels of accuracy involved in an account of understanding. Now that we have translated these twin goals for knowledge over to understanding, we will quickly look at how understanding is in a better position to reach these goals than knowledge.

2.3 The motivation for understanding

We saw that, with the virtue-epistemic approach to knowledge, the accuracy and value goals are closely tied together. The key move in this approach is to classify knowledge as a kind of achievement. This move ended up helping with the goal of accuracy because in Gettier cases—cases involving intervening luck—the agent’s success is not because of her ability. As a result, the virtue-epistemic approach to knowledge can handle these cases for free (that is, without adding any additional constraints or criteria to the account). Since we are taking a virtue-epistemic approach to understanding, we are already on equal footing with knowledge on this front.

One worry for knowledge on the virtue-epistemic account was environmental luck—cases in which an agent is in an epistemically unfriendly environment. We used the barn façade case to help illustrate this point. In cases involving environmental luck, an agent arrives at a true belief because of her cognitive abilities—after all, in the barn façade case you are looking at a real barn and form the relevant true belief. These sorts of cases motivated the introduction of an anti-luck condition, which ends up defeating the claim that knowledge has unique value as an epistemic state qua achievement. Instead, when environmental luck is in play, we have cases of true belief that fall short of knowledge but would nonetheless be classified as achievements. Thus, knowledge loses its unique value as a cognitive achievement, which would have set it apart from epistemic states that fall short of knowledge.

What gives understanding an advantage here is that, unlike knowledge, it is not defeated by environmental luck. In other words, agents in epistemically unfriendly environments are still in the market for understanding. We will need to spend time further developing and defending this claim, but if the argument goes through then we will have no need for an anti-luck condition to handle cases of environmental luck. The upshot is that we can preserve the unique value of understanding as an epistemic state; agents with understanding will have demonstrated a finally valuable cognitive

achievement that is not present in cases where an agent fails to understand.

There is a worry for the value understanding that will be a bit more difficult to avoid. Earlier, we mentioned the problem of easy knowledge—cases in which an agent has knowledge but hasn't demonstrated an achievement (at least, of the sort that would be valuable for its own sake). Much of our ordinary knowledge may fall into this category; for example, perceptual and memory-based knowledge don't seem to require much (if any) cognitive effort on the part of the agent. By starting with the claim that knowledge has final value *qua* achievement, the worry is that these ordinary cases of knowledge would be excluded from a virtue-epistemic approach. There is a similar worry—developed by several different authors—that certain instances of understanding are too easy to qualify as an achievement.⁴

We will explore this issue in detail in Chapter 4, but I want to quickly touch on the move we'll be making to address this worry. When we are talking about knowledge as an achievement, the success in question is the agent arriving at a true (rather than false) belief. We've already noted that ampliative understanding involves on a network of beliefs, and the relevant success here will *not* be acquiring a true belief. Instead, we should focus on whether the agent has made the right sorts of connections in this body of information. In short, this involves an active cognitive engagement on the part of the agent, rather than the passive belief formation we find in knowledge based on perception and memory.

We can also use this move to address a potential worry about meeting the accuracy goal for understanding. Recall that we listed two additional criteria (in addition to getting the cases right) that a satisfying account of understand should fulfill: (1) accounting for varying degrees of understanding, and (2) being able to distinguish between a failure to understand and a misunderstanding. Since our focus will be on whether the agent has made the right sorts of connections, we have a lot more leeway when it comes to ascribing or withholding understanding. This is because of (1) above: some agents may have a very low degree of—but still be in the market for—understanding. Additional flexibility on this front has the potential to offer more opportunities than it does challenges. The same line of thought can also work for (2); more ways to fail to meet the epistemic standards of understanding provides us with more information for what can go wrong, but can also help inform ways in which an agent

⁴ See Morris (2012), and Whiting (2012), and DePaul & Grimm (2007) for a some good discussions of this general worry.

can get things right. In other words, these additional criteria do offer unique challenges, but they also provide us with more freedom, flexibility and information in this project.

Added flexibility, on the other hand, can have its drawbacks. In particular, we will need to make sure there is a clear cutoff point between an agent who has a very poor understanding of a subject or concept and an agent who has simply failed to understand the relevant material. At this point in the project, it isn't clear whether understanding is on a sort of sliding scale with a flawless understanding at the top of the slope that slowly degrades to a complete lack of understanding. An alternative to this view would be a principled distinction between agents who barely qualify as understanding and agents who lack understanding. Once our account is more developed, we will argue for the latter view, which provides for the clear cutoff in question.

2.4 Concluding remarks

The primary goal of this section has been to motivate a shift in focus away from epistemological accounts of knowledge and to understanding. We began by identifying two important goals that a theory of knowledge should meet—the accuracy and value goals—and showing the difficulty that many theories of knowledge have in meeting both goals. Again, my argument here is not that the issue is intractable or that epistemologists should stop working to develop satisfying account of knowledge—I simply want to show that an account of understanding may be in a better position to meet these goals. Besides being a topic that has received more attention in recent years, understanding as an epistemological concept provides us with moves that can account for its value and can shed more light on a topic that is still in its infancy (at least, relative to other, more traditional topics in epistemology).

We now have some clear criteria that a satisfying account of understanding should meet as well as some general approaches on how ampliative understanding can meet these goals and criteria. The next move from here will be to spend some time discussing the nature and value of achievements, as well as their relationship to luck. We've mentioned both intervening and environmental luck in this chapter, but only briefly. Because ampliative understanding is motivated by the virtue-epistemic approach, it will be absolutely critical to have a clear idea of achievements (specifically, those achievements which are finally valuable), how they relate to both environmental and epistemic luck and how, in turn, this relates back to understanding.

Fortunately, the literature on achievements and final value is quite well-developed. Unfortunately, there are some flaws in the discussion that we will need to address that will be especially important for a value-based account of understanding. There are a number of moves and views to consider, so the next chapter is dedicated to furthering the dialectic on achievements and final value.

3

Two Criteria for Achievement

In the previous chapter, we spent time developing the virtue-theoretic approach to knowledge and relating those goals to understanding. In particular, we saw that accounts of knowledge (whether virtue-based or not) would have a difficult time balancing the twin goals of accuracy (getting the cases right) and value (accounting for the value of knowledge). From what we've seen, understanding has at least a slight advantage over knowledge on this front; but we need a more thorough discussion of achievements in order to show that understanding doesn't face the same pitfalls as knowledge. The aim of this chapter is to get a clear idea about the nature of achievements, their relation to epistemic luck, and their applicability to knowledge and understanding. Remember, the key move behind the virtue-epistemic approach to understanding is developing and defending the claim that understanding has final value *qua* achievement.

We will begin by developing two key aspects of achievements. The first, what we'll call the ability-success relation, ensures that the agent's success is connected to her abilities in the right way. The second, what we'll call the difficulty condition, is intended to narrow the scope and identify only those successes which are, in a sense, difficult enough to intuitively qualify as achievements. Once we meet both of these key aspects, we will have an account of achievements that can capture both their nature and value. We'll then use this discussion to help clear up some muddy points in the current literature. The goal here is not to resolve disputes that are outside the ambit of the project at hand; instead, I want to use some of the debates to highlight moves that will be important for our purposes—specifically when it comes to defending the final value of understanding.

3.1 The ability-success relation

So far, our discussion of what an achievement is has been brief; we relied on the idea that an agent who is successful because of luck—rather than her abilities—is not in the market for achievement. This is the line that Pritchard (2010: 70) takes to develop what

he calls the *achievement thesis*.

Achievement thesis: Achievements are successes that are because of ability.

The key feature of the achievement thesis is that the agent's success is connected to her abilities in the right way. This allows accounts of knowledge (and understanding) to avoid a certain kind of epistemic luck (what is called Gettier-style intervening luck). Thus, achievements require what we'll call an *ability-success relation*—that is, an agent's success must be connected to her abilities in the right way. Relating this back to our two goals of accuracy and value, the Achievement Thesis falls squarely in the accuracy goal. After all, it's what allows us to get the cases right when intervening luck is involved (that is, to withhold knowledge or understanding in those cases).

The primary issue at hand is what to make of this 'because of' relation. In the last chapter, for the sake of brevity, we held off on unpacking this relation because not much hung on it at the time. But we're at the point now where clarity and precision are vital, especially when we relate these principles to ampliative understanding. There has been some debate on how to read the 'because of' relation—or, what we're calling the ability-success relation. This debate has focused solely on knowledge (i.e. how to read the 'because of' relation in a way that gets the cases right for knowledge), but we can use that debate to help inform the ability-success relation as it relates to understanding. With that in mind, let's look at a few of the more prominent suggestions on this front.

One approach is to say that the agent is *creditable* with the success; on this view, the agent, because of her abilities, is at least partly responsible for the success in question. This is an extremely weak approach because an agent can be almost completely uninvolved in a success and still be creditable with the success, at least to a degree. Imagine a tourist who, upon arriving in a new city for the first time, asks the first person he comes across directions to the nearest butcher shop.⁵ He forms a true belief about which way to go, but in this case, the stranger had a much bigger role to play in the tourist's true belief. We should be wary of any account of achievement that allows someone to demonstrate an achievement when they are barely involved in the process.

A stronger relationship between success and ability can be found in John Greco's (2003) approach, which argues that that the ability-success relation entails credit—that

⁵ This is a slightly altered version of Jennifer Lackey's (2007) case against credit-worthiness and knowledge.

is, the agent is praiseworthy for the success in question. Using the tourist example, we may be able to make the case that our tourist is praiseworthy, but it would be a difficult battle. The tourist did, after all, approach a person, rather than a dog or a child or someone else who was obviously a tourist. But we might rightfully worry here that the tourist can hardly be said to be deserving of credit, and certainly isn't praiseworthy. To be fair, the tourist could have been *blameworthy* for arriving at a false belief, had he ended up doing something silly (like asking a child or a lamppost for directions). But a lack of blameworthiness does not imply praiseworthiness, nor does it imply credit. Remember, we are thinking about the nature of achievements—even if the tourist is deserving of credit in some sense, it is not the sense that is relevant to achievements.

We can respond to this worry by increasing the level of involvement required from the agent in bringing about the success. Specifically, an agent's success should be *primarily creditable* to the agent. This is the line that Greco (2007a; 2007b) takes later on, followed by Pritchard (2010). Viewing the ability-success relation in this way makes intuitive sense and works in a number of cases. Take the case of an archer who hits the bull's-eye because of her abilities; it's natural to read this as saying that the archer is primarily creditable with her success. We also get the case of the lost tourist right. His true belief is primarily creditable to the stranger rather than the tourist. The upshot is that, if there's an achievement in play, it is not the tourist's achievement—it's the stranger's. For our purposes, this reading of the ability-success relation ends up being very close to what we need. As a quick aside, this view of achievement doesn't do much to help with the claim that knowledge is an achievement. In the lost tourist case, we have a putative case of knowledge (albeit, testimonial knowledge) that fails to be an achievement. This ends up being a serious worry for the virtue-epistemic approach and one without a clear solution. In fact, it's cases like this that motivate Pritchard (2010) to shift his focus away from the final value of knowledge and focus on the value of understanding. But for our purposes—capturing the notion of an achievement in terms of the ability-success relation—the view endorsed by Greco and Pritchard works quite well. An agent who is primarily creditable with a success will (at least, intuitively) have a sufficient degree of involvement in a particular success to be in the market for achievement.

I only have one complaint about the above analysis: it focuses on the relationship between the success and the *agent*. In other words, it's an agent-success relation, rather than an ability-success relation. This may seem like more of a semantic worry than a serious worry, but so much of this project hangs on getting the ability-success relation

right that it's worth risking being slightly pedantic in order to ensure an accurate and precise account. Perhaps we can make a simple modification to the above account and say that an achievement is a success that is primarily creditable to the agent's abilities. The worry with a move like this is that we're now giving credit to a set of abilities, which is an odd way to think about credit-worthiness. When we think about something being creditable with (or deserving credit for) a given success, we credit the *agent* with that success. Instead, we would need to say that the agent is primarily creditable with the success *because of her abilities*.

The above modification is very close to Ernest Sosa's (2007) triple-A account of performance-assessment. Sosa's account focuses on three aspect of performances: accuracy, adroitness and aptness. *Accurate* performances are successful ones, such as an archer hitting her target or Tiger Woods sinking a putt. *Adroit* performances are ones that manifest competence. Competence, in this context, is an ability (or set of abilities) that is relevant to the agent's successful outcome. Finally, Sosa brings in *apt* performances, which are accurate because they are adroit. In other words, they are successful because they are competent. Once we put it all this together, the outcome of an apt performance, on Sosa's view, will be an achievement. Put another way, an agent who is in the market for achievements will be successful because she manifested competence. Sosa's account does a nice job of connecting an agent's success with her abilities, rather than with the agent herself. But there is an objection to this account developed by John Turri (2013) that we can utilize to further close in on an adequate notion of the ability-success relation.

Turri has us imagine the battlefield of the Trojan War, with the legendary archer Paris shooting down his enemies from afar. Each successful hit is certainly an achievement—Paris's shots are successful because they manifest competence. But a lot can happen in between Paris shooting at an enemy and the arrow arriving at its target. Suppose, during the arrow's flight, something gets in the way of his intended target (a falling soldier or chariot, for example). This unforeseen obstruction would cause Paris to fail to hit his target. But the god Apollo, who favors Paris, decides to intervene on the condition that Paris's shot manifests competence. In other words, Apollo will guide the arrow to its target if and only if—and because—Paris's shot manifests competence but would have failed to hit the target because of some unforeseen circumstance. In the cases where Apollo intervenes, Paris's success is nonetheless *because* he manifested competence. After all, this is why Apollo intervenes in shots that would otherwise miss. Yet we wouldn't call Paris's success an achievement since Paris is not responsible for his

success—the god Apollo is.

Remember, the claim we started with was that achievements are successes that are because of ability; this is Pritchard's Achievement Thesis. Each of the above accounts were attempts at spelling out this 'because of' relation—whether through some notion of creditability, credit-worthiness, or a causal or explanatory relation. The move that Turri makes, which I endorse (though for completely different motivations) is to do away with the 'because of' relation entirely. Instead, Turri introduces the notion of *adeptness*: “An adept performance is one whose accuracy manifests competence (as opposed to being accurate *merely because* competent)” (2013, p. 4 *italics original*).

Rather than get bogged down with jargon, we can amend Pritchard's Achievement Thesis (AT) to accommodate our findings so far:

(AT*): Achievements are successes that manifest relevant abilities.

Notice that we've removed the mention of competence during our revision of the Achievement Thesis. This move is also to cut down on the 'thickness' (the amount of jargon) in our present analysis. Sosa's notion of competence, which is included in Turri's account of adeptness, is simply those abilities that are relevant to the success in question. So we haven't lost anything of substance in making this change. The main question we have now is how to cash out this idea of manifesting abilities. I want to treat the term 'manifests' here as primitive; but with my focus on developing a clear and accurate account of achievements, a healthy dose of skepticism with respect to this move is warranted. Let's quickly look at why we can take a term like 'manifests' as primitive but not make the same move with the 'because of' relation.

The first thing to note about the 'because of' relation is that there are plenty of plausible ways of cashing out this term that can be both meaningful and accurate—at least within a given domain. There is a wide class of cases, for example, in which we can use 'because of' in an explanatory sense. As long as our domain is one of explanation, then we can make claims of the general form “X happened because of Y” without any ambiguity in the 'because of' relation here. In other words, there are clear and legitimate ways of cashing out 'because of' that can be sorted through when thinking about the Achievement Thesis. This is precisely what this section has been trying to do—find the correct way (amongst many legitimate interpretations) of cashing out the 'because of' relation when it comes to the Achievement Thesis. To manifest an ability, on the other hand, does not have the same reductive features we see in the 'because of' relation (even

in context-specific domains) and yet we seem to have a reliable intuition about when an ability is manifested and when it's not. As Turri puts it, "we almost certainly cannot specify in general when an outcome manifests a disposition. But case by case we find it relatively easy to tell..." (2013:3). While our overall assessment of an achievement, then, will be informed by our intuitions about whether an ability is manifested, we still end up with a clear analysis of achievements that, because of our reliance on intuitions, will have a high degree of accuracy when it comes to whether a given success qualifies as an achievement.

With the modification of the Achievement Thesis, we have our ability-success relation. This relation is one in which the agent's success manifests relevant abilities. But our analysis so far has only provided half the story—an accurate account of achievements must also be one in which those achievements are, in fact, finally valuable (i.e. valuable for their own sake). The ability-success relation only tells us how an agent's success is connected to her abilities, but with just this criteria in place our account of achievements will be too broad. It doesn't exclude those successes which are too easy to have final value. We will need to spend some time narrowing in on just those achievements that are finally valuable.

3.2 The difficulty requirement

Suppose I want a beer, so I get up, walk to the refrigerator, and grab one. This is a success that manifests my relevant beer-grabbing abilities. In other words, it meets the criterion of the ability-success relation. But we would hardly call this sort of success an achievement—or, at least, not one that is valuable for its own sake. This isn't to say that the state of affairs has no value whatsoever. Depending on one's views on value, my getting a beer could have value in terms of increasing my welfare (I'm happier now that I have a beer) or fulfilling my desires. But it lacks the kind of final value as an achievement that we're looking for—it's not valuable for its own sake.

Ordinary successes like this motivate us to bring in a second criterion for achievements, what we will call the *difficulty requirement*. The idea behind this move is that achievements—or, at least the ones that are valuable for their own sake—are going to be sufficiently difficult. Pritchard (2010) identifies these kinds of achievements as *strong achievements* and amends his previous Achievement Thesis so as to exclude successes that are too easy:

Strong Achievement Thesis: Achievements are successes that are because of ability where the success in question either involves the overcoming of a significant obstacle or the exercise of a significant level of ability (2010, p. 70).

The first thing to notice is that his original Achievement Thesis is still present. We've already made a modification to this in order to develop the ability-success relation. The key difference here comes in the latter part of the Strong Achievement Thesis, which requires that the success involves a significant obstacle or a significant level of ability. This added requirement, which we will develop into the difficulty requirement, is a disjunction: a significant obstacle *or* a significant level of ability. Furthermore, we should read this as an *inclusive* disjunction (we can have an achievement that satisfies both disjuncts), for reasons we will see shortly. First, let's take a look at some cases that illustrate what Pritchard has in mind for each of these disjuncts.

Let's go back to the case we started with: someone going to the fridge to get a beer. Under normal circumstances, this success involves neither a significant obstacle nor a significant level of ability. But suppose the individual is recovering from a severe spinal injury, and these are the first steps unaided steps they have taken. Now the success in question (getting a beer from the fridge) involves a significant obstacle-overcoming effort. In this new scenario, we get back the intuition that this is an achievement that is valuable for its own sake. Another way of putting this is that there is a sufficient degree of difficulty in play to elicit the intuition that our injured beer-getter has demonstrated a finally valuable achievement.

On the other side of the disjunct, we find agents who have demonstrated a high degree of ability. A nice example of this sort of achievement is a major-league baseball pitcher who is able to throw a 90-mph fastball. For many pitchers, this is not something they would classify as particularly difficult. Given the number of times in a game that a pitcher is able to consistently do this, there doesn't seem to be any sort of significant obstacle in play. Instead, these are top-level athletes who are demonstrating a significant level of ability. Whereas most of us would find such a feat extremely difficult (if not impossible)—even with the appropriate level of instruction—these pitchers make it look easy, so to speak. But it's because of their significant abilities that such feats do not involve an obstacle-overcoming effort.

All of this suggests a general relationship between the two disjuncts of the difficulty requirement. Adam Carter and Emma Gordon (2014) also notice this

relationship. As they put it, “there is a kind of *inverse* correlation at play between obstacle-overcoming and skill ...” (p. 4 *italics original*). This inverse correlation doesn't hold for all cases; even agents with a high degree of skill can find themselves in situations which still require an obstacle-overcoming effort. A highly-skilled mountain climber will still need to display an obstacle-overcoming effort in order to climb a challenging route. Since most of our everyday achievements, however, don't involve summitting Mt. Everest, we can use this inverse correlation as a general rule of thumb when thinking about the relationship between an agent's skill and the obstacle-overcoming effort in play.

Pritchard recognizes this relationship as well when it comes to the cognitive achievement of understanding:

Typically, after all, one gains understanding by undertaking an obstacle-overcoming effort to piece together the relevant pieces of information. Moreover, where understanding is gained with ease, this will be because of the fact that one is bringing to bear significant cognitive ability (2010: 82-3).

To see how this might work, think about Sherlock Holmes and Dr. Watson investigating a murder. Holmes is able to immediately—and seemingly effortlessly—piece together the information needed to solve the crime. Meanwhile, Dr. Watson (and presumably, the reader of the mystery novel) must struggle to follow Holmes's line of reasoning and, only after a considerable amount of effort, come to understand the state of affairs (e.g. that Mr. White must be the killer).

With this inverse relationship between ability and effort, we may wonder whether there is a point at which an agent's ability is at such a high degree that the success no longer intuitively qualifies as an achievement. In other words, what we're calling the difficulty requirement won't be met by agents with a sufficiently high level of ability. This is an objection that Daniel Whiting brings against Pritchard's Strong Achievement Thesis. Whiting has us imagine a golfer with an amazingly high level of ability when it comes to putting.

One can imagine, as he sinks putt after putt with the most casual swing of the club, his sense of achievement diminishing, perhaps precisely because he misses the challenge. This, in effect, is to suggest that [the strong achievement thesis] is false—a success that is because of ability might count as an achievement if the success in question involves the overcoming of a significant obstacle, but it does not count as an achievement if it

only involves the exercise of a significant level of ability (2012:221)

At first glance, Whiting's objection seems to be focused on the phenomenology of achievement. After all, the golfer in the above scenario has lost the sense of achievement when sinking a putt. But our discussion of achievements makes no mention of any phenomenological requirements, so it may seem that Whiting's objection isn't relevant.

A more charitable reading of the objection is that the above case challenges our intuitions about what is or isn't an achievement. Because so much of our discussion relies on our intuitions about achievements, the worry is that there is a flaw in the difficulty requirement. Remember that this requirement is a disjunction involving an obstacle-overcoming effort or a significant level of ability. If we lose the intuition that our golfer's successful putts are achievements, then it puts pressure on the second disjunct. In other words, the requirement that a significant level of ability is on display isn't doing any work. Or, at the very least, there is a point at which the ability requirement does no work in capturing those achievements with final value. Call this the *ability objection*.

The key feature in the ability objection is that the golfer is unable to appreciate the achievement in his successful putts. So, in a sense, this objection does come down to the phenomenological aspects of achievement, albeit in a way that challenges our intuitions. The first thing to note about the difficulty requirement is that it is going to be both context-sensitive and relativized to the agent(s) involved. Suppose that eight-time Olympic gold medalist, Usain Bolt, decides to show up to a field-day race for 10-year-olds. Bolt, widely considered to be the fastest man alive, obliterates the competition and comes in first place. Besides ruining these kids' race, Bolt manifested a significant level of ability in winning the race. We can add to this scenario that Bolt also feels a huge sense of achievement, taking great pride in his stunning victory.

From an outside perspective (say, as a parent watching this strange set of events unfold), we would rightly think that, in this context, Bolt is being kind of a jerk. His sense of achievement is not deserved, despite the fact that he manifested a significant level of ability. One takeaway from this scenario is that the feeling of achievement can be present even if there is not achievement on offer. Bringing Whiting's objection back in, we can make the case that the phenomenology and ascription of achievement can come apart in both directions. Bolt feels a sense of achievement when it is undeserved and our golfer lacks that sense when there is an achievement in play. In short, relying on a successful agent's frame of mind is not going to be a reliable way to determine whether an achievement is in play. Recognizing this does lessen the force of the ability objection, but

we will need to look a little deeper into the difficulty requirement in order to properly address this worry.

The upshot of these cases is that we have motivation to make a slight revision to the difficulty requirement. So far, we've been saying that, in order to meet this requirement, the success must either involve an obstacle-overcoming effort or a significant level of ability. The key word here is 'involve'. Bolt's success in the race did involve a significant level of ability—we built that into the case. Yet we wouldn't intuitively view his success as an achievement. The reason for this is because Bolt didn't *need* to manifest his athletic abilities in order to win a race against children; pretty much any adult, in virtue of their longer legs and more-developed muscles, could have won a race against children. When thinking about the difficulty requirement, we want something stronger than the mere involvement of the agent's abilities—we want these significant abilities or obstacle-overcoming effort to be *required* for the agent to succeed. Thus, we can revise the difficulty requirement to say that the agent's success requires an obstacle-overcoming effort or significant level of ability.

By revising the difficulty requirement, we can explain what's going on in the case of Usain Bolt racing against children—his abilities weren't required in order to win the race. But it's important to keep in mind that the abilities or effort in play are going to be sensitive to the agent(s) involved and the overall context, a point we alluded to earlier. If we look at a race just amongst children we get the intuition back that the winner will have demonstrated a finally valuable achievement. The worry here is that these are children who are going to lack a significant level of ability. We can make the move to say that what's required now is an obstacle-overcoming effort, which seems intuitively plausible. But the waters are getting muddy now. Keep in mind that our two criteria for achievements are meant to rely on intuition, but a little more precision in these cases can't hurt. What's missing from the above discussion and, in particular, from our response to the ability objection is a clearer idea of the relevant notion of success in play for a given performance.

3.3 Locating the success

So far, we've developed two criteria that a success must meet to be considered an achievement that is valuable for its own sake. The ability-success relation ensures that an agent's success manifests her abilities, rather than resulting from luck or some other

agent's help or intervention. The difficulty requirement ensures that the success in question requires (rather than merely involves) an obstacle-overcoming effort or a significant level of ability. We've already noted that the difficulty requirement is going to be sensitive to the state of affairs in which the success takes place. This ties in to the second crucial point regarding the difficulty requirement: it acts as a constraint on the success component of the ability-success relation. In short, we want the agent's success to be sufficiently difficult (relativized to the agent) so that the success qualifies as an achievement.

As a quick aside, the reader might have noticed that we're focusing heavily on the success component of the ability-success relation without saying much about the ability component (e.g. what constraints should we expect there). There are two reasons for this. First, we have already noted the difficulty in developing a general account for when an agent has manifested her abilities. Instead, this is something we would assess on a case-by-case basis. Second, there is a similar difficulty when it comes to identifying any sufficient or necessary abilities for a given success. Take, for example, a major-league baseball player getting a hit. This is a success that manifests the hitter's abilities, but there are a number of abilities that may or may not be in play. Some hitters are able to 'see the ball' really well (i.e. what type of pitch is being thrown and whether it will be a ball or a strike). Others rely on power, while some have a quick bat and are able to hit the ball into defensive gaps. But none of these abilities by themselves are sufficient or necessary to get a hit in baseball. In short, we need to assess what abilities are in play on a case-by-case basis.

Returning to the ability objection, let's think back to Whiting's case of a golfer who is effortlessly sinking putt after putt. At the core of the objection is that the golfer's high level of ability takes away a sense of achievement. We've already seen that the phenomenological aspects of achievement can come apart in both directions from actually being in the market for achievement. But we may still worry that the ability-disjunct of the difficulty requirement isn't doing any work—or, at least, stops working once an agent reaches a high degree of ability. Put another way, the inverse correlation between ability and obstacle-overcoming effort might, at some point, render the agent's level of ability otiose when assessing a successful performance.

The key aspect of the ability objection is that we have an agent who has become a highly skilled golfer (or, at least, putter). But notice that becoming a highly skilled golfer is itself a success that is fraught with obstacles. Put more broadly, reaching the status as an agent with a significant level of ability is itself a finally valuable achievement. Once we

locate this particular success, answering the ability objection is a straightforward matter. An agent with a high degree of ability may not encounter any significant obstacles in succeeding at a given performance. But this is precisely because the agent reached this level of ability through an obstacle-overcoming effort. If we do away with the ability aspect of the difficulty requirement, we lose out on this explanation for what's going on when highly-skilled agents are successful without facing the same obstacles that a lesser-skilled agent would face. In short, we need the ability requirement to keep the intuition that top-level performers (whether they are athletes, artists or astrophysicists) have demonstrated achievements—even in the absence of significant obstacles.

Whiting's objection has helped us make a novel move when it comes to assessing performances—one that we can apply to ampliative understanding. Recall the key feature of this kind of understanding is that an agent is able to acquire true beliefs. There is a similarity, then, between agents who are proficient or highly-skilled and those with ampliative understanding: both are states that enable an agent to have more successful performances. Athletes, for example, who are proficient will see more success, whether that is sinking a putt, getting a hit in baseball, hitting the bull's-eye, etc. For ampliative understanding, an agent will be in a position to have epistemic success—that is, true beliefs. We will return to this point to develop it further, but we need to say more about properly locating the success when it comes to assessing performances. The real novelty behind our reply to the ability objection is the realization that we can take a more fine-grained approach when looking at a performance to see what other successes or achievements might be in play.

Imagine a soccer player—a striker—whose job it is to push towards the opponent's defensive line in order to set up a breakaway attempt on goal. A player in this role typically has a lot of speed, is able to predict defensive maneuvering well, and is able (along with his teammates) to spot these breakaway opportunities. An ideal situation would be for one of his teammates to send a ball over his head towards the opponent's goal and, using his speed and skill on the ball, the striker would score a goal. There are a number of things that can go right (or wrong!) in this scenario. But let's imagine our striker has a breakaway opportunity, beats the defense and takes a shot on goal. He takes a shot on goal, which is a beautifully struck ball into the upper-left corner of the goal. But the goal keeper makes an amazing dive and saves the ball, which keeps the striker from scoring. So far, we have been looking at performances through a very broad lens—a kind of hit-or-miss analysis (e.g. either the archer hit the target or she didn't). But using a more fine-grained approach, we are in a position to take this scenario and parse it out in such

a way that allows us to identify many achievements, despite the fact that the striker was ultimately unsuccessful.

First, we can look at the striker's breakaway. He had to be in just the right position, reading the defense and timing his run so that he could be at full acceleration the moment the ball was played to him. In our scenario, the player was able to beat the defense, which could be considered an achievement in its own right—independent of whether it resulted in a goal. We could also look at how the player struck the ball. Compare two scenarios: one in which the player places a fairly limp shot directly at the keeper and one in which he hammers the ball in the upper corner. In the scenario we've just considered, the player was able to strike the ball well, forcing the keeper to dive in order to make a save. Notice that, in both scenarios, the striker has taken a shot and the keeper has saved it. But in the latter scenario (where the ball is well-struck), there seems to be more value in the overall state of affairs, despite the outcome being the same. Using our fine-grained approach, we can explain this intuitive difference in value.

Let's start by looking at the keeper blocking the shot. This is his role in the game and in both scenarios (the weak shot straight at him and the strong shot in the upper corner) he is successful. But in the first scenario, just about anyone could have kept the ball out of the goal, while the second clearly involves a high degree of ability or an obstacle-overcoming effort (or both). What's interesting to notice is the relationship between how well the striker kicked the ball and the value that attaches to the keeper's save. A well-struck ball allows room for an achievement on the keeper's part, while a poorly-struck ball doesn't. The upshot is that, with a more fine-grained way of locating the success component of the ability-success relation, we can better account for and explain the value of a given state of affairs. This, in turn, allows us to have a clearer picture of assessment performances and identifying achievements.

The reader may worry at this point that we're just cherry-picking certain outcomes and calling them achievements without principled means of locating a success. Why, for example, did we shift focus from the striker to the keeper and only focus on the latter's success? This is an especially pertinent worry given that our assessment of a performance and our discussion of achievements relies so heavily on our intuitions, which takes place on a case-by-case basis. But there is a guiding principle we can use to identify relevant success in a complicated situation like the one we're considering. By looking at a particular agent's intentions—i.e. what their goal or aim was in a given performance—we can locate the relevant successes and assess these performances to determine whether they are achievements.

There is an obvious sense in which we need to look at an agent's intentions in order to locate the relevant success. For example, an archer who is trying to hit target X but ends up hitting target Y has not demonstrated an achievement. Her goal was to hit target X and, with respect to this goal, there is no relevant notion of success to be found; the archer isn't even in the market for achievement. But what we're talking about here is much narrower than this. Take our (ultimately unsuccessful) striker—in a broad sense we can say that his goal or intention was to score a goal. But in a narrower—and more precise—sense, his goal was to kick the ball hard into the upper corner; and with respect to this goal, the striker was success, even though the outcome did not result in a goal. Viewing his performance in this way, we have room to say that the player has demonstrated a finally valuable achievement. He was successful in meeting this more narrowly defined goal and his doing so manifested his ability and met the difficulty requirement. So, both the striker and keeper have demonstrated finally valuable achievements in this case. Both players met their specific goals and both did so in a way that met the criteria for an achievement.

When we first started thinking about this soccer scenario, we considered two unsuccessful outcomes: an easy shot right at the keeper and a hard shot in the corner of the goal. In both cases, we noted that the striker was unsuccessful in scoring a goal and that the keeper was successful in blocking the shot. Yet there is an intuition that, in the hard shot scenario, there is more value in the state of affairs. We would have a difficult time explaining this intuitive difference in value without using the fine-grained method of performance assessment. Instead of looking at the broad goals of the agents involved (i.e. scoring a goal), we focused in on the intentions of the agents involved and what goals they had in mind in order to reach those broader goals. The striker's achievement in placing the ball right where he wanted provided the opportunity for the keeper to also demonstrate an achievement, which wouldn't be present in the easy-shot scenario.

There may be some hesitation in accepting the above assessment of the performance. After all, we can still say that the keeper has demonstrated a finally valuable achievement in making an amazing save that wouldn't be present in the easy-shot scenario. Thus, we don't need the fine-grained approach to explain the value difference in these two scenarios. But suppose the striker hit the ball hard, but nowhere near his intended target. The wild shot, however, ricocheted off a player and got driven towards the upper corner. In this new scenario, the player is still just as unsuccessful on the broad view of performance—he didn't score a goal. And in both scenarios, the keeper has made the same amazing save. Yet this ricochet scenario intuitively has less value than one in

which the striker hit the ball well and where he wanted to. Again, this brings us back to the fine-grained approach to assessing performances. In our scenario where the striker is successful in the narrow sense, he has demonstrated a finally valuable achievement, which explains the value difference in this ricochet-scenario compared to the hard-shot scenario.

The above will be an especially important move when it comes to the value of understanding. With a more fine-grained way of located successes, we can make sense of an agent coming to understand a false theory like astrology. In short, we will be able to draw a distinction between the aim of understanding the information with which an agent is engaging (which would be their narrow goal) and the broader goal of understanding the world around us. Someone whose aim is to understand astrology might be successful with respect to the former goal but certainly not the latter (since astrology does very little—if anything—to help us understand the world around us). By properly locating the success, we have room to say that an agent is successful in one respect—and perhaps even in the market for achievement—while keeping distinct a broader aim that the agent falls short of.

Besides identifying two criteria of achievements—the ability-success relation and the difficulty requirement—we have made two important moves. First, we have identified an achievement in acquiring the status as an agent with a high level of ability. Even the Tiger Woods, Bobby Fischers and Michael Jordans of the world, who are born with certain innate talents, must train incredibly hard to reach this status. Given the ratio of professionals to amateurs in these activities, it's entirely reasonable to say that the road to acquiring such a status has its fair share of significant obstacles. Along with this move came a more fine-grained way of assessing performances. Both moves are important for identifying achievements and, as we will see, are critical for explaining and assessing the value of an agent's ampliative understanding. In the next chapter, we will fully develop this idea and defend the claim that understanding has final value *qua* achievement. But before exploring the value of understanding, we need to first take a look at the relationship between achievements and luck. We touched on some of these points briefly in the previous chapter, and the literature on this topic is already quite well-developed, so I won't have much to add to the basic discussion. But once we start applying the lessons to understanding, we will have plenty to add to the dialectic. With this goal in mind, let's quickly survey the main points in the field of epistemic luck.

3.4 Achievements and luck

In the previous chapter, we touched on the topic of epistemic luck and how this relates to the twin goals of accuracy and value for knowledge. The discussion there was meant to motivate a virtue-theoretic approach for understanding, but it will be important to have a firm grasp of the relationship between achievements and luck. We discussed two varieties of epistemic luck in that discussion: intervening luck (which is found in standard Gettier-style cases) and environmental luck (where an agent is in an epistemically unfriendly environment, such as the barn façade case). At this point, we need to expand the discussion a bit to include luck in general, rather than just epistemic luck. For the most part, we'll find shifting the discussion from epistemic luck to luck *simpliciter* straightforward, though there may be some translation issues that we will need to work through.

Intervening luck translates quite well, especially with the ability-success relation in place. Recall that cases involving this kind of epistemic luck were ones in which the luck intervenes between the agent and her belief. The example we gave here involved an agent (whom we called Rocky) who saw what he believed to be a sheep on a hill. What he was actually looking at was a sheep-shaped rock but just behind that rock, out of view, was a sheep. So Rocky's belief that there is a sheep on the hill was true, but only luckily so. The standard example of intervening luck when we are talking about achievement in general is an archer who hit the bull's-eye, but only because a gust of wind blew the arrow on target; otherwise, she would have missed the target. Intervening luck, then, breaks down the ability-success relation—the archer's success does not manifest her relevant abilities. Because this success fails to meet the ability-success relation, our archer is not in the market for achievement.

We noted in the previous chapter that the virtue-theoretic approach for knowledge gets these kinds of cases right, though the explanation for why this is the case was a bit quick. Under the view that knowledge is a kind of achievement, we now have a clear explanation for why cases involving intervening epistemic luck fail to count as knowledge on this view: the agent's success—her true belief—does not manifest her relevant cognitive abilities. And the story is the same for achievements in general. Intervening luck seems to translate well from the epistemic realm to luck and achievements that are outside the epistemic realm. There are two important takeaways from this: cases involving intervening luck fail to meet the criterion of the ability-success

relation, and the notion of success that is relevant for knowledge is the agent's true belief.

Environmental luck is perhaps a little trickier to translate than the current literature would suggest. One of the standard examples of (non-epistemic) environmental luck is an archer with a field of ten targets in front of her. She picks one of the targets at random, takes careful aim, and hit the bull's-eye. As it turns out, nine of the ten targets were protected by an invisible force field so that, had she picked any other target, her arrow would have simply bounced off. Our archer is in an environment in which she could have easily failed, but got lucky in picking the right target. Despite the luck involved in this case, our archer's success does manifest her abilities. And we can build the case so that the shot was challenging enough to meet the difficulty requirement. So, while achievements are *incompatible* with intervening luck, they are compatible with environmental luck. Going back to virtue-theoretic approaches to knowledge, this compatibility with environmental luck is what motivated the inclusion of an anti-luck condition. If we want to preserve the intuition that knowledge cannot be gained in epistemically unfriendly environments, then we would need to include this anti-luck condition to avoid ascribing knowledge in these cases.

The above case with our lucky archer who chose the target unprotected by a force field, however, seems strange—especially when we have a wealth of real-life cases of achievements in unfriendly environments. When we think about achievements in sports—and, in particular, team-based sports—we find agents succeeding in environments that aren't just unfriendly—they're outright hostile! Consider what it takes to get a hit in baseball. The pitcher's mound is 60 feet, 6 inches away from home plate, where the batter is standing. If we do the math, a 90-mph fastball takes just under *half a second* to reach home plate. With human reaction time at around 0.25 seconds, a hitter simply doesn't have enough time to properly react to a pitch. Combine this with the fact that a pitcher often utilizes a variety of pitches that can, because of their spin, change direction after being released and it seems like a near miracle that a batter could get a hit at all. And keep in mind that making contact with the ball isn't sufficient for a hit. The ball must be hit into the field of play, and the batter must run the 90 feet to first base before the nine top-level athletes field the ball and throw it to first, which would result in the batter being out.

With all this in mind, it should come as no surprise that most batters are unsuccessful more often than not. In fact, the typical major-league player only gets a hit in around 1-in-4 attempts. In other words, for every successful hit they have, the average major-league batter has three unsuccessful outings. Given the level of difficulty, we

should have no problem making the case that getting a hit in baseball is a finally valuable achievement. But this is where (what we're calling) the translation problem from the realm of achievements to the realm of the epistemic comes in. Turri (2013; 2017) takes these kinds of statistics to suggest that knowledge, as an achievement, need not be reliable. After all, if we are willing to consider getting a hit in baseball when failure is more common than success, then surely knowledge can be gained in environments that are far less unfriendly.

This may seem like a move that the virtue-epistemologist would embrace. If knowledge is compatible with environmental luck, then there would be no need to introduce an anti-luck condition. This would effectively solve the value problem for knowledge, since we won't have cases of true beliefs that are achievements yet fall short of knowledge. But Turri's claim is stronger than this. It's not just that knowledge is compatible with environmental luck; Turri's suggestion is that knowledge need not be produced by cognitive faculties that are reliable. Epistemologists (reliabilist or not) would find this claim unacceptable—it's standardly accepted that knowledge is the sort of thing that is reliably produced. For our purposes, Turri's argument puts pressure on the ability-success relation. Achievements are successes that manifest an agent's abilities; but if those abilities only produce successes 25 percent of the time, there's a worry that the agent in question doesn't actually possess the relevant abilities. After all, if you are successful less than half the time you attempt something, there is a worry that your abilities either aren't entirely relevant to your success or that you simply lack these abilities altogether and are only getting lucky when you do, in fact, succeed.

Ultimately, this argument hinges on an asymmetry between successes in competitive sports and those in the epistemic realm. If we think back to our fine-grained approach to assessing performances, we find achievements throughout a competition. Much of the time, for each unsuccessful plate appearance for a batter, we can find an achievement on the opposing team. Perhaps the pitcher successfully struck out the batter or forced him to swing at a pitch that resulted in a fielder being able to make a play. In some cases, the batter does hit the ball well, but a fielder makes a great play to get the batter out—like our soccer player who makes a terrific shot on goal but is thwarted by an equally amazing save by the keeper. The point here is that sports in general are set up in such a way to make it difficult to succeed. After all, if these activities weren't difficult to begin with, we wouldn't have the need for professional athletes and it certainly wouldn't be nearly as much fun to watch. This asymmetry makes for a bad comparison when trying to make the case for unreliable epistemic abilities—especially when relying

on statistics in activities that are set up in such a way to place participants in environments where they are likely to be unsuccessful. Thus, we can recognize an agent's low success rate while still preserving the idea that they are manifesting relevant abilities and as reliably as can be expected, given that the environment is stacked against them.

3.5 Concluding remarks

Our main focus in this chapter has been to develop the two main criteria for achievements: the ability-success relation and the difficulty requirement. In the process, we have addressed a few worries about each of these criteria. The important lessons here are, first, that we need to properly locate the success in the ability-success relation and to make sure that any analogies we try to make between general achievements and cognitive achievements are properly analogous. Saying that an epistemic state has value as an achievement does not imply that it behaves in the same way as achievements in general. This will be especially important as we defend the claim that understanding has final value *qua* achievement—what we'll call the *value of understanding thesis*.

4

The Value of Understanding Thesis

Our focus in this chapter is to develop and defend the claim that understanding has final value as a cognitive achievement, what we've called the *value of understanding thesis*. Like the virtue-theoretic approach for knowledge, we are going to focus on the value of understanding and develop a full analysis from there. Most of the pieces of the value of understanding thesis are already in place, so we will spend most of our time looking at objections to this claim. Along the way, we will place the lessons learned with each objection into the broader dialectic so that, by the end of this chapter, we will be able to jump straight in to giving a complete account of ampliative understanding.

Before starting, let's take a quick look at how ampliative understanding fulfills the criteria for achievement: the ability-success relation and the difficulty requirement. One key thing to keep in mind is that the success component for ampliative understanding is not a true belief like it is for knowledge. Instead, an agent with understanding will have successfully pieced together a body of information. This is an admittedly vague way of cashing out the success component, but we will get a clearer idea of what this means as we move through the chapter. As for the difficulty requirement, our primary focus will be on the obstacle-overcoming effort involved in understanding. After all, the onus will always be on the agent to piece together the information in the right way and this piecing together will itself be an obstacle to overcome. But let's start by looking at a general challenge to the ability-success relation and whether it holds in all cases of ampliative understanding.

4.1 The Hand-Holding Objection

Recall that the ability-success relation looks at whether the agent's success manifests her abilities. We used this relation to help explain the compatibility between achievements and luck, though there is a worry that some cases of understanding may not meet this criterion for achievements. Stephen Grimm (2012) has us consider a story from the *Meno* in which Socrates holds the intellectual hand of a young slave, guiding him ever so slowly through a geometric proof. He suggests that in cases like this, "it is plausible to think that

the main reason why anything was grasped at all ... had more to do with the abilities of the teacher than the student" (2012: 112). To put this worry in terms of the ability-success relation, the thought is that the student's understanding is an achievement that is creditable to the teacher rather than the student. In other words, the slave boy's success in grasping the geometric proof manifests Socrates's pedagogical skill rather than his own cognitive abilities. Call this the *hand-holding objection*.

Taken at face value, Grimm's worry does not pose a direct threat to the ability-success relation. The story he is offering seems more an explanatory one—that is, what explains why the student has understanding is the teacher's skill and patience in explaining the material. This explanation seems right. Had Socrates not been so patient or so gifted as a teacher, the slave boy likely would not have understood the geometrical proof. What's more, we could also credit Socrates with an achievement in helping his student understand the geometrical proof. But this sort of explanation—why an agent came to understand such-and-such—is always going to be in the background of any case. Perhaps the agent read a book on the subject, did their own research, or acquired their information via testimony.

What matters is whether the agent has an appreciation of how the information hangs together. And, as we noted at the beginning of this chapter, the onus is always going to be on the agent to do this piecing together. A teacher can do a lot to facilitate this process, even to such a degree that the teacher can also be credited with an achievement in presenting the material. Teaching comes with its own set of obstacles and abilities, and presenting material in a clear way can sometimes be difficult. But a teacher cannot intercede in the active cognitive processes that are required for understanding on behalf of the student. In other words, there's nothing a teacher can do to form the connections necessary for understanding for the student. That process must take place internally such that the agent will meet the criterion of the ability-success relation—even in cases involving intellectual hand-holding.

However, this objection does help emphasize the notion of success that is relevant to understanding: the successful piecing together of information in the right way. For ease of discussion, we will call this success an *appreciation*. Most of the literature on understanding recognizes an appreciation as at least a necessary condition for understanding. Kvanvig uses this idea of piecing together information to help draw a distinction between understanding and knowledge:

Understanding requires the grasping of explanatory and other coherence-making

relationships in a large and comprehensive body of information. One can know many unrelated pieces of information, but understanding is achieved only when informational items are pieced together by the subject in question. (2003: 192)

Pritchard (2008: 334) echoes this sentiment:

Understanding requires ... a genuine conception of how the cause and effect are related ...

As does Linda Zagzebski (2001: 242):

Understanding ... involves grasping relations of parts to other parts and perhaps the relations of part to wholes.

The theme of the above quotes is that, when thinking about knowledge, we tend to focus on the agent's relationship to a true belief. But with understanding, we are looking at whether an agent has what we've called an appreciation.

In the above quote, Kvanvig explicitly recognizes acquiring an appreciation as a necessary condition for understanding. Other authors recognize the importance of an appreciation in slightly different ways but all of them, I take it, would accept acquiring an appreciation as the relevant notion of success for understanding. Well, maybe. There is a particular debate regarding the relationship between knowledge and understanding in which the success component for understanding is identified as a true belief. At the center of this debate is the claim that understanding is a species of knowledge. Call this the *USK thesis*. While I don't want to come down on one side of this debate, we do need to look at the arguments put forward by the opponents of the USK thesis. This is because the next objection I want to respond to relies heavily on some of the moves within the debate. Without some background information on what's going on here, my response to this objection won't make much sense.

4.2 Interlude: The USK Thesis

As I mentioned, our goal in this section is not to come down on the USK thesis, but simply to rehearse the main argument given by opponents of this claim. In the process, we will add further support to the claim that the success component of understanding is an appreciation rather than a true belief. The main argument against the USK thesis is that understanding is compatible with environmental epistemic luck, whereas knowledge is not. So, a case in which an agent has understanding but lacks the relevant knowledge

would be enough to show that the former is not a species of the latter.

Pritchard, an opponent of the USK thesis, develops just such a case involving environmental epistemic luck.

(FIRE OFFICER): Inquirer arrives home to find his house has burnt down. On the scene are several people who appear to be fire officers. Inquirer asks one of them what happened and is told that his house burnt down because of faulty wiring. Unbeknownst to Inquirer, the person with whom he is speaking is the only genuine fire officer on the scene. The rest are merely dressed as fire officers on their way to a costume party. Had Inquirer asked any of these party-goers, he would have been told a false but completely believable story about the cause of the fire.⁶

The argument here is that in cases like FIRE OFFICER, the environmental luck defeats the agent's knowledge but not their understanding. What explains this difference is (what we're calling) the value of understanding thesis. Since achievements in general are compatible with environmental luck, understanding, as a cognitive achievement, is compatible as well. The important feature about this case for our purposes is the notion of success in play: Inquirer's true explanatory belief about the cause of the fire.

It's unclear why Pritchard identifies a true explanatory belief as the success component for Inquirer's understanding in FIRE OFFICER. Pritchard may simply want to hold fixed the relevant notion of success for Inquirer's knowledge and his understanding to make the case clearer. But whatever the reason, this case has muddied the waters and has opened up room for objections to the value of understanding thesis that otherwise might not have gained any traction. What's more, using a true belief as the success component for understanding risks undermining the value of understanding thesis.

To see this, we just need to look at whether acquiring a true belief can meet the criteria for achievement; the answer here is a pretty straightforward 'no'. First, we can put pressure on whether the ability-success relation holds in cases of true beliefs gained via testimony. It's not clear that acquiring a true belief manifests the recipient's cognitive abilities, rather than the testifier's. To be fair, there may be some relevant abilities on offer. Inquirer did go ask a fire officer rather than, say, a child or a dog. So he has some degree of ability in identifying typically reliable informants. But when it comes to the difficulty requirement, there isn't any significant level of ability or obstacle-overcoming effort involved.

⁶Adapted from Pritchard (2010) and Morris (2012).

In Chapter 2, we discussed the problem of testimonial knowledge for the virtue-theoretic approach. We noted that Pritchard (2010) argues that cases of testimonial knowledge cannot meet the demands of the strong achievement thesis (the claim that we modified to our twin criteria for achievement), so knowledge does not have final value *qua* achievement. This ended up motivating his shift in focus away from knowledge and towards understanding. By holding fixed the success component for Inquirer's knowledge and his understanding, Pritchard risks undermining his motivation for this shift in focus. In Chapter 3, we saw that the difficulty requirement is actually a constraint on the success component of the ability-success relation. For a given successful performance, the success in question must require a significant level of ability or an obstacle-overcoming effort. Inquirer's true belief, which he acquired via testimony, cannot meet the difficulty requirement. In short, identifying the success component for understanding risks undermining the value of understanding thesis, which is used to get cases like FIRE OFFICER off the ground in the first place. But we also would lose focus on the notion of success that is at the heart of understanding—what we've identified as acquiring an appreciation. This is not to suggest that there may be cases of understanding in which the success component is, in fact, a true belief. But for any kind of understanding that would have final value *qua* achievement—including ampliative understanding—our focus should be on whether or not the agent has an appreciation.

Once we put the focus back on acquiring an appreciation, we end up with a much more natural reading of FIRE OFFICER. The way the case is presented, Inquirer has a quick conversation with the fire officer and forms the true belief that his house burnt down because of faulty wiring. It would seem, then, that Inquirer already had an understanding of how faulty wiring can cause house fires and applied his broader understanding to this particular situation. In other words, Inquirer's appreciation—the success that is relevant to his understanding—took place well before Inquirer found himself in this situation. After all, there is no mention of the fire officer explaining to Inquirer how faulty wiring causes houses to burn down.

So, the correct way to view FIRE OFFICER is not that Inquirer's understanding is *compatible* with the environmental luck in play; instead, we should say that the environmental luck is *irrelevant* to his understanding. It's like making the claim that penguins can survive if Goldbach's conjecture is true. There isn't some feature of penguins or the conjecture that explains this compatibility—it's just that the two things have nothing to do with one another! In order to have a case of environmental luck that is relevant to an agent's understanding, we would need for the agent to be in an

environment such that they could easily fail to acquire an appreciation. In this respect, it's not clear that FIRE OFFICER fits the bill.

This result ends up opening moves for both sides of the debate on the USK thesis. Proponents of the thesis could argue that the value of understanding thesis isn't doing any work in explaining why Inquirer's understanding is preserved in FIRE OFFICER. And without any explanation for the compatibility between understanding and luck in this case, there's not much blocking a move to insist on an anti-luck condition for understanding similar to what we find for knowledge. On the other hand, opponents of the USK thesis could argue that this result only strengthens their position. After all, if the environmental luck in play is irrelevant to whether Inquirer has acquired an appreciation, then we have no reason to think that Inquirer would lack understanding.

What's important for our purposes is that we can use the above reasoning to respond to an objection which is based in part on FIRE OFFICER. At the center of this objection is the claim that understanding can be gained in cases of intervening epistemic luck, which looks to threaten the ability-success relation in a different way than the hand-holding objection.

4.3 The Lucky Understanding Objection

The first objection we looked at—the hand-holding objection—suggested that an agent could come to understand something because of the abilities of their teacher. In these cases, the ability-success relation is still in place, it's just that the success purportedly manifested the teacher's abilities rather than the student's. This is similar to a professional archer standing right behind an amateur, helping them hold the bow and aim the shot. In *this* case, there does seem to be a sense in which a successful performance (i.e. hitting the target) manifests the teacher's abilities and not the student's. But once we identified acquiring an appreciation as the notion of success in play for understanding, we could see that the internal connections needed for understanding must be made by the agent trying to understand the material.

The debate surrounding the USK thesis, however, has opened up room for an objection that does directly threaten the ability-success relation. Kevin Morris (2012) makes a compelling case that understanding can be gained in cases involving intervening luck. In the previous chapter, we saw that achievements are incompatible with this kind of luck. So if Morris's argument goes through, we will have a case of understanding

without an achievement. In other words, the agent's success won't manifest her cognitive abilities; instead, she will succeed because of luck. Call this the *lucky understanding objection*.

Morris focuses on a case, developed by Pritchard, that is very similar to FIRE OFFICER. The key difference here is that intervening, rather than environmental, luck is in play.

(FIRE INVENTOR): Inquirer arrives home to find his house has burnt down. On the scene is a man dressed as a fire officer who tells Inquirer that his house burnt down because of faulty wiring. Unbeknownst to Inquirer, the man he is talking to is not a real fire officer, but is merely dressed as one on his way to a costume party. This party-goer has completely fabricated a story about the cause of the fire. As it turns out, everything this fire inventor has said is true – Inquirer's house did burn down because of faulty wiring.⁷

The standard approach to FIRE INVENTOR is to say that Inquirer lacks both knowledge and understanding because of the presence of intervening luck. Pritchard makes it quite clear that Inquirer should not have understanding in the above case. As he puts it, "One cannot gain an understanding of why one's house burnt down by consulting someone who, unbeknownst to you, is not a real fire officer but instead merely someone in fancy dress" (2010: 78).

Morris puts pressure on this idea, suggesting that Inquirer does, in fact, understand why his house burnt down. He claims that "to be able to explain something, and to have an explanation, one does not need to worry much about how one came to believe the explanatory proposition" (2012: 360). In other words, if an agent can explain why such-and-such is the case, then the etiology of this explanation isn't important. I am sympathetic to Morris's point here and share the intuition that Inquirer does understand why his house burnt down. But we need to keep distinct the question of whether Inquirer understands why his house burnt down from the question of whether this is a challenge to the value of understanding thesis. It turns out we can ascribe understanding to Inquirer without any risk to the value of understanding thesis.

Just like in FIRE OFFICER, the notion of success in play is Inquirer's true belief about the cause of the fire. With this in mind, we can run the same argument we used in the previous section against the lucky understanding objection; the result is that the intervening luck in play is irrelevant to the success component for Inquirer's

⁷ Adapted from Pritchard (2010) and Morris (2012).

understanding. We still have a situation in which Inquirer has a brief conversation with a person—this time a fake fire officer—and forms a true belief about why his house burnt down. But just like before, it seems that Inquirer already had an appreciation of how faulty wiring can cause house fires and is simply applying his understanding to this particular case.

We could alter the case slightly so that Inquirer doesn't already have a general understanding of how faulty wiring can cause a fire. In this altered case, the fake fire officer would have to invent some explanation of how faulty wiring causes house fires—an explanation that would turn out to be correct. We can again ask whether Inquirer understands why his house burnt down seems relevant, but our focus is on the value of understanding thesis and whether this altered case presents a challenge to the claim. I argue that it doesn't. Even though we have a case in which Inquirer is working to acquire an appreciation, the luck in play has nothing to do with this notion of success. So long as Inquirer can correctly piece together the information with which he is presented, he will have demonstrated a finally valuable cognitive achievement.

The salient feature of FIRE INVENTOR is that Inquirer is lucky to have true information with which to engage; the fake fire officer could have provided a false, but believable, explanation about how faulty wiring can cause a fire. But once we shift our focus away from Inquirer's true explanatory belief to Inquirer's appreciation of the information with which he is presented, its etiology is no longer relevant to Inquirer's success on this front. Pritchard may be right that one cannot gain an understanding from a fake fire officer; but what explains the failure to understand has nothing to do with the value of understanding thesis. If Pritchard wants to get this result, we would need an anti-luck condition for understanding that would exclude cases like FIRE OFFICER from being classified as cases of understanding. A move like this, however, introduces a new set of challenges for the value of understanding thesis.

To address these worries, let's take a quick inventory of our results so far. We've seen that, in order to stay consistent with the value of understanding thesis, the relevant notion of success for understanding must be the agent's appreciation of how a body of information hangs together. If we focus on the agent's true explanatory belief (and its etiology), we cannot account for the final value of understanding *qua* achievement. After all, this true belief can be reached via testimony with little-to-no cognitive engagement from the agent working to understand why such-and-such is the case. We've already discussed the strong parallels between understanding and knowledge vis-à-vis the difficulty requirement for finally valuable achievements. Acquiring a true belief through

testimony involves neither a significant level of ability nor an obstacle overcoming effort.

All this seemingly leaves us without a principled means of ascribing or withholding understanding for agents in cases involving epistemic luck. Pritchard, in no uncertain terms, says that an agent cannot gain an understanding of why their house burnt down by talking to someone who isn't actually a fire officer. Since we've shown that the intervening luck in FIRE INVENTOR is irrelevant to Inquirer's understanding, we would need some other way to explain why Inquirer lacks understanding on Pritchard's view. One way of capturing Pritchard's intuition is to note that Inquirer is lucky to have true information available to him with which he can engage. This, I take it, is a point about epistemic risk in general. We don't want to ascribe knowledge in such an epistemically risky scenario, and this would presumably hold for ascriptions of understanding as well. The other way we can take Pritchard's point is that Inquirer's understanding of why his house burnt down has seemingly come 'out of thin air'. After all, the fake fire officer doesn't understand why the house burnt down—this is stipulated in the case (remember he is making up the cause of the fire with no idea that the explanation he's offering is the correct one). Thus understanding is not transferred from one agent to another (such as from a teacher to a student); Inquirer just ends up with this understanding of why his house burnt down. This seems odd, to say the least.

We are left with a dilemma, of sorts. Everything we've said so far points strongly to an appreciation (rather than a true belief) as the relevant notion of success for understanding. But as we've seen, this move makes environmental and intervening luck irrelevant to the ability-success relation for understanding. This, in turn, would seem to motivate the introduction of an anti-luck condition for understanding that is similar to what we might find in virtue-theoretic accounts of knowledge. But introducing an anti-luck condition could offer a direct challenge to the value of understanding thesis. Recall that this move, when applied to knowledge, blocked the available avenue for capturing the value of knowledge. An anti-luck condition confers no additional value to an agent's true belief that manifests her cognitive abilities. As a result, one would need to construct an entirely different story that explains the value of knowledge over these environmentally lucky true beliefs. We have no reason to think that similar problems wouldn't arise when introducing an anti-luck condition for understanding. Introducing an anti-luck condition for understanding would leave us with a class of cases in which the relevant success for understanding is in place (i.e. the agent has acquired an appreciation) yet fail to qualify as cases of understanding. Since the anti-luck condition confers no additional value on these lucky cases, we face the same value problem as we

did with knowledge.

Clearly, the first horn of this dilemma is not an attractive line to take if we want to preserve the value of understanding thesis. This leaves us facing the second horn: that understanding can be created ‘out of thin air’. Recall that in FIRE INVENTOR, the fake fire officer did not understand why Inquirer’s house burnt down. But after talking with the fire inventor, Inquirer came away understanding why this happened. Our challenge in taking this line will be to explain this unintuitive result and to diagnose what’s going on in cases like this. On the plus side, we have successfully defended the value of understanding thesis against two prominent objections in the literature. We’ve been able to identify the success component of understanding that makes it immune to standard cases of epistemic luck that would defeat an agent’s knowledge. Rather than a true belief, which is vulnerable to both the hand-holding objection (a belief gained via testimony) and the lucky understanding objection (when intervening epistemic luck is in play), the ability-success relation for an agent’s appreciation remains unaffected. This move also matches up with the value of understanding thesis much better than focusing on an agent’s true belief.

As it stands, we are looking at a disagreement on the intuitive response to cases involving epistemic risk. Common sense seems to be on Pritchard’s side when he suggests that one cannot gain an understanding from someone who is making up information (what we’ve been referring to as an understanding appearing out of thin air). While we won’t be able to settle this disagreement here—after all, this project is meant to set up a *framework* for a theory of understanding—we can look at other cases in which an agent we would intuitively withhold ascribing understanding to an agent. Perhaps one of the clearest examples of this would be a case in which an agent is presented with false information, rather than information that is invented but is luckily true. As we will see, agents in these cases can nonetheless demonstrate a finally valuable cognitive achievement—even when engaging with a body of information that is mostly false.

4.4 Factivity and Appreciation

The general worry I want to address here is that attributing a cognitive achievement to agents who acquire an appreciation of false information may seem incredibly unintuitive. But making this move comes with several advantages and ties back into the point we developed in the previous chapter about having a more fine-grained view of successful

performances. There, we considered a soccer player who breaks away from the defense and shoots the ball into the upper corner, which is saved by the keeper. We were able to identify a number of successes within that performance that could be considered strong achievements, despite the fact that the player was ultimately unsuccessful.

Viewing a given performance in a more fine-grained way allowed us to better account for the overall value of a performance or state of affairs. We also identified a narrow and a broad notion of success based on the intentions or aims of the agent that will play an important part here. In the soccer player scenario, we noted that, in a narrow sense, his aim was to strike the ball in the upper corner; but in a broader sense, his aim was to score a goal for his team. Although he was successful in placing the ball where he wanted, his broader aim was unsuccessful since the keeper blocked the shot. We can use the same strategy here to help identify what's going on in cases involving false information.

Let's start with a case in which an agent is working to acquire an appreciation of a body of information that is, for the most part, false.

(SEMINAR): Sally wants to understand why some stars become black holes while others don't. To that end, she attends a seminar on that very topic. The lecturer, angry about not getting tenured, decides to present a completely coherent and believable – but entirely false – explanation of why some stars become black holes. Sally diligently takes notes, asks questions, and studies the material. After the seminar she has pieced together the information she has been given and takes herself to understand why some stars become black holes.

Using the distinction of narrow versus broader aims, we can diagnose what's going on in SEMINAR. Sally has a broader aim to understand why some stars become black holes. On that front, she is clearly unsuccessful. But she also has a more narrow aim to understand the material that the lecturer is presenting. And it's here that we can locate a strong cognitive achievement.

To be clear, Sally does not understand why some stars become black holes. For that, the information with which she is engaging would need to be true. Yet it still seems that Sally has demonstrated a finally valuable cognitive achievement. After all, she has acquired an appreciation and has done so in a way that manifests her cognitive abilities. The key here is that, while Sally doesn't understand why some stars become black holes, she does understand *something* – namely the explanation offered by the lecturer. And

even though her understanding is not factive, it behaves in the same way a factive understanding would. We can compare Sally's understanding to another seminar attendee, Sam, who had trouble following the lecture and failed to acquire an appreciation. In other words, he failed to understand the lecture. We could even assess Sally's appreciation to determine to what degree she understands the material.

One might worry here that I've stopped focusing on cases of understanding that are of epistemological import. Put another way, if we attribute understanding to agents in cases like SEMINAR, we aren't getting the cases right. But we can avoid this worry by distinguishing between an agent who has *an* understanding versus one who has a *factive* or *correct* understanding. This is an important distinction for several reasons. First, it allows us to say what's going on in cases where an agent is working to understand something like astrology. An expert astrologer would, for example, have a higher degree of understanding with respect to astrology than would someone who just reads their horoscope every day. Even though much of the information contained within the field of astrology is false, an agent can still piece together that information in the right way. So, for example, there may be significance to the location of Jupiter with respect to some constellation. An agent with a strong understanding of astrology would be able to recognize this significance and tell us what sorts of things this means, according to astrology. By putting factivity to one side, we can explain what's going on when an agent comes to understand a false theory.

This move also allows us to accommodate Catherine Elgin (2009)'s concerns about the factivity of understanding. Here is Elgin's worry about imposing a strict factivity requirement:

An 8-year-old's understanding of human evolution might include as a central strand the proposition that human beings descended from apes. A more sophisticated understanding has it that human beings and the other great apes descended from a common hominid ancestor who was not strictly an ape. The child's opinion displays some understanding of evolution. It is clearly cognitively better than the belief that humans did not evolve. But it is not strictly true (2009: 325).

To be clear, Elgin is objecting to the factivity requirement suggested by Kvanvig (2003) that an understanding cannot contain any falsehoods that are central to the understanding. In other words, any falsehoods must be, as he puts it, peripheral to the understanding. But we can broaden Elgin's point here to help make another distinction.

Using the SEMINAR example, we can assess Sally's understanding of E – the explanation offered by the lecturer as to why some stars become black holes. But we can also assess E *as an explanation*. Put more broadly, we can assess S's understanding of P or we can assess P as an understanding of the world around us. This seems to capture the spirit of what Elgin is after. The 8-year-old may have been told that humans evolved from apes from her teacher, and we can ask to what degree she understands this information. But, as Elgin points out, an understanding that has it that humans and apes evolved from a common ancestor is a better understanding of the world around us.

By using a fine-grained approach to cognitive achievements, we can accommodate the intuition that there is *something* an agent understands, even in cases of a false theory or information. We can also distinguish between our epistemic aims of understanding some theory or subject matter and understanding the world around us. As for a theory of understanding, we can see that we will need some way to account for the different degrees of factivity two bodies of information may contain. Since an agent can acquire an appreciation of information that is false, we need to be able to pick out some feature of a factive understanding that makes it better *qua* understanding.

But as far as the value of understanding thesis goes, we are still in good shape. Remember, we can have cognitive achievements that fall short of understanding; what we can't have is understanding without a cognitive achievement. With this in mind, there are two more objections with which we need to engage. So far, we have only been looking at objections that challenge the ability-success relation. But there are some compelling arguments that challenge the difficulty requirement instead. We will also look at an objection which suggests that locating the final value of understanding as a cognitive achievement is a misguided approach.

4.5 The Easy Understanding Objection

In the literature on the value of understanding are cases in which the understanding in play is so easy to acquire that we would balk at the idea of classifying it as a strong achievement. These are cases in which there seems to be no obstacle-overcoming effort or significant ability required. If this is right, then we have cases of understanding that fail to meet the difficulty requirement and, in turn, can't be classified as strong achievement. Call this the *easy understanding objection*.

Whiting (2012) offers several cases in which a so-called easy understanding is in

play. Here is one such case and Whiting's analysis of it.

I see my son knocking a glass of milk and spilling its contents on to the floor. I thereby understand why there is milk on the floor – because my son knocked the glass and spilled the milk on to the floor. In this case, I do not even have to put two and two together; that is, I do not have to find the link between my son's actions and the spilled milk; I simply see the spilling of the milk. The degree of cognitive ability exercised here is negligible (2012: 220).

Whiting's goal here is to force a dilemma for the difficulty requirement and, in particular, for what we mean by a 'significant' level of ability. Pritchard is understandably vague on what constitutes a significant obstacle or level of ability. We noted, while talking about achievements, that a principled approach to these sorts of issues (such as whether an agent is manifesting her abilities in a given performance) isn't very plausible. But in the case above, Whiting wants the proponent of the value of understanding thesis to come down on whether or not there is a significant level of ability in play. If we say that there is, then we run the risk of an unintuitive result that is at odds with the strong achievement thesis. On the other hand, if we deny that there is a significant level of ability in play, then we may end up undermining the value of understanding thesis entirely.

At least, this is how Whiting views the dilemma. We do have some other options available, however. We could make the same sort of move we did in FIRE OFFICER and FIRE INVENTOR and point out that there is already an understanding in place before the son even spills the milk. In other words, the dad already understands why fluids, when spilled, can end up on the floor. But this isn't as compelling of a move as it was in previous cases; understanding why fluids can end up on the floor still doesn't seem to meet the difficulty requirement. So pushing the case back still leaves us with the same kind of dilemma.

Recall that we began this project with a guiding intuition that understanding has a special, or distinctive, value. One of the consequences of this view of understanding is that it is something over and above knowledge. What this means is that we can have cases in which an agent *knows* (in this case) why such-and-such is the case but can fall short of *understanding* why. Here's another way of putting this point: to ascribe understanding to an agent is to say something about that agent's epistemic state that is distinct from knowledge. Whether or not understanding is a species of knowledge or even if understanding turns out to lack final value *qua* achievement, we need the

epistemological notion of understanding to do some sort of work.

I suggest that what is going on in the above case is an infelicitous replacement of the word 'knows' with the word 'understands'. This isn't a grammatical concern—we can typically make this sort of replacement and still get a properly formed sentence. Changing 'Sally knows geometry' to 'Sally understands geometry' is grammatically permitted. What we want to avoid is making such a replacement when the change doesn't carry any epistemological weight. Looking back at the spilled milk case, it's clear that the dad *knows* why there is milk on the floor. But to claim that he *understands* why doesn't carry any epistemological distinction. In other words, this isn't a case in which the agent could know why yet fail to understand why.

This approach meshes up well with the current thinking on understanding; in particular, the thought that understanding involves engaging with a network of body of information. We even used this thought to identify the success component for understanding as an appreciation of how the information hangs together. By Whiting's own analysis of the case, however, there is no piecing together required by the dad. As he puts it, "I do not even have to put two and two together" (ibid.). Thus the easy understanding objection is, at best, an objection based on cases that aren't giving a charitable reading of the literature on understanding. At worst, the objection rests on an outright strawman of the value of understanding thesis.

That being said, there may well be certain kinds of understanding that are in play. In the introductory chapter, we identified several kinds of understanding that, while important, would not be the focus of this project. My hunch is that there is some *conceptual* understanding on offer; the dad needs to have a basic understanding of causation and the movement of fluids, for example. But this sort of understanding is a prerequisite for his knowledge of why there is milk on the floor. If we want knowledge to be anything more than the mere 'parroting back' of information, then the dad should have this conceptual understanding in order to know why there is milk on the floor.

Whatever kinds of understanding are on offer – whether it is conceptual or linguistic or something we haven't yet identified – ampliative understanding is not in the picture. In these cases of easy understanding, there is no information with which the agent can engage. Placed in the broader dialectic, then, ampliative understanding will involve a large body or network of information. After all, if there is no appreciation to be had of how the information hangs together, then the agent isn't even in the market for ampliative understanding.

With the easy understanding objection out of the way, we can now turn to the

final objection we will look at. In our discussion on achievements, we identified the status of acquiring a significant level of ability as a finally valuable achievement. We have used this approach to argue that acquiring an understanding is a finally valuable achievement. But we also want to get the result that having an understanding is also finally valuable. As we will see, this objection looks to threaten that move.

4.6 A Finally Valuable State

In defending the value of understanding thesis, we have so far been focused on the value of acquiring an understanding. We have worked to show that this success meets the criteria of the strong achievement thesis. But Whiting (2012) offers another objection that – even if everything we've said here is right – our overall approach is wrong. Here is Whiting's worry:

...[A]ppealing to achievement could at most account for the value of *arriving at* an epistemic state, not of *being in* it, of *acquiring* an epistemic status, not of *having* it. Compare the following. Climbing to the summit of Mount Everest is a great achievement and to that extent valuable. It does not follow that *sitting* on the summit of Mount Everest is an achievement ... even if one could only sit on the summit of Mount Everest having climbed there (2012: 229).

To put the worry another way, winning a race may well be a finally valuable achievement. But the resulting status – that of being a race winner – doesn't seem to have value *qua* achievement.

The point here is not that these statuses lack all-things-considered value. Being a race winner may bring with it endorsement deals or open up qualification paths for other, more prestigious races. But this sort of value is *pragmatic*, rather than ethical, and the value of being a race winner is only instrumental relative to these other goods. Since understanding is a cognitive state, we need to show that being in this state has final value, and that it is valuable *qua* achievement – that is, the value is ethical rather than, say, pragmatic. We touched on this point in the previous chapter when looking at whether Tiger Woods sinking putt after putt was an achievement. There, we identified the state of being an agent with significant ability (in this case, the status of being a proficient golfer) as a strong achievement.

While making this move was an important step in locating strong achievements

in a given state of affairs, we are still missing a piece of the puzzle. In order to fill in this gap, we need to show that being in this state has final value *qua* achievement. We will need to be careful because there is a disanalogy between acquiring an understanding and, say, becoming a proficient golfer. The latter is a state of being an agent with a high level of ability which, in turn, explains why an agent might not need to demonstrate an obstacle-overcoming effort in order to be successful at a given performance. But understanding doesn't confer such a status to the agent.

Where the similarity lies is in the fact that both states put the agent in a position to demonstrate other finally valuable successes within that domain. A proficient golfer, for example, is in a better position to succeed at performances within the domain of golfing (e.g. putting or driving the ball). Similarly, an agent with understanding – or, at least, with ampliative understanding – will be able to use this state to acquire true beliefs within that domain. Someone who understands astrophysics, for example, can use her understanding to acquire true beliefs that are relevant to astrophysics. Put more broadly, an someone with understanding can use that understanding and apply it to new scenarios or problems. This is really the distinctive feature of ampliative understanding – the ability to acquire relevant, true beliefs within a given domain.

This is a feature of understanding that has been recognized throughout the literature. Catherine Elgin makes this point very eloquently:

The understander should be able ... to use that information – for example, to reason with it, to apply it, to perhaps use it as a source of working hypotheses about other related matters... [S]omeone who understands geometry can reason geometrically about new problems, apply geometrical insights in different areas, assess the limits of geometrical reasoning for the task at hand, and so forth (2009: 323).

Henk de Regt (2004) and James Woodward (2003) make very similar points in recognizing that an agent with understanding has an ability to address and answer counterfactual questions (e.g. what if things had been *this way*). Stephen Grimm (2006; 2014) also takes a similar line in his grasping account of understanding. An agent who grasps how certain information hangs together can see how a change in one aspect of a given state of affairs may lead to other changes. At the core of all of these proposals is the idea that an agent with understanding can acquire true beliefs because they are in a particular state – i.e. they have an understanding of the information with which they are engaging. Thus our approach focuses on this particular kind of understanding that puts

an agent in a position to be successful (i.e. acquire true beliefs) in future epistemic performances. Much like a proficient golfer, an agent with understanding can demonstrate finally valuable achievements because they are in a particular state.

This view fills in part of the puzzle – the future goods of having an understanding are valuable *qua* achievement rather than having value of a different kind. But we might think that understanding has only instrumental value relative to these future goods. We can point out here that, while being a race winner and having an understanding are both persistent states, an agent must actively attend to her understanding in order to maintain that state. In other words, being a winner of the Boston Marathon is a persistent state that doesn't require active engagement. The winner of this race could never run again and still have this status. But understanding, much like being a proficient runner, requires something on the part of the agent in order to maintain. In other words, acquiring an understanding and holding on to that understanding are both strong achievements that require having (and maintaining) an appreciation.

We can also use our fine-grained approach to locating success to help fill in the gap here. An agent may have a broad goal of winning the Boston Marathon, but their more narrow goal would be to finish the race under a certain time. In other words, the agent is working to become a proficient runner and will take whatever future successes come with that status. Viewing the state of affairs in this way, we can also see that succeeding in a way that manifests these abilities allows for those future successes to have final value *qua* achievement. The status of being a proficient runner actually confers ethical value to these future success, as opposed to winning a race because of cheating or through luck.

This point is easier to see in the cognitive realm. When we look at the goal of inquiry, there is a broad sense in which we're concerned with acquiring true beliefs. In cases of purely practical inquiry, such as when the next bus is arriving or what time you have to be at work, acquiring a true belief may be all that you're looking for. But when we engage in inquiry for its own sake – that is, when we just want to learn about some event, concept, or subject matter – what we're after is understanding. We don't want to just know (or truly believe) a bunch of related pieces of knowledge – we want to understand the material. There may even be cases of practical inquiry in which understanding is a suitable goal. Understanding the bus schedule, for example, would allow one to acquire true beliefs about when the next bus is arriving. In short, ampliative understanding – as an epistemic aim – puts one in a position to acquire true beliefs that are strong achievements precisely because of the final value of being in that state of understanding. The soccer player who strikes the ball into the upper corner has already demonstrated a

finally valuable achievement and will happily take whatever success results from that achievement. And it is the status of the shot itself as a finally valuable achievement that confers value on a resulting goal.

Notice, though, that we are now talking about acquiring true beliefs as finally valuable achievement. This may seem strange given that so much of this chapter has been focused on arguing against the idea that a true belief can be the success component for understanding. What's important to keep in mind here is that acquiring an appreciation is still the relevant notion of success for coming to understand some state of affairs or subject matter. Once an agent has an understanding, however, we can identify acquiring true beliefs as successes that manifest the agent's cognitive abilities. Recall that the worry before was that acquiring a true belief in the cases we've looked at so far weren't strong achievements. This was either because the ability-success relation failed to hold (we saw this cases of testimony, where the testifier's abilities were more relevant to the success) or because the difficult requirement wasn't met. But when an agent uses her understanding to acquire true beliefs, this is clearly a success that manifests her cognitive abilities. And because one must actively engage with the information in order to acquire these true beliefs, there will be an obstacle-overcoming effort in play, much as there is when an agent acquires an appreciation. Just like someone can see (or fail to see) how a body of information hangs together, one can see (or fail to see) what follows from this information.

4.7 Concluding remarks

In this chapter, we have looked at and responded to a number of challenges to the value of understanding thesis. Along the way, we have placed the lessons learned within a broader dialectic that has helped us shape the account of understanding on offer. We began by identifying the relevant notion of success for ampliative understanding, which is acquiring an appreciation of how the information hangs together. We have also seen that this success can be met even in cases involving environmental or intervening epistemic luck that would defeat the agent's knowledge. An agent can also acquire an appreciation of false information, such as in SEMINAR. These results will help to frame the account of understanding on offer. But the foundation of ampliative understanding rests on the agent's ability to acquire true beliefs. Not only does this move match up well with the literature – it also helps explain the distinctive value of understanding both as a

goal of inquiry and as an epistemic state. With the general structure of ampliative understanding in place, we are now ready to provide a proper analysis of this kind of understanding.

5

Ampliative Understanding

This chapter will begin with an analysis of ampliative understanding. We have already unpacked most of the aspects of this analysis in the previous chapters, so most of our time will be focused on considering and responding to objections.

Ampliative Understanding (AU): S understands why P if and only if S can acquire true beliefs that are relevant to P in a way that manifests S's cognitive abilities.

The first thing to note is that the above analysis is for explanatory understanding. This is simply to keep the account from being too unwieldy or awkward. We can replace the 'why' with a 'that' to talk about straightforward propositional understanding. And we can remove the 'why' to capture objectual understanding (understanding some topic or subject matter such as the French Revolution or astrophysics). It's also important to keep in mind that we are working to build up a framework for analyzing understanding. As has been the case throughout this project, the key feature of ampliative understanding is the agent's ability to acquire true beliefs in such a way that manifests the agent's cognitive abilities. In particular we want these true beliefs to manifest her appreciation of how the information hangs together. As we unpack this analysis and respond to a variety of objections, we will look at some potential revisions. While I think that (AU) is on point, the analysis carries with it some intuitions that I have about the nature and structure of understanding. In order to make this view as attractive as I can to those who may not share these intuitions, we will consider some revisions that are in the same vein. While these revisions will vary in the details, they are all based on this core idea of acquiring true beliefs.

As I mentioned, we have already unpacked this analysis for the most part. In our discussion of achievements, we made the move to cash out achievement as successes that manifest ability. We have also seen why understanding is valuable as an epistemic state: it puts the agent in a position to have more epistemic success. In particular, an agent with understanding will be able to acquire true beliefs. As a quick aside, we don't need to require that these true beliefs that the agent acquires be *new* true beliefs (i.e. beliefs that

are new to the agent). While it's reasonable that in most cases these true beliefs would be new, including that as an explicit requirement would be a misstep. It's unreasonable to think that we have every piece of information that we take ourselves to know readily available. With this in mind, we should allow for an agent to use her understanding in order to navigate and retrieve information that comprises an understanding. As an example, I take myself to know when the American Civil War ended; but suppose that, for whatever reason, I can't immediately recall the date. I can use my historical understanding to navigate my way to this piece of information. I know, for example, that Abraham Lincoln was assassinated near the end of the Civil War and I know that he was shot in 1865. I can put this information together to get the correct year for the end of the Civil War: 1865. While this is a simple example, it's meant to illustrate the general point of navigating the information we have as well as the seamlessness with which this can take place.

The only claim within (AU) that we haven't explicitly discussed is the idea of relevance. In other words, what it means for a true belief to be relevant to an explanation, topic, field of study, etc. The thought here is that an understanding is something that already comes as compartmentalized. Information about pigs isn't going to be relevant to astronomy. But giving a comprehensive story about relevance would be a project in its own right. My hope is that we have enough of an intuitive idea about what it means for a belief to be relevant to a body of information that we don't need to unpack this claim. And it may be the case that we don't actually need this constraint for (AU) to be successful. When thinking about relevance, it's hard to see how we could acquire a belief about pigs (much less a true belief) by engaging with information on astronomy. Since the information that comprises an understanding is already compartmentalized, we take this notion of relevance as primitive and develop the rest of the story from our intuitions.

5.1 The Factivity Requirement

In the previous chapter, we captured the notion of success that is relevant to understanding as acquiring an appreciation of how the information hangs together. While this move fits better with the value of understanding thesis than having a true belief, it misses out on an easy way to satisfy the intuition that understanding, like knowledge, is factive. Call this intuition the *factivity requirement*. If we had held on to an agent's true belief as the relevant notion of success, we may have been able to meet the

factivity requirement for free. But the tradeoff wouldn't have been worth it; as we've seen, we would have been unable to defend the final value of understanding as a cognitive achievement. This is a much more attractive route for the project at hand, and it will give us an opportunity to engage with and make clear this factivity requirement.

We've already seen that an agent is still in the market for acquiring an appreciation—even when she is engaging with information that is mostly false. The example we looked at was a student, Sally, who wanted to learn about black holes. She attends a class a seminar on the topic, takes careful notes, follows the lectures, and takes herself to understand why some stars become black holes. But the lecturer presents his students with believable, but false, information about black holes. The conclusion we arrived at was that, while Sally can't be said to understand why some stars become black holes, she was nonetheless able to acquire an appreciation of how the information she was presented with hung together. Since there is a cognitive achievement on offer that is relevant to understanding, the factivity requirement is what we need to explain why we wouldn't ascribe understanding to Sally in this case.

We have two fairly broad options for imposing a factivity requirement, which I call *strong* and *weak*, and each option allows some room for variation. Let's begin by looking at some possibilities for a strong factivity requirement. In general, this option would not allow for any falsehoods in an agent's understanding. At the extreme end of this option would be to deny understanding to an agent who has any falsehoods whatsoever in the body of information that comprises her understanding. This move is far too extreme. All of us have false information sprinkled throughout various domains—even those in which we are experts. The extreme version of the strong requirement returns the implausible result that understanding a topic of any real complexity would be virtually impossible. As imperfect epistemic agents, we are almost certainly going to make mistakes when it comes to gathering and processing information on a difficult or complicated subject. It should not follow that we therefore lack understanding of these topics.

Kvanvig (2003) endorses a variety of the strong factivity requirement that excludes false information from an agent's understanding, rather than having a false belief defeat her understanding. But he also makes an important distinction between false information that is central to an understanding and that which is peripheral:

When the falsehoods are peripheral, we can ascribe understanding based on the rest of the information grasped that is true and contains no falsehoods. In such a case, the false beliefs

are not part of the understanding the person has, even though they concern the very material regarding which the person has understanding (Kvanvig 2003: p. 201-2).

Unlike the extreme version, an agent can have false beliefs on Kvanvig's view and still be in the market for understanding. We simply set those false beliefs to the aside and assess an agent's understanding based on the remaining true information. Kvanvig's more moderate view also gives us room to deny understanding when those falsehoods are central to agent's understanding. We will return to this point, but a quick example may help here. Going back to our hapless student Sally, suppose that one of her central beliefs is that black holes are the result of galactic dust clouds that cover up a star, which makes it black. We don't need to come down on what distinguishes a central falsehood from a peripheral one to see that this particular falsehood is central to Sally's understanding of black holes.

While Kvanvig's strong factivity requirement seems to do well in getting the cases right (i.e. excluding those cases in which the agent is clearly not in the market for understanding without excluding too much), there are a few worries with his view. Catherine Elgin develops one of these worries:

... [S]ome bodies of information, even though they are not true, nonetheless display a measure of understanding. The growth of understanding often involves a trajectory from beliefs that, although strictly false, are in the right general neighbourhood to beliefs that are closer to the truth. The sequence may terminate in true beliefs. But, I contend, the earlier steps in the sequence should fall within the ambit of epistemology (Elgin 2009: p. 325).

Elgin's main point is that Kvanvig is excluding too much. When we are assessing a body of information, it's reasonable to include earlier stages when developing an understanding. Consider the Copernican model of our solar system, which gets a number of things right—especially when compared to earlier, geocentric models. But the model was incomplete (Copernicus was unaware of the planets beyond Saturn) and showed the planets' orbits as perfectly circular, rather than slightly elliptical. Elgin's example is that of a student who incorrectly believes that humans evolved from monkeys (the correct belief is that humans and monkeys evolved separately from a common ancestor). The student's understanding of evolution is nonetheless far better than someone who believed that humans evolved from birds or that we didn't evolve at all.

It may be that the above examples aren't a problem for Kvanvig since it's not obvious (to me, at least) that these beliefs are central to understanding our solar system or evolution, respectively. However, Kvanvig's move to exclude false information from our assessment introduces a more serious worry when assessing the *degree* of an agent's understanding. Recall that part of what it means to get the cases right for a theory of understanding is to be able to account for the degrees of understanding. Consider two different bodies of information which comprise two cases of understanding of the same subject matter. The first, U_1 , has only true information. The second, U_{1+n} , contains all the true information of U_1 but also has n additional pieces of information, all of which are false. On Kvanvig's view, we would exclude those n pieces of information when assessing the understanding associated with U_{1+n} . This would make these two cases of understanding equivalent in degree, which doesn't match up with the intuition that an understanding which contains falsehoods is worse off than an understanding with all the same true information but without the falsehoods. In other words, we would say that an agent with understanding U_{1+n} does not understand the subject as well as an agent with understanding U_1 . It's not enough that falsehoods don't count towards an agent's understanding—we need false information to count *against* an understanding. So we must allow for falsehoods when assessing an agent's understanding.

Once we're in the arena of allowing for falsehoods in an understanding, we are on the side of the weak factivity requirement. Like the strong version, there are several varieties from which we can choose. Some of these versions might look at how many false beliefs we can allow while others would focus on how central these falsehoods are to the agent's understanding. We can imagine these views along a spectrum with an extremely weak version at one end, which would allow a large proportion of falsehoods (including central ones) into an understanding, and a stronger version at the other end that allows for only a few peripheral falsehoods. The worry with this view may already be apparent; by endorsing a weak factivity requirement, we run the risk of introducing a threshold problem. There is no clear factivity threshold that delineates agents who are in the market for understanding from those who aren't. To be fair, not all threshold problems are problematic. After all, some concepts can admit of vagueness. But vagueness on a key epistemological issue like factivity, while not devastating to a theory of understanding, would make the view very unattractive.

We are left with a dilemma. The strong factivity requirement may be too strong in that it excludes cases of understanding which are of epistemological import. We noted an additional worry that Kvanvig's factivity requirement (the only tenable version of the

strong requirement) has trouble account for the degrees of understanding in some cases. The weak factivity requirement, on the other hand, introduces a threshold problem that would make the resulting theory of understanding unacceptably vague. If forced to take a horn, we should definitely take the first. The cases in which Kvanvig's requirement would have trouble assessing the degrees of understanding are actually pretty limited. And Elgin's worry about including earlier versions of an agent's understanding could be accommodated, so long as the falsehoods involved are not central to the agent's understanding. Fortunately, ampliative understanding provides a third option that can accommodate the factivity requirement while avoiding the above dilemma.

The above dilemma only arises when we try to place a factivity constraint on the body of information comprising an understanding. But the key feature of ampliative understanding is the agent's ability to acquire a true belief. The factivity requirement is thus imposed on what beliefs an agent can acquire by engaging with the information, rather than directly on the body of information. Naturally, this move will raise questions of its own, which we will now address.

5.2 Objections to the Factivity Requirement

There are two major objections we need to examine. The first is whether we've just shifted the threshold problem for the weak factivity requirement onto the true beliefs that the agent can acquire. Our worry with the threshold problem was that it created an unacceptable level of vagueness for an account of understanding, so providing a clear factivity requirement will be key to answering this objection. The second objection is how we can be sure that a true belief is the result of engaging with true information. Many different processes can result in true beliefs and not all of these processes are going to be representative of proper epistemic practices. We will take each of these worries in turn.

Our new threshold objection asks how many true beliefs an agent must be able to acquire in order to meet the factivity requirement. The answer, while *prima facie* implausible, is that so long as the agent is able to acquire *any* true beliefs, she has met the factivity requirement. The obvious response is to point out that, on this view, an agent who is only able to acquire a single true belief after engaging with a body of information is in the market for understanding. We have several responses to this worry available, depending on the framework of understanding that we want to develop.

First, we can work to debunk this objection, which is the response I prefer. Notice that this objection makes two assumptions: (1) the body of information that comprises an understanding is numerable, and (2) there is an understanding that would allow for an agent to acquire only one true belief. Another way to view this objection is that it would rely on some contrived body of information that could comprise an agent's understanding such that she could only acquire one true belief from that set. Let's think about the plausibility of developing such a set. We could begin by returning to whatever information Sally, our black hole student, received from her lecturer. In that case, we stipulated that the information she received was entirely false. As a result, when she engages with this information, it seems apparent that she is going to wind up with false beliefs as a result. We could alter some of the information so that it includes a small set of true information. That way, when she engages with this smaller chunk, she will acquire a true belief. But in changing small pieces of information, we've changed the way that the information hangs together as a whole. In other words, we will lose the connections that were in place before we made the change and, as a result, Sally will no longer be able to acquire an appreciation.

At this point, however, the Sally case is becoming less and less useful. I introduced it to motivate the disconnect between understanding and the success component for knowledge of having a true belief. As a genuine case to consider, however, Sally's understanding of black holes is also contrived and assumes that the body of information involved is numerable. When we're talking about numerability, I am not suggesting that the set of information is uncountably infinite; rather, it's not obvious that the information and necessary connections that would comprise an understanding can be listed in any meaningful or thorough way. Textbooks may come close to a list like this, but it seems clear that the information presented in the textbook is going to be structurally quite different from the information that comprises a student's understanding of that textbook.

Perhaps the above discussion seems too handwavy or perhaps the reader doesn't share my intuitions about the information contained in an understanding. Fortunately, we can offer a different response by altering our factivity requirement along reliabilist lines. Instead of requiring that the agent be able to acquire any true belief, we can make the move to say that an understanding is factive when the agent can *reliably* acquire true beliefs. This would be cashed out as an agent who is able to acquire more true beliefs than false ones; typically, the cut-off point is a 50 percent. I am amenable to this move as it ties in well with the virtue-theoretic project as a whole—agents who are in the market for achievement are typically going to be reliably successful. John Turri (2015) puts some

pressure on this idea by noting cases of achievement in which agents are unreliable. A good example of this is an average (or even above average) major league baseball player, where a batting average of .400 (that's an average of 4 hits over 10 at bats) is an excellent success rate. We need not worry about a disconnect between reliability and success for understanding. Unlike in a professional sports arena, we don't usually find ourselves in an unfriendly—or, rather, outright hostile—environment in which other agents do everything they can to keep us from succeeding.

To sum up, ampliative understanding will meet the factivity requirement. But we should not impose that requirement on the body of information with which the agent is engaging. Instead, our requirement is focused on the beliefs that the agent can acquire. And we have two ways of spelling out this requirement. The first is to require that the agent is able to acquire any true beliefs based on the information. Keep in mind that this line carries the implicit claim that an understanding in which only one true belief could be acquired is a contrived and implausible scenario. For those who consider this too weak or that we could develop a counterexample to the above line (i.e., an understanding in which the agent could plausibly be able to acquire only one true belief), we can make a move towards reliabilism. This would require that the agent can acquire more true beliefs than false ones based on the information that comprises her understanding.

The second objection we need to consider is whether we can have a case in which the agent is engaging with information that is false and yet ends up acquiring a true belief. This objection comes as a result of our focus on what beliefs the agent can acquire, rather than the body of information with which she is engaging. My initial response is that it's hard to imagine (for me, at least) how such a scenario might happen. Take, for example, the so-called flat earth society. This is a marginalized (but apparently growing) group of people who reject the globe model of Earth in favor of a flat, dinner-plate-shaped Earth. As a result of engaging with this—and other—false claims, every belief they acquire is patently false. The mental gymnastics these individuals must perform in order to hold on to their ridiculous beliefs results in this group claiming that gravity is not a real force, that the space beyond Earth is an illusion and that the Sun and Moon are located inside a dome that encloses the planet. While this is admittedly an extreme example, it demonstrates the claim on offer: individuals that engage with false information are going to end up acquiring false beliefs. There may be cases in which an agent acquires a true belief from false information because of luck, but this result would no longer manifest the agent's cognitive abilities or her appreciation of the information.

Another way in which an agent might come to acquire a true belief based off of

false information is if that information contained a contradiction. Since a contradiction logically implies (at least, through material implication) any statement—whether true or false—this is a case worth considering. We have several available here, one of which we’ve just examined. The reliabilist approach would handle a case like this very well, since one is more likely to acquire a false belief than a true one. In other words, this is not a reliable method for acquiring true beliefs. This response also ties in with our previous point about acquiring a true belief because of luck, since the acquired belief would not manifest the agent’s cognitive abilities or appreciation. Finally, it’s worth noting that contradictory claims are not amenable to acquiring an appreciation in the first place. One cannot appreciate how contradictory claims hang together since they are, by definition, at odds with one another—i.e., they can’t hang together.

5.3 Grasping and Ampliative Understanding

For Grimm, grasping is a cognitive ability that allows us to do two things.⁸ First, grasping some concept allows us to manipulate the variables within that concept; in other words, the agent would be able to see how things could have been different. An agent who grasped Newton’s second law of motion ($F = ma$) could tell us how the acceleration of an object would change if the force applied to that object or its mass was different. Second, an agent who grasps a concept could apply that concept in other situations. So understanding, insofar as it involves grasping, requires that the agent be able to reason counterfactually and apply their understanding to different scenarios or situations.⁹

In addition, Grimm requires that what is grasped is the correct explanation; understanding why a plane can fly is to have a grasp of Bernoulli’s Principle. This in turn means that one could figure out how the plane might fly if the wings had a different shape or if the plane were heavier. In requiring that an agent grasp the correct explanation, we can account for (at least the spirit of) the factivity required for understanding without running into the problems we saw in the previous section. Grasping can also account for the degrees of understanding and to illustrate this, Grimm has us consider a scenario.

[W]hen I start chopping onions and my eyes begin to water, I think I understand why my eyes are beginning to water, namely, because I am chopping the onions. I don’t think it is

⁸ See Grimm (*forthcoming*)

⁹ There is a genuine worry (cf. Pritchard, *forthcoming*, §4) about the intelligibility of grasping. I will set these worries aside, however, because it is the spirit of the account that I find most plausible.

because of the time of day, or the color of the shirt I am wearing, or anything like that; it's because of the onions. But obviously someone with a greater understanding of onion (and eyeball) chemistry would be able not just to identify the onions as the cause but would be able to say what it was about the onions that was bringing this about – in this case, the particular sulfur compounds that were being broken down and released into the air when I did the chopping (Grimm *forthcoming*: p. 11).

For Grimm, the key difference between understanding-why and knowing-why his eyes are watering comes down to the ability to provide correct answers to questions that ask how things might have been different. In the onion example, one with a greater understanding would be able to answer more of these kinds of questions. So, for example, one could answer what might happen if Stephen were to chop up a different vegetable that contained this sulfur compound; or what would happen if the sulfur compound were removed from the onion. In short, a higher degree of understanding implies an ability to correctly answer a greater variety (and number) of these types of questions.

Grimm's notion of grasping can also make sense of when an agent comes to understand a false theory, such as alchemy or phlogiston theory. Grimm actually identifies two kinds of grasping: *simple* grasping, where an agent believes the explanation she grasps is the correct one; and *conditional* grasping, where the agent does not believe the explanation grasped is correct. In principle the idea is sound – in trying to understand a false theory, we are nonetheless subject to epistemic scrutiny. The problem comes in when we think that a false theory is true or that a true theory is false. Suppose an agent has successfully grasped some particular explanation X; we can generate the following possible outcomes based on their doxastic attitude toward X and the actual truth of X.

	X is correct	X is not correct
S believes X is correct	Simple grasping	Type I error
S does not believe X is correct	Type II error	Conditional grasping

When S is right about whether or not a particular explanation is correct, then we have either simple or conditional grasping. But grasping has nothing to say about when an agent believes an explanation is correct when it is not (which I will call a Type I error) or

when an agent believes that an explanation is incorrect when it is, in fact, correct (a Type II error). In both of these cases, however, it seems unreasonable to say that the agent in question fails to understand anything at all.

Thinking back to the onion case, we can imagine two agents who have an explanation of why onions make their eyes tear. Nancy, who has a very good understanding of phlogiston theory, believes that onions contain trace amounts of phlogiston which, when released, causes her eyes to burn and tear up. Sally, on the other hand, understands the theory of why the sulfur chemicals in the onion causes her eyes to tear but believes that some other theory (perhaps Nancy's) is the correct explanation. In this case, Nancy has committed a Type I error by believing that an incorrect theory is correct while Sally has committed a Type II error. Although neither of these agents is in a particularly admirable epistemic situation, the fact that grasping has nothing to say about either of them is telling. While Nancy does not have a correct explanation of what is going on, it seems unreasonable to not count her as understanding phlogiston theory. Sally is in a very interesting epistemic situation as far as grasping is concerned. She could answer questions correctly about cutting other vegetables or about onions that lacked the sulfur compounds – she just would not believe that any of her answers were correct. But again it seems that, according to Grimm, Sally would lack understanding.

There are two important lessons we can take from the above example. First, there is a separation between grasping and being able to reason counterfactually evinced by Type I errors. Sally clearly has this ability to the same extent that someone would who did believe their explanation was correct. Yet Sally fails to grasp anything because she does not believe the answers she is giving or the explanation behind them. Nancy is also able to answer questions correctly about what a phlogiston theorist would say about onions. The only difference between Nancy and an agent with a conditional grasp of phlogiston theory is that Nancy believes the theory is correct. This separation between grasping and what it enables an agent to do jeopardizes the value of understanding on the grasping account. Nancy and Sally are both agents who can answer questions correctly with respect to a particular domain (Nancy's responses are limited to explanations on behalf of phlogiston theory) yet they both lack grasping, and thus understanding. This is especially problematic in Sally's case, since she can correctly answer questions about the actual world. Instead of the traditional problem of explaining the value of knowledge over mere true belief, we have run into a difficulty of explaining the value of understanding over merely true statements.

The main takeaway from this account (for our purposes, at least) is that we can capture Grimm's notion of grasping with an agent's ability to acquire true beliefs. The two main features of grasping—reasoning counterfactually and applying one's understanding to new situations—are both cases of acquiring true beliefs. And while grasping offers more precision than the broader claims within (AU), it's not clear that grasping is necessary for understanding. When discussing factivity, we touched upon the possibility of someone who is just beginning to understand a concept or explanation. In these early stages of development, one may not be able to meet the conditions required for grasping. Of course, this would be a very basic level of understanding; but so long as the agent has an appreciation of how the information hangs together and is able to acquire true beliefs based on that appreciation, she is in the market for understanding. Put another way, grasping is clearly sufficient for understanding. But it may be too strong a requirement to be included as a necessary condition. It seems clear that grasping is an important goal for an agent who is aiming to understand something, but there are early stages (as Elgin points out) in which grasping may not be present and yet the understanding is nonetheless of epistemological import.

Another nice feature of Grimm's grasping account is that it helps us start to construct a story about how we might assess the degree of an agent's understanding. On the grasping view, once one is able to start reasoning counterfactually and apply their grasp to new situations, they have understanding. Someone who is able to apply their grasp to more situations and counterfactuals would have a higher degree of understanding. For (AU), the story will be very similar, except that we are looking at the amount and variety of true beliefs that an agent can acquire. Counterfactual reasoning gives us a somewhat limited criterion along which to assess the degree of an understanding. By looking at the variety of true beliefs an agent can acquire, we are capturing this counterfactual reasoning along with other kinds of reasoning that could help indicate a stronger understanding. A broader criterion in this case allows us to give a more nuanced assessment of the degrees of understanding.

5.4 Concluding remarks

We started this chapter with an account of ampliative understanding, but it's important to view this as a concrete instance of a more general framework. We may not all share the same intuitions about the nature and structure of understanding. But the key

feature of ampliative understanding—an agent’s ability to acquire true beliefs—is built up from a virtue-theoretic approach to understanding that preserves the value of understanding thesis. Understanding, as an epistemic state, puts one in a position to enjoy more successes, much like being a skilled athlete. This is a state that is valuable to achieve and, because we can acquire (often new) true beliefs, it is a valuable state in which to be.

We’ve looked at two prominent discussions of understanding in the literature: Kvanvig’s factivity requirement and Grimm’s grasping account. On both fronts, ampliative understanding is able to accommodate the spirit of these discussions. By requiring that an agent acquire *true* beliefs, we are able to introduce a factivity requirement for understanding while avoiding some potential worries for Kvanvig’s view. The acquisition of true beliefs is also able to capture Grimm’s conditions for grasping: counterfactual reasoning and applying one’s understanding to different situations. We also saw how (AU) can account for the degrees of understanding, which is an important feature in getting the cases right for understanding.

Along the way we also considered some possible revisions to (AU), depending on the intuitions one has about the nature and structure of understanding. These revisions are still motivated by the virtue-theoretic project and are able to preserve the value of understanding thesis. The end result is a framework built around an agent’s cognitive abilities and that is rigorous enough to facilitate further discussion on understanding. In the final chapter, we will look at what responses are available to some of the main questions in the philosophical landscape on understanding.

6

Further Research

We started this project with some important questions about understanding that seemed intractable without a clear enough understanding of understanding we could bring to bear on the discussion. With a framework for developing an account of understanding—and an analysis based on framework—in place, we can get a cursory look at what a principled answer to these questions might look like. This discussion is not intended to be exhaustive; the goal here is to simply help inform and guide future research on understanding.

One of the central questions in the literature is whether understanding is a species of knowledge. The most common arguments we see in the literature focus on cases that involve epistemic luck. In particular, the debate is built around cases involving environmental epistemic luck. Since achievements in general are compatible with environmental luck, understanding, as a cognitive achievement, is compatible with this variety of luck as well.

In Chapter 4, we put pressure on this argument. We came away with the conclusion that the debate is fundamentally misguided. These cases hinge on identifying the success component of understanding as a true explanatory belief, which is the same success component for knowledge. But we saw that holding fixed the relevant notion of success between understanding and knowledge risks undermining the value of understanding thesis. By focusing on whether the agent has a true explanatory belief, the opponent of the USK (Understanding is a Species of Knowledge) thesis opens themselves to the same objections that put pressure on the claim that knowledge has final value as an achievement. When this true belief is gained via testimonial knowledge (which is precisely what happens in the cases used to argue against the USK thesis), an agent's true belief manifests the informant's cognitive abilities. Another way of expressing this worry is that this success cannot meet the difficulty requirement for achievements. This discussion motivated the move away from having a true explanatory as the success component for understanding and towards the agent acquiring an appreciation of how the information hangs together. But this move made the environmental luck irrelevant to this success. In other words, an agent in these cases is not in an epistemically unfriendly

environment with respect to acquiring an appreciation.

More generally, we seen that the etiology of the information comprising an understanding does not have any principled impact on whether or not we ascribe understanding to an agent. Even in cases involving Gettier-style intervening luck (a case in which the information received is only luckily true), the luck involved does not affect the ability-success relation that is relevant to understanding. It's at this point where intuitions diverge with some authors taking it as obvious that one cannot gain an understanding in these sorts of cases while others suggesting that the etiology of the information genuinely does not matter. The framework we have develop for ampliative understanding does not explicitly come down on this issue. According to (AU), however, the presence of intervening luck (in the cases we've looked at) would have no impact on an agent's understanding. So long as one is able to acquire true beliefs in a way that manifests cognitive ability, ampliative understanding will be preserved.

It's understandable why there might be resistance to this result. After all, cases involving intervening luck are standard examples of what we can call *epistemic risk*. This risk alone could be enough for some authors to insist on including an anti-luck condition on the framework we have established here. There is nothing in the framework of ampliative understanding that would preclude this move, but those who insist on introducing an anti-luck condition risk undermining the value of understanding thesis. We see this worry with virtue-theoretic accounts of knowledge: by including an anti-luck condition, these theories of knowledge allow room for cases of cognitive achievements that fall short of knowledge. As a result, such a theory can no longer accommodate the intuition that knowledge is more valuable than other epistemic states that fall short of knowledge. The value in play rests on the value of the achievement; no additional value is conferred by the anti-luck condition. We can run the same kind of argument against a theory of understanding that includes an anti-luck condition.

The motivating factor for including an anti-luck condition is the presence of epistemic risk. In other words, the epistemic luck in play creates this sense of epistemic risk. But this sense may just be an artifact of long-held knowledge-centric thinking in epistemology. We've already seen that the epistemic luck is not relevant to an agent's understanding, so there is no reason to suspect that an agent's understanding is vulnerable to the same epistemic risk as knowledge. In short, the introduction of an anti-luck condition could be a knee-jerk reaction to cases involving epistemic luck that is based off of our intuitions about knowledge. This reaction risks undermining the value of understanding thesis without a clear, articulable, or principled reason to do so. With

the framework of ampliative understanding in place, we can start to make decisions about how to deal with understanding and epistemic luck that aren't informed by artifacts from a knowledge-centric bias. A shift in epistemological focus away from knowledge and towards understanding should ideally carry with it as few of these biases as possible. With a framework for understanding in place, we can meet this goal and begin to think about understanding on its own terms, so to speak.

Shifting focus to understanding brings with it additional benefits for handling problem areas of knowledge. A prime example of such a benefit is the problem of trivial knowledge.¹⁰ The worry here is that some truths seem to have no value, epistemic or otherwise. A standard example of this sort of trivial knowledge is someone who spends their time meticulously counting the number of dust motes on a surface. Such a belief intuitively carries no epistemic value and, because of the time commitments involved, might even have *disvalue* since we could spend our time doing other, more productive things.

An epistemology built around understanding, which is a project I am working to build from the research here, has a principled means of explaining why some true things aren't worth knowing. Counting dust motes or grains of sand, or memorizing phone numbers doesn't contribute to one's understanding in any significant way. In other words, incorporating these beliefs into a relevant body of information (assuming there is, in fact, such a body of information) would do nothing to aid the process of acquiring true beliefs. On this view, what explains the value of knowledge is its role in our understanding of the world around us.

Clearly, there is a lot of work left to do; the framework we have developed is just a starting point for this endeavor. We have said nothing about other kinds of understanding, and the questions that we started with on understanding remain open. With the framework of ampliative understanding in place, we can start to make principled responses to these questions that can accommodate the variety of intuitions and other comprehensive theories of understanding that are sure to arise in the literature.

¹⁰ Alston (2005) and Goldman (1999) both work to deal with the problem of trivial knowledge. Boylu (2010) considers some responses that are similar to the line we take here.

References

- Achinstein, P. (1983). *The Nature of Explanation*, Oxford: Oxford University Press.
- Alston, W. (2005). *Beyond "Justification"*, Ithaca: Cornell University Press.
- Baehr, J. (2009). 'Is There a Value Problem?', *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, p. 42-59, Oxford: Oxford University Press.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard University Press.
- Boylu, A. (2010) 'How Understanding Makes Knowledge Valuable' *Canadian Journal of Philosophy*, vol. 40, no. 4: p. 591-610.
- Brogaard, B. (2007). 'I Know. Therefore, I Understand', *typescript*. Available on the World Wide Web: <http://www.philosophy-dev.stir.ac.uk/staff/duncan-pritchard/documents/BrogaardOnKvanvig2.pdf>
- DePaul, M. & S. Grimm (2007). 'Review essay on Jonathan Kvanvig's 'The Value of Knowledge and the Pursuit of Understanding'', *Philosophy and Phenomenological Research*, vol. 74, no. 2: p. 498-514.
- Elgin, C. (2007). 'Understanding and the Facts' *Philosophical Studies* 132 (1): 33 - 42.
- Elgin, C. (2009). 'Is Understanding Factive?' *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, p. 322-330, Oxford: Oxford University Press.
- Fogelin, R. (1994). *Pyrrhonian Reflections on Knowledge and Justifications*, Oxford: Oxford University Press.
- Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis*, 23: p. 121-3.
- Goldman, A. (1999). *Knowledge in a Social World*, Oxford: Oxford University Press.
- Goldman, A. & Olsson, E. (2009). 'Reliabilism and the Value of Knowledge', *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, p. 19 - 41, Oxford: Oxford University Press.
- Greco, J. (1999). 'Agent Reliabilism', *Noûs*, vol. 33, no. s13: p. 273-96.
- Greco, J. (2003). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul and L. Zagzebski, p. 111-34, Oxford: Oxford University Press.
- Greco, J. (2007). 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17, p. 57-69.
- Greco, J. (2009a). 'Knowledge and Success from Ability', *Philosophical Studies* 142: p. 17-26.
- Greco, J. (2009b). 'The Value Problem', *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, p. 313-21, Oxford: Oxford University Press.
- Grimm, S. (2005). '*Understanding as an Epistemological Goal*', Ph. D. Dissertation, University of Notre Dame: U.S.
- Grimm, S. (2006). 'Is Understanding a Species of Knowledge?', *British Journal for the Philosophy of Science* 57, p. 515-35.
- Grimm, S. (2008). 'Explanatory Inquiry and the Need for Explanation', *British Journal for the Philosophy of Science* 58, p. 482-97.
- Grimm, S. (2010). 'Understanding', *The Routledge Companion to Epistemology*, (eds.) S. Bernecker & D. H. Pritchard, p. 84-94, London: Routledge.
- Grimm, S. (*forthcoming*). 'Understanding as Knowledge of Causes', *Virtue Scientia*, (ed.) A. Fairweather, Dordrecht, Holland: Springer.
- Hempel, C. (1965). 'Aspects of Scientific Explanation', (ed.) C. Hempel, *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.
- Kim, J. (1994). 'Explanatory Knowledge and Metaphysical Dependence', *Philosophical Issues* 5: p. 51-69.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Kvanvig, J. (2009a). 'Responses to Critics', *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, p. 339-352, Oxford: Oxford University Press.
- Kvanvig, J. (2009b). 'The Value of Understanding', *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, p. 95-111, Oxford: Oxford University Press.

- Kvanvig, J. (2011). 'Coherentist Theories of Epistemic Justification', *The Stanford Encyclopedia of Philosophy (Summer 2011 Edition)*, (ed.) E. N. Zalta, Available on the World Wide Web: <http://plato.stanford.edu/archives/sum2011/entries/justep-coherence/>.
- Lackey, J. (2007). 'Why We Don't Deserve Credit for Everything We Know', *Synthese* vol. 158 no. 3: p. 345-61.
- Lycan, W. (1996). 'Plantinga and Coherentisms', *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*, (ed.) J. L. Kvanvig, p. 3-24, Totowa, NJ: Rowman & Littlefield.
- Morris, K. (2012). 'A Defense of Lucky Understanding', *British Journal for the Philosophy of Science*, vol. 63, no. 2: p. 357-71.
- Pritchard, D. (2005). *Epistemic Luck*, Oxford: Oxford University Press.
- Pritchard, D. (2007), 'Recent Work on Epistemic Value', *American Philosophical Quarterly* 44: p. 85-110.
- Pritchard, D. H. (2009). 'Knowledge, Understanding and Epistemic Value', *Royal Institute of Philosophy Supplement* 64: p. 19-43.
- Pritchard, D. H., Millar, A. & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Riggs, W. (2009). 'Understanding, Knowledge, and the *Meno* Requirement', *Epistemic Value*, (eds.) A. Haddock, A. Millar and D.H. Pritchard, p. 331-8, Oxford: Oxford University Press.
- Riggs, W. (forthcoming). 'The Value Turn in Epistemology', *New Waves in Epistemology*, (eds.) V. Hendricks & D. H. Pritchard, Ashgate.
- Scott, D. (ed.). (2006). *Plato's Meno (Cambridge Studies in the Dialogues of Plato)*, Cambridge: Cambridge University Press.
- Sosa, e. (1993). 'Proper Functionalism and Virtue Epistemology', *Noûs*, vol. 27, no. 1: p. 51-65.
- Sosa, E. (2003). 'The Place of Truth in Epistemology', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul and L. Zagzebski, p. 155-79, Oxford: Oxford University Press.
- Swinburne, R. (2009). *Providence and the Problem of Evil*, Oxford: Oxford University Press.
- Turri, J. (2013). "Knowledge as achievement, more or less". In *The Present and Future of Virtue Epistemology*. Ed. M. F. Vargas. (forthcoming).
- Turri, J. (2015). "Unreliable Knowledge". *Philosophy and Phenomenological Research* vol. 90, no. 3: p. 529-45.
- Whiting, D. (2012). "Epistemic Value and Achievement". *Ratio* 25(2): 216-30.
- Williamson, T. (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.
- Zagzebski, L. (1996). *Virtues of the Mind*. New York: Cambridge University Press.
- Zagzebski, L. (1999). 'What is Knowledge?', *The Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, p. 92-116, Blackwell: Oxford University Press.
- Zagzebski, L. (2001). 'Recovering Understanding', *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility and Virtue*, (ed.) M. Steup, Oxford: Oxford University Press.
- Zagzebski, L. (2004). 'The Search for the Source of Epistemic Good', *Metaphilosophy*, vol. 34, no. 1-2: p. 12-28.