Knowledge and Understanding

By

Blake Myers

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Philosophy)

at the

UNIVERSITY OF WISCONSIN-MADISON

The dissertation is approved by the following members of the Final Oral Committee:
Peter Vranas, Professer, Philosophy
John Bengson, Associate Professor, Philosophy
Michael Titelbaum, Associate Professor, Philosophy
John Mackay, Assistant Professor, Philosophy
Mike Koenigs, Associate Professor, Psychiatry

Introduction

We spend significant amounts of time, energy, and resources in pursuit of understanding why certain facts obtain. While a number of these facts are empirical, others are not. Regarding the former, we seek to understand why life evolved on Earth, why caffeine increases alertness, why worker bees engage in self-sacrifice, and why humans wage wars, to give just a few examples. Regarding the latter, by contrast, we seek to understand why certain theorems in logic and mathematics hold, why a given theory of planetary motion is incompatible with another theory, and why two hypotheses that have been treated as rivals are in fact equivalent.

Much of the literature on understanding-why focuses on one of two relationships: (a) the relationship between understanding-why and knowing-why, and (b) the relationship between understanding-why and explanation. Keeping in line with this approach, my dissertation consists of three parts. Chapter 1 examines the relationship between understanding-why and knowing-why. Chapter 2 looks at the relationship between understanding why some event proposition (i.e., a proposition representing an event) is true and having a *causal* explanation of why the proposition is true. Chapter 3 then looks at the relationship between understanding why some proposition is true and having an explanation of why the proposition is true in terms of logical grounding. Given this, consider a more detailed overview of these three parts, along with how each part is related to the others.

In Chapter 1, I begin by considering the relationship between understanding-why and knowing-why, and I argue that understanding-why is what I call a "quantitative species" of knowing-why. By way of comparison, other relations in which one state (or activity, entity, or property) is a quantitative species of another state (or activity, entity, or property) include: (i) the relation between exhaustion and tiredness, (ii) the relation between proficiency and competency, and (iii) the relation between being an elder and being a person. In all such cases, any essential difference between the two relata can be wholly accounted for in terms of some single magnitude *m* (where the value of *m* required by the first relatum is greater than the value of *m* required by the second relatum, and yet both relata require that *m* have at least some positive value). For instance, in the case of (i), the single magnitude that can account for the difference between exhaustion and tiredness is the need for rest, while in the case of understanding-why and knowing-why, the single magnitude is how well (or fully) some proposition *P* is explained by the set of explanatory propositions (pertaining to *P*) that the person in question knows to be true.

Such cases stand in contrast to the many types of species relations that are not instances of quantitative species; e.g., the relation between being a violinist and being a musician, the relation between being a biologist and being a scientist, and the relation between surfing and moving. While these are all species relations (since, in each case, every instance of the former relatum is an instance of the latter relatum), they are not quantitative species relations (since, in each case, there are essential differences between the two relata that cannot be wholly accounted for in terms of some single magnitude *m*, as *m* is defined above).

In Chapter 2, I move to develop a more robust account of understanding-why by considering the relation between understanding why some event proposition *P* is true and having a causal explanation of why *P* is true. As many authors have suggested, there seems to be a close relationship between understanding why some fact is so and having an explanation of why it is so. I develop a novel account of understanding-why that is a refinement of this basic idea. To do this, I look largely to the literature on dispositions and higher-level properties, since problems discussed in that literature can be used to motivate such an account. I begin by providing a broad overview of several closely related notions, including dispositions, multiple realizability, causation, and explanation. I then introduce a notion that I call "*causal guiding explanation*," and I contrast it with the notion of *objective causal explanation*.

Suppose, for example, that Jamie purchases a small statue while on vacation, and when she arrives home, she opens her suitcase and finds that the statue is broken. Jamie has no idea why the statue broke, so she asks her roommate (who happens to be an avid collector of statues) why the statue broke. After examining the statue, Jamie's roommate replies that it broke because it was made of a very fragile material. Given this, it seems plausible to say that while Jamie initially had no understanding of why the statute broke (prior to talking with her roommate), Jamie acquired at least some understanding of why the statue broke by acquiring an understanding (via her roommate) that the statue broke because it was fragile.

Some argue, however, that once all of the causal roles of the relevant structural and microphysical properties of the statue are provided, any appeal to the statue's fragility is causally superfluous, and thus can be left out when describing the full causal nexus that is responsible for the fact that the statue broke. But even granting this, I argue that there remains a sense in which the statue's fragility might still be relevant to *guiding* a person to an objective causal explanation of why the statue broke. In particular, by acquiring an understanding that the statue broke because it was fragile, a person might acquire the ability to identify some set of propositions Φ, such that some member or other of Φ is at least partially causally responsible for the fact that the statue broke. To illustrate, Φ might contain every proposition of the form, <The statue has property *x*>, where *x* is a potential realizer of fragility. For, by understanding that the statue broke because it was fragile, a person might acquire the ability to understand that *the statue's possession of some property (or other) that realizes fragility* figures into the full causal nexus that was responsible for the fact that the statue broke. Similarly, Φ could also (or instead) contain every proposition of the form, <*x* triggered the fragility of the statue>, where *x* is a potential trigger of the statue's fragility (e.g., the striking, twisting, or dropping of the statue). For, by understanding that the statue broke because it was fragile, a person might acquire the ability to understand that *some fact (or other) that triggered the statue's fragility* figures into the full causal nexus that was responsible for the fact that the statue broke.

Given this, I argue (roughly) that a person has at least some understanding of why *P* (where *P* is a true event proposition) if and only if the person has a causal guiding explanation of why *P*. And a person has a causal guiding explanation of why *P* if and only if there is some proposition *C*, such that: (i) the person understands that *C*; and (ii) the truth of *C* provides information, which is not provided by *P*, about the causal nexus that was responsible for the fact that *P*. Finally, I provide a broad outline of how one might extend the proposed account of what it is to have at least some understanding-why (given in terms of causal guiding explanations) to further account for the *gradability* of understanding-why.

Lastly, in Chapter 3, I turn to examine the relationship between understanding why some proposition *P* is true and having an explanation of why *P* is true in terms of logical grounding. It is widely held that the notion of *grounding* is primitive; i.e., unanalyzable in terms of other independently intelligible notions. I argue, however, that at least one prominent type of grounding (i.e., logical grounding) can indeed be analyzed in terms of other independently intelligible notions. I begin by considering various types of grounding in order to situate the relation of logical grounding among these various types. Then, I examine the relationship between *logical grounding* and *minimal entailment*, and I introduce a notion that I call "*sentential representation*", which serves to connect *sets of facts* to *sets of sentences representing those facts*. Next, I consider how certain sets of sentences are *more basic* than others, and I proceed to explore how this idea might be employed in constructing an account of logical grounding with respect to logically contingent facts. Then, I extend the account so that it also applies to logically necessary facts, and I examine a few final instances of logical grounding and show how the final account applies to such cases. Finally, I considered how the account pertains to the notion of understanding-why.

CHAPTER 1

Understanding-Why and Knowledge-Why

Overview

Over the past decade, there has been much debate over the significance of various features in differentiating understanding from knowledge. Despite what has often been suggested, I argue that there is a close relationship between understanding and knowledge. I examine three prominent proposals that have been made in attempt to differentiate understanding from knowledge, and I argue that once relevant distinctions are made, the purported differences disappear. I then introduce a notion that I call "*quantitative species*", and I argue that while understanding-why cannot be reduced to knowledge-why, understanding-why is a quantitative species of knowledge-why. Finally, I consider how this proposal differs from the mere claim that understanding-why is a species of knowledge-why, and I examine various ways in which the proposal accounts for phenomena regarding the relationship between understanding-why and knowledge-why.

Introduction

There has recently been much debate over the nature of understanding and its relationship to knowledge.[1] In spite of several features that have been proposed to differentiate understanding from knowledge, I argue that there remains a close relationship between them. When considering the relationship between understanding and knowledge, it is important to keep in mind that there are at least three types of understanding featured in the literature (each of which corresponds to a distinct type of knowledge): understanding-why, objectual understanding, and propositional understanding. It is widely held that for any person *S*, and any proposition *P*, *S* understands *why P* only if *S* understands that *Q*, where *Q* is an explanatory proposition

---

[1] Zagzebski 1996, 2008; Kvanvig 2003; Brogaard 2005; Elgin 2007; Pritchard 2009; Grimm 2010; Hills 2010; Baumberger 2011; Gordon 2012.

with respect to *P* (e.g., a proposition of the form, <*P* because *R*>, or <*P* in virtue of the fact that *R*>, or <*R* explains *P*>, or <*R* grounds *P*>).[2] For instance, Mick might have at least some understanding of why his car will not start by understanding that it will not start because the distributor rotor is broken. Objectual understanding, by contrast, is the type of understanding that one has of a particular thing, such as an object, relation, process, event, subject matter, or body of information: e.g., in the sense that Jamie might understand contemporary evolutionary theory, and Jill might understand the behavior of gases. Finally, propositional understanding is the type of understanding that one has with respect to an individual proposition: e.g., in the sense that Caitlyn might understand that mass is a measure of energy.

These three types of understanding correspond to three distinct types of knowledge: knowledge-why, objectual knowledge, and propositional knowledge. One path towards determining the relationship between understanding and knowledge is to determine, for each type of understanding, how that type relates to the corresponding type of knowledge. In this paper, I focus primarily on the relationship between understanding-why and knowledge-why, though I occasionally consider aspects of the relationship between objectual understanding and objectual knowledge, and the relationship between propositional understanding and propositional knowledge, since such considerations are at times relevant to my proposal on the relationship between understanding-why and knowledge-why.

---

[2] Riggs 2003; Brogaard 2005; Grimm 2006; Pritchard 2009; Baumberger 2011; Carter & Gordon forthcoming.

I begin, in Sections 1 - 3, by examining three prominent proposals that have been made in attempt to differentiate understanding from knowledge, and I argue that once relevant distinctions are made, the purported differences disappear. Then, in Section 4, I consider a phenomenon that *does* seem to suggest a difference between understanding and knowledge, or, more precisely, a difference between understanding-why and knowledge-why; and I use this phenomenon to motivate a new account of the relationship between understanding-why and knowledge-why. In brief, I propose two ways in which some state *x* might be a species (or type) of another state *y*: i.e., *x* might be a *quantitative species* of *y*, as exhibited by the relationship between *being a long novel* (as *x*) and *being a novel* (as *y*); or, *x* might be a *non-quantitative species* of *y*, as exhibited by the relationship between *being a Russian author* (as *x*) and *being an author* (as *y*). In Section 5, I employ this distinction between quantitative and non-quantitative species, and I argue that understanding-why is a *quantitative species* of knowledge-why, in that any essential difference between *S understanding why P* and *S knowing why P* can be accounted for in terms of a single magnitude: i.e., the magnitude of how fully *P* is explained by the set of explanatory propositions (pertaining to *P*) that *S* knows to be true.

To begin, consider three views that have been proposed to differentiate understanding from knowledge, namely: (a1) the view that understanding but not knowledge is *non-factive*, (a2) the view that understanding but not knowledge is immune to *epistemic luck*, and (a3) the view that understanding but not knowledge is *transparent*. While some of these views are offered in connection with the relationship between objectual understanding and propositional knowledge (i.e., knowledge-that), it is not hard to imagine similar arguments that might be given with respect to the relationship between understanding-why and knowledge-why. Thus, I examine each of these three views, in turn.

1. Factivity

First, consider factivity. Linda Zagzebski (2001) and Catherine Elgin (2007, 2009) argue that understanding, unlike knowledge, is non-factive. Elgin (2007, 2009) motivates this position by arguing that understanding, unlike knowledge, can involve falsehoods. Her point is perhaps best seen by considering the use of idealizations in science. According to Elgin, certain scientific claims, such as the ideal gas law, are strictly speaking false. However, they are "felicitous falsehoods", in that they play a key role in understanding certain phenomena, such as the behavior of *actual* gases. Elgin (2009) writes:

> We understand the behavior of actual gases by reference to the alleged behavior of a so-called ideal gas. There is no such gas. So how can it figure in our understanding of the world? I suggest that effective idealizations are felicitous falsehoods. That they are false is evident. They are felicitous in that they afford epistemic access to matters of fact that are otherwise difficult or impossible to discern (Elgin 2009).

Since scientific understanding often involves such felicitous falsehoods, Elgin concludes that it cannot accurately be described as factive. What matters most for understanding, according to Zagzebski and Elgin, is that the coherence-making relations among the relevant pieces of information (i.e., pieces of information comprising the subject matter or body of information) are sufficiently grasped.

While this view regarding the non-factivity of understanding seems to hold for certain types of objectual understanding (e.g., understanding a body of information), it does not seem to hold for understanding-why (e.g., understanding why a particular gas has a pressure of 11.2 atmospheres). Consider first the understanding that one might have of a body of information. Here, it seems relatively clear that a person can understand a body of information even if the body of information is not a correct representation of reality. For instance, Joshua might understand Ptolemaic astronomy even though the Ptolemaic system fails to provide a correct model of the cosmos. However, such examples do little to distinguish understanding from knowledge. For, if we compare this type of understanding with the corresponding type of knowledge (i.e., knowledge of a body of information), then it seems that both understanding and knowledge have this feature of non-factivity. To illustrate, just as Joshua might understand Ptolemaic astronomy, he might know Ptolemaic astronomy or know the Ptolemaic system. In order to have knowledge of some body of information, one need not have knowledge of the propositions that constitute the body of information. There are many propositions that figure into Ptolemaic astronomy that one cannot know, since the propositions are false. Nevertheless, one can still know Ptolemaic astronomy by knowing that *such-and-such propositions are part of the Ptolemaic system and that they relate to one another in such-and-such ways* (regardless of whether the propositions themselves are true).

Next, consider the understanding that one might have of *why* a given proposition holds, such as a proposition of the form <real gas *G* has a pressure of *x* atmospheres>. While I agree with Elgin (2009) that one might understand why such a proposition holds by way of understanding the ideal gas law, I disagree that such understanding should be deemed as non-factive. To illustrate, consider the ideal gas law:

$$PV = nRT$$

(where P is the pressure of an ideal gas, V is the volume of the ideal gas, n is the number of moles of the ideal gas, R is the ideal gas constant, and T is the temperature of the ideal gas).

Now, suppose that for some real gas *G* and some value *x*, *G* has a pressure of *x* atmospheres. How might the ideal gas law figure into whether a person understands why some proposition pertaining to this fact holds? The answer appears to be twofold, depending on the type of proposition in question. First, suppose that the proposition is <the pressure of *G* is *approximately* *x* atmospheres>. Call this proposition, "*APROX*". Then, a person might understand why *APROX* by understanding that the pressure of *G* is approximately *x* atmospheres because:

(b1) *G* is a gas that is composed of unbound atoms, has a relatively high temperature and volume, and a relatively low pressure;

(b2) for any gas *G* that is composed of unbound atoms, has a relatively high temperature and volume, and a relatively low pressure, the pressure of *G* is *approximately* equal to nRT/V; and

(b3) when the correct values (with respect to *G*) are plugged in for n, T, R, and V, the value of nTR/V is *x*.

Thus, given that a person might understand why *APROX* by way of understanding (b1) – (b3), a person might understand why *APPROX* by way of understanding some scientific idealization (i.e., by way of understanding that PV = nRT). But since (b1) – (b3) are all true, this fails to give reason to think that understanding-why is non-factive.

Similarly, suppose that the proposition in question is <the pressure of *G* is *exactly* *z* atmospheres>. Call this proposition, "*EXACT*". Then, a person might understand why *EXACT* by understanding that the pressure of *G* is exactly *z* atmospheres because:

(b1) *G* is a gas that is composed of unbound atoms, has a relatively high temperature and volume, and a relatively low pressure;

(b2) for any gas *G* that is composed of unbound atoms, has a relatively high temperature and volume, and a relatively low pressure, the pressure of *G* is approximately equal to nRT/V;

(b3) when the correct values (with respect to *G*) are plugged in for n, T, R, and V, the value of nTR/V is *x*; and

(b4) the difference between *x* and *z* can be accounted for by incorporating more terms into the ideal gas law to correct for features that the ideal gas law excludes, such as spatial extension and non-elastic interactions of real (as opposed to ideal) gas particles.[3]

[3] To roughly illustrate what such an incorporation of additional terms into the ideal gas law might look like, consider the van der Waals equation. This equation can be seen as providing a closer approximation (at least in certain cases) of the behavior of a real gas by modifying the ideal gas law, i.e., PV = nRT, to correct for the ideal gas law's exclusion of non-elastic interactions among real gas particles (i.e., by changing the pressure term "P" to "$P + n^2a/V^2$") and the ideal gas law's exclusion of spatial extension of real gas particles (i.e., by changing the volume term "V" to "V – nb"), where "a" denotes the force of attraction between the particles of gas under consideration, and "b" denotes the volume of one mole of the gas under consideration, resulting in the revised equation, $(P + n^2a/V^2)(V - nb) = nRT$.

So, again, although one might understand why *EXACT* by way of understanding (b1) – (b4), and thus by way of understanding that PV = nRT (a scientific idealization), none of (b1) – (b4) is false. Hence, as with *APROX*, the case of *EXACT* provides no reason to think that the use of idealizations in science suggests that understanding-why is non-factive.

In sum, if we are consistent in our comparisons, so that objectual understanding is compared with objectual knowledge, and understanding-why is compared with knowledge-why, there is no longer any path to distinguish understanding from knowledge by appealing to the claim that understanding is non-factive.

2. The anti-luck condition

A second feature that is proposed to differentiate understanding from knowledge is an anti-luck condition. Pritchard (2009) provides a scenario that is meant to show that understanding, unlike knowledge, is immune to certain types of luck. Suppose that Samuel's house burns down and that he is trying to figure out why it burned down. Samuel therefore asks one of the fire officers why his house burned down, and the officer provides a fairly detailed explanation of why it burned down, which Samuel accepts. Unbeknownst to Samuel, however, there happens to be several other people standing around who are dressed as fire officers for a costume party but are not actual fire officers; and by mere luck, Samuel happens to receive the detailed explanation of why his house burned down from an actual fire officer. Pritchard argues that in such a case, Samuel understands why his house burned down but does not know why it burned down.

I do not find Pritchard's conclusion about this scenario to be intuitively obvious. In particular, I find it plausible to say that Samuel both understands why his house burned down and knows why his house burned down. But in order to avoid an intuition stalemate, I explore another path by which to evaluate Pritchard's conclusion. Namely, I propose that if we grant Pritchard's claim that Samuel does not know why his house burned down, then given three widely held claims, it follows that Samuel does not understand why his house burned down (and thus Pritchard's conclusion is false). To illustrate, consider the following argument:

(1) Samuel does not know why his house burned down. [Pritchard's claim]

(2) For any subject *S*, and any proposition *P*, *S* knows why *P*, if for some *R*, *S* knows that *P* because *R*.[4]

(3) There is no *R*, such that Samuel knows that his house burned down because *R*. [From (1) and (2)]

(4) For any subject *S*, and any proposition *Q*, *S* understands that *Q* only if *S* knows that *Q*.[5]

(5) There is no *R*, such that Samuel understands that his house burned down because *R*. [From (3) and (4)]

(6) For any subject *S*, and any proposition *P*, *S* understands why *P* only if, for some *R*, *S* understands that *P* because *R*.[6]

---

[4] Achinstein 1983; Kitcher 2002; Woodward 2003; Lipton 2004; Brogaard 2005.

[5] Brogaard 2005; Elgin 2007; Pritchard 2009; Grimm 2010; and Gordon 2012.

[6] Kvanvig 2003; Brogaard 2005; Hills 2010; Baumberger 2011.

(7) Therefore, Samuel does not understand why his house burned down. [From (5) and (6)]

Thus, in sum, if a person does not know why *P*, then she does not understand why *P*.

To see why this argument is sound, consider the support for (2), (4), and (6). Regarding (2), suppose that Samuel's neighbor, Susan, knows that Samuel's house burned down because the house had faulty wiring. Then, if Susan were asked whether she knows why the house burned down, she could truthfully respond, "Yes, I know why Samuel's house burned down. It burned down because it had faulty wiring." Similarly, suppose that Omar knows that his class is cancelled because the professor is ill. Then, if someone were to ask Omar whether he knows why his class is canceled, Omar could truthfully respond, "Yes, I know why my class is canceled. It is canceled because the professor is ill." These examples suggest that knowing a proposition of the form <*P* because *R*> is sufficient for knowing why *P*.

Next, consider claim (4). One reason to think that (4) is true comes from arguments used in support of the widely held view that propositional understanding is just propositional knowledge; a view that Pritchard (2009) himself endorses. Consider, for instance, the following passages:[7]

> If the proposition 'I understand that Athens defeated Persia in the battle of Marathon' is supposed to be a stand-alone proposition…it is hard to see how it differs from 'I know that Athens defeated Persia in the battle of Marathon' (Elgin 2007).

[7] Similarly, see Brogaard 2005; Gordon 2012.

> [W]here understanding is of a proposition, it does seem to be pretty much synonymous with knows. On discovering that my train has been cancelled, I may well say to the person at the ticket office that I understand that the train is cancelled in such a way that I could just as well have used 'know' without any loss (Pritchard 2009).

> As several authors have noted…ascriptions of understanding along the lines of (1) [i.e., "Mary understands that her class starts in an hour"] seem to be more or less synonymous with corresponding ascriptions of knowledge. Thus on most occasions it seems that we can substitute "S knows that p" for "S understands that p" with little loss of meaning (Grimm 2010).

While there are cases in which "understands that" is not synonymous with "knows that", such cases seem to be best described by saying that "understands that" is merely being used idiomatically as a hedge to indicate the possibility of error, as in the claim, "I understand that the bus will arrive at 3:15PM, but it might arrive at 3:30PM".[8] Thus, granting that this view is correct, claim (4) is true, since (4) follows from the claim that propositional understanding is just propositional knowledge.

Lastly, with respect to (6) (i.e., the claim that *S* understands why *P* only if, for some *R*, *S* understands that *P* because *R*), suppose that Mick understands why his car will not start. Then, at the very least, Mick must understand that his car will not start because

[8] See Elgin 2007; Pritchard 2009; Grimm 2010.

R (where *R* is a correct answer to the question, "Why will Mick's car not start?"). Accordingly, if Mick were to claim that he understands why his car will not start, but he could not provide any *R* such that he understands that his car will not start because *R*, then it would be hard to see what Mick means by claiming that he understands why his car will not start.

3. Transparency

A third feature proposed to differentiate understanding from knowledge is transparency. Zagzebski (2001, p. 246) argues that understanding, unlike knowledge, is transparent. She writes:

> Understanding … is a state that is constituted by a state of conscious transparency. It may be possible to know without knowing one knows but it is impossible to understand without understanding one understands.

The claim that it is "impossible to understand without understanding one understands" might mean at least one of four things (depending, in part, on whether one is speaking of understanding-why or objectual understanding). Regarding understanding-why, it could mean either:

(c1) understanding why *P* entails understanding ***that*** one understands why *P*; or

(c2) understanding why *P* entails understanding ***why*** one understands why *P*.

Regarding objectual understanding, it could mean either:

(c3) understanding *x* entails understanding ***that*** one understands *x*; or

(c4) understanding *x* entails understanding ***why*** one understands *x*.

In order to show that each of (c1) – (c4) is false, it is sufficient to show that (c1) and (c3) are both false. To illustrate, suppose that (c1) is false. Then, there is a possible situation in which some subject *S* understands why *P* but does not understand *that* she understands why *P*; and, since *S* does not understand *that* she understands why *P*, *S* does not understand *why* she understands why *P* (for one cannot understand *why* something is the case without understanding *that* it is the case). Hence, supposing that (c1) is false, there is a possible situation in which *S* understands why *P* but does not understand *why* she understands why *P*, and thus (c2) is also false.

Next, suppose that (c3) is false. Then, there is a possible situation in which a subject *S* understands some thing *x* but does not understand *that* she understands *x*; and, since *S* does not understand *that* she understands *x*, *S* does not understand *why* she understands *x* (for, again, one cannot understand *why* something is the case without understanding *that* it is the case). Hence, supposing that (c3) is false, there is a possible situation in which *S* understands *x* but does not understand *why* she understands *x*, and thus (c4) is also false. Therefore, if (c1) and (c3) are both false, then each of (c1) – (c4) is false.

With this in mind, consider both (c1) and (c3). First, regarding (c1), is it the case that understanding why *P* entails understanding *that* one understands why *P*? It seems that the answer is "no". To see why, suppose that Timothy is a neuroscience professor who has recently developed an accurate and comprehensive account of why humans dream. Moreover, he both believes and has good reason to believe that his account explains why humans dream. Thus, it seems plausible to say that Timothy understands why humans dream. Does it then follow that Timothy understands that he understands why humans dream? I submit that it does not. For, Timothy might fail to have any belief-ish type state (or cognitive-pro attitude) pertaining to the proposition <Timothy understands why humans dream>; and, just as knowing that *P* is typically thought to require some type of cognitive pro-attitude towards *P*, so too understanding that *P* seems to require some type of cognitive pro-attitude towards *P* (e.g., if Julie understands that the meeting starts at 3:00PM, then she must have some type of cognitive pro-attitude towards the proposition <the meeting starts at 3:00PM>).[9] Given this, it is not hard to imagine that Timothy understands why humans dream (for the reasons mentioned above) but yet fails to understand that he understands why humans dream: namely, by failing to have any type of cognitive pro-attitude towards the proposition <Timothy understands why humans dream>. This failure could result from a number of reasons: e.g., perhaps Timothy has never considered the proposition <Timothy understands why humans dream>, or perhaps even if he were to consider it, he would remain noncommittal about its truth-value due to having little interest in the truth conditions for statements of the form "*S* understands why *P*"; or perhaps he is strongly committed to a view of

[9] I use the phrases "belief-ish type state" and "cognitive pro-attitude" to remain noncommittal about the precise type of mental state that is required for knowledge. While some have provided reason to think that knowledge does not entail the state that is typically picked out by the term "belief" (Beebe 2013; Murray et al. 2013), others have provided reason to think that knowledge at least entails some type of bare cognitive pro-attitude. For instance, Buckwalter et al. (2013) note, "A thin belief is a bare cognitive pro-attitude. To have a thin belief that P, it suffices that you represent that P is true…. Thinly believing P involves representing and storing P as information. It requires nothing more. In particular, it doesn't require you…to *explicitly avow* or *assent to* the truth of P."

understanding according to which he does *not* understand that he understands why humans dream (e.g., it might be a view of understanding on which it is nearly impossible for anyone to understand that she understands why *P*), thus leading Timothy to strongly hold that he does *not* understand that he understands why humans dream.

Therefore, it seems that neither (c1) nor (c2) is true: that is, understanding why *P* does not entail understanding *that* one understands why *P* (as illustrated in the case with Timothy), and given that understanding-why entails understanding-that (as argued above), understanding why *P* does not entail understanding *why* one understands why *P*.

A similar argument can be made with respect to (c3). Suppose that Julie attends a short series of lectures on the redundancy principle: a relatively straightforward principle in educational psychology which states that in multimedia learning environments, people learn better when graphics are presented with a corresponding spoken narration than when they are presented with both a corresponding spoken narration and corresponding written text (Mayer 2009). Given this, suppose that while attending the series of lectures, Julie learns all of the key claims regarding the redundancy principle, as well as how the principle has been used to both design educational assessments and guide new research. Accordingly, it seems plausible to say that Julie understands the redundancy principle. However, for reasons similar to those mentioned in the case with Timothy, it does not follow that Julie understands that she understands the redundancy principle. For, she might have no cognitive pro-attitude towards the proposition <Julie understands the redundancy principle>, which again might result from the fact that she has never considered the proposition, or from the fact that even if she were to consider it, she would remain noncommittal about its truth-value due to having little interest in the truth conditions for statements of the form "*S* understands *x*".

Given this, it seems that neither (c3) nor (c4) is true: that is, understanding *x* does not entail understanding *that* one understands *x* (as illustrated in the case with Julie), and given that understanding-why entails understanding-that (as argued above), understanding *x* does not entail understanding *why* one understands *x*. Thus, in sum, each of (c1) – (c4) is false, and hence neither understanding-why nor objectual understanding is transparent.

At this point, one might object that although (c1) and (c3), and thus (c2) and (c4), are false, there are weaker versions of (c1) and (c3) that are true, namely:

(c1′) understanding why *P* entails *being in a position to understand* that one understands why *P*; and

(c3′) understanding *x* entails *being in a position to understand* that one understands *x*.

In response, however, I submit that (c1′) and (c3′) fare no better in differentiating understanding from knowledge than do (c1) – (c4), since (c1′) and (c3′) seem to be no more or less plausible than the knowledge version of each claim; that is, they seem no more or less plausible than either of the following claims:

(c1″) knowing why *P* entails *being in a position to know* that one knows why *P*; or

(c3″) knowing *x* entails *being in a position to know* that one knows *x*.

To illustrate, while it may seem plausible to claim that the proposition <Timothy *understands* why humans dream> entails the proposition <Timothy is *in a position to understand* that he understands why humans dream>, it seems equally plausible to claim that the proposition <Timothy *knows* why humans dream> entails the proposition <Timothy is *in a position to know* that he knows why humans dream>. Similarly, while it may seem plausible to claim that the proposition <Julie *understands* the redundancy principle> entails the proposition <Julie is *in a position to understand* that she understands the redundancy principle>, it seems equally plausible to claim that the proposition <Julie *knows* the redundancy principle> entails the proposition <Julie is *in a position to know* that she knows the redundancy principle>. Thus, it is difficult to see how either (c1′) or (c3′) could be used to differentiate either objectual understanding from objectual knowledge, or understanding-why from knowledge-why.

4. Quantitative species

Given the above arguments, one might conclude that understanding-why is just knowledge-why. However, this conclusion faces an important problem. Namely, it does not account for the fact that it makes sense to ask, “I realize that *S* knows why *P*, but does *S* understand why *P*?” To illustrate, Pritchard (2009) offers an example that is meant to show that understanding-why can indeed be differentiated from knowledge-why. He asks us to imagine that his house burns down, that he knows and understands why it burned down, and that he knows that it burned down due to faulty wiring. Subsequently, Prichard tells his young son why the house burned down. Prichard (2009) then writes:

> [My son] has no conception of how faulty wiring might cause a fire, so we could hardly imagine that merely knowing this much suffices to afford him understanding of why his house burned down. Nevertheless, he surely does know that his house burned down because of faulty wiring, and thus also knows why his house burned down.

This accords with the fact that while one might intelligibly ask, "I realize that his son knows why the house burned down, but does he understand why it burned down?", it would be quite awkward to ask, "I realize that his son understands why the house burned down, but does he know why it burned down?" The reason appears to be that, if his son understands why the house burned down, then he also knows why it burned down; thus there is no reason to further inquire into the matter.

Given the above line of thought, I propose that while all instances of understanding-why are instances of knowledge-why, there may be instances of knowledge-why that are not instances of understanding-why. Or, in other words, understanding-why is a *species* (or type) of knowledge-why. There are, however, at least two ways in which some state (or activity, entity, or property) can be a species of another state (or activity, entity, or property). Compare the states of surfing and moving. Surfing is a species of moving since every instance of surfing is an instance of moving. The question then arises as to which feature(s) might be used to differentiate surfing from moving. One such feature is that of *standing on a board.* While every instance of surfing is an instance of standing on a board, not every instance of moving is an instance of standing on a board. Indeed, there are many instances of moving that are not instances of standing on a board (e.g., instances of walking, jumping, climbing, or twirling).

Next, compare the relation between surfing and moving to that between exhaustion and tiredness. Just as surfing is a species of moving, exhaustion is a species of tiredness, since every instance of exhaustion is an instance of tiredness. Thus, the question arises as to which feature(s) might be used to differentiate exhaustion from tiredness. One plausible candidate is the feature of *needing rest beyond some threshold*. There is an important difference, however, between the way in which *needing rest beyond some threshold* helps to differentiate exhaustion from tiredness and the way in which *standing on a board* helps to differentiate surfing from moving. In particular, unlike the relationship between surfing and moving (where every instance of surfing, but not every instance of moving, is an instance of standing on a board), the difference between exhaustion and tiredness is that exhaustion involves a *greater need of rest* than that which mere tiredness involves (and yet both exhaustion and tiredness involve at least some need of rest). I call any relation thus resembling that between exhaustion and tiredness a relation of "being a *quantitative species* of". Other examples in which one state is a quantitative species of another state include *being massive* and *being large*, *longing for* and *desiring*, *being an old person* and *being a person*, and *being proficient* and *being competent*.

How might the notion of *quantitative species* be more formally described? I propose the following account (call it "QSPEC"):

QSPEC For any x and any y, x is a *quantitative species* of y just in case:

(q1) there is some magnitude m and some set of features ϕ, and there are thresholds $t1$ and $t2$, such that: something is an instance of x if and only if it both has the features that are members of ϕ and has *m* beyond *t2*, and something is an instance of *y* if and only if it both has the features that are members of ϕ and has *m* beyond *t1* (where $t2 > t1$).

According to QSPEC, exhaustion is a quantitative species of tiredness (as it should be), since both exhaustion is a species of tiredness, and there is some magnitude (i.e., needing rest) that satisfies (q1) with respect to exhaustion (as *x*) and tiredness (as *y*).

Secondly, given the use of ϕ in QSPEC, there are at least two ways in which one state might be a quantitative species of another state. In certain cases, such as the relationship between exhaustion and tiredness, ϕ is just the empty set, since there are no features (other than that of needing rest beyond some threshold) that are needed to describe the relationship between exhaustion and tiredness. By contrast, consider the relationship between *being an old person* and *being a person*. According to QSPEC, being an old person is a quantitative species of being a person. However, the relationship between being an old person and being a person cannot be wholly described in terms of some magnitude *m*. For instance, it is false that both: something is an instance of being an old person if and only if it is an instance of having an age beyond threshold *t2*, and something is an instance of being a person if and only if it is an instance of having an age beyond threshold *t1* (where $t2 > t1$); for, many rocks and trees are of an age beyond any such threshold, and yet no rock or tree is either an old person or a person. Rather, something is an instance of being an old person if and only if it both is an instance of some feature (i.e., being a person) and has an age beyond *t2*, while something is an instance of being a person if and only if it both is an instance of being a person and has an age beyond *t1* (where $t2 > t1$).

Thirdly, according to QSPEC, surfing is *not* a quantitative species of moving. To see why, consider the following two statements from (q1):

(q1.1) something is an instance of *x* if and only if it both has the features that are members of ϕ and has magnitude *m* beyond threshold *t2*; and

(q1.2) something is an instance of *y* if and only if it both has the features that are members of ϕ and has magnitude *m* beyond threshold *t1*.

For the purpose of reductio, suppose that surfing is a quantitative species of moving. Then, there is some magnitude *m* and some set of features ϕ, and there are thresholds *t1* and *t2*, such that both (q1.1) is true with respect to surfing (as *x*) and (q1.2) is true with respect to moving (as *y*). First consider (q1.1). If (q1.1) is true with respect to surfing (as *x*), then ϕ has as one of its members the feature of standing on a board, since standing on a board is a necessary feature of surfing. Given this, (q1.2) can be reformulated with respect to moving (as *y*) in the following manner:

(q1.2′) something is an instance of moving if and only if it both has the features that are members of ϕ (one of which is the feature of standing on a board) and has magnitude *m* beyond threshold *t1*.

This creates a problem for the claim that surfing is a quantitative species of moving. In particular, it follows from (q1.2′) that all instances of moving are instances of standing on a board. But given that some instances of moving are *not* instances of standing on a board (e.g., jogging), (q1.2′) is false. Therefore, although surfing is a species of moving, it is not a *quantitative* species of moving.[10]

5. Understanding-why is a quantitative species of knowledge-why

In light of the above discussion, I propose that understanding-why is a *quantitative* species of knowledge-why. To illustrate, consider again the case in which Pritchard's son knows why, but does not understand why, the house burned down. In order for Pritchard's son to move from a state of merely knowing why the house burned

[10] According to QSPEC, the extension of *being a quantitative species of* excludes certain pairs of states that one might think it should include. For instance, given QSPEC, neither being a teenager nor being a young person is a quantitative species of being a person. Thus, suppose that we prefer an alternate account of quantitative species (call it "ALT") on which: (i) both being a young person and being a teenager *are* quantitative species of being a person, and (ii) every member of the extension of *being a quantitative species of* as determined by QSPEC remains a member of the extension when determined by ALT. Then, such an account can be constructed by beginning with the account described in terms of QSPEC and then replacing all references to *thresholds* with references to *ranges*, as follows:

ALT For any x and any y, x is a *quantitative species* of y just in case:

(q1′) there is some magnitude m and some set of features ϕ, and there are ranges $[r_x, r_y]$ and $[r_w, r_z]$, such that: something is an instance of x if and only if it both has the features that are members of ϕ and has m within range $[r_x, r_y]$, and something is an instance of y if and only if it both has the features that are members of ϕ and has m within range $[r_w, r_z]$ (where $[r_x, r_y]$ is a proper sub-range of $[r_w, r_z]$).

Given that none of the examples discussed in the main text requires any consideration of magnitudes *within some range*, but only magnitudes *beyond some threshold*, I restrict my discussion of quantitative species so that it centers on thresholds rather than ranges, for ease of reading. However, all of the arguments made in the main text using QSPEC can likewise be made using ALT.

down towards a state of understanding why it burned down, Pritchard's son needs only to move from knowing fairly basic explanatory propositions (e.g., propositions of the form <*R* explains *P*>) to knowing more complex propositions that integrate new explanatory information with explanatory information that has already been obtained (e.g., by moving from knowing that <*R* explains *P*> to knowing that <*R* explains *P*, and *Q* is an explanatory intermediate between *R* and *P*>).

To see how this pertains to the above account of quantitative species, let Σ*SP* be the set of all explanatory propositions (pertaining to *P*) that some subject *S* knows to be true. Then, the relevant magnitude *m* is just how well Σ*SP* explains *P*. Thus, the view that understanding-why is a quantitative species of knowledge-why can be described more formally as follows:

> **The Quantitative Account:** For any subject *S*, and any proposition *P*, there is some magnitude *m* (i.e., how well Σ*SP* explains *P*), and there are thresholds *t1* and *t2*, such that both: something is an instance of *S understanding why P* if and only if it has *m* beyond *t2*, and something is an instance of *S knowing why P* if and only if it has *m* beyond *t1* (where *t2* > *t1*).

To illustrate, contrast the Quantitative Account with an account proposed by Stephen Grimm. Grimm (2006) argues that understanding-why differs from knowledge-why in that there is a necessary psychological component of understanding-why that is not necessary for knowledge-why. Specifically, while both understanding-why and knowledge-why require an "act of assent (a 'saying Yes' of some kind)", Grimm (2006, p. 515) builds on ideas from Brian Skyrms (1980) and James Woodward (2003) to argue that understanding-why (but not knowledge-why) requires the further “act of ‘grasping’”. Grimm writes:

> [W]hen trying to offer an account of understanding the notion of *grasping* arises almost irresistibly. Moreover, when we grasp some claim we are apparently doing something significantly different from merely saying Yes to it. In reply, it might be said that this attitude of grasping is nothing significantly different from an act of assent: it just involves *more* assent. But it seems clear that one can pile up assents as high as you like without getting a grasping.
>
> [G]rasping involves having an ability to anticipate that “wiggling” one variable will characteristically lead to a “wiggling” of another variable. Moreover, if you lack the ability to see how things are connected in this way, however— specifically, if you lack the ability to see how a change in the value of one variable will lead to a change in the value of another variable—then although you might justifiably assent to the fact (based on testimony, etc.) that the two variables depend on one another, there remains an important sense in which you will not have grasped how they depend on one another (Grimm 2006, pp. 532-533).

To see how Grimm’s (2006) proposal differs from the Quantitative Account, consider again the case in which Pritchard’s house burns down due to faulty wiring (Pritchard 2009). Pritchard argues that if he were to tell his young son why the house burned down, his son would know why the house burned down (i.e., by knowing that it burned down because of faulty wiring), but he would not understand why it burned down (since he "has no conception of how faulty wiring might cause a fire"). Following Grimm's proposal, one might argue that the reason that Pritchard's son knows why, but does not understand why, the house burned down is that although his son assents to the relevant proposition (i.e., that the house burned down because of faulty wiring), he does not grasp the relation between the two variables, i.e., *the faulty wiring* and *the house burning down*. For, to grasp this, Pritchard's son must be able "to see how a change in the value of one variable will lead to a change in the value of another variable". Hence, according to Grimm's view, the key difference between understanding-why and knowledge-why can be described in terms of some feature (i.e., that of grasping) such that every instance of understanding-why, but not every instance of knowledge-why, is an instance of grasping. Or, in other words, the psychological state that is required for mere knowledge-why (i.e., assenting) is qualitatively different from one of the psychological states that is required for understanding-why (i.e., grasping).

Is this the correct conclusion to draw? I submit that it is not. For, in order to differentiate a case in which one merely knows why the house burned down (e.g., by merely knowing that it burned down because of faulty wiring) from a case in which one understands why it burned down, there is no need to appeal to some notion of grasping that is qualitatively different from any feature that is necessary for knowledge-why. To see this, let *HOUSE* be the proposition <the house burned down>; let *WIRE* be the proposition <the wiring was faulty>; let *INSUL* be the proposition <the wires melted through the insulation and ignited the bedroom wall>; and let *RESIST* be the proposition <there was a high resistance to the electrical current in the wires>. Given this, a subject *S* might move from a state of merely knowing why the house burned down to a state of understanding why it burned down in the following way. First, *S* might merely know why *HOUSE* (by merely knowing that *WIRE* explains *HOUSE*). Second, *S* might come to know, not only *that WIRE* explains *HOUSE*, but also *why WIRE* explains *HOUSE* (e.g., by knowing that *WIRE* both explains *INSUL* and explains that *INSUL* explains *HOUSE*). Third, *S* might come to know, not only *that WIRE* both explains *INSUL* and explains that *INSUL* explains *HOUSE*, but also *why* this is the case (e.g., by knowing that *WIRE* both explains *RESIST* and explains that *RESIST* explains *INSUL*, and that *RESIST* explains that *INSUL* explains *HOUSE*). Thus, in sum, *S* knows that *WIRE* both explains *RESIST* and explains that *RESIST* explains *INSUL*, and that *RESIST* explains that *INSUL* explains *HOUSE*. Given this, it is hard to see how *S* does not also *understand why* the house burned down.

Thus, contrary to Grimm's (2006) claim, the move from merely knowing why *P* to understanding why *P* does not require a move from one type of mental state (e.g., merely assenting) to a qualitatively different type of mental state (e.g., grasping). Rather, it only involves a move from merely knowing some fairly basic explanatory proposition to knowing more complex explanatory propositions that integrate new information with information that has already been obtained.

One reason that might motivate Grimm's claim that "one can pile up … as high as you like" instances of the relevant state that is required for knowledge-why (which according to Grimm is mere assenting)[11] and never reach any instances of the relevant

[11] I use the phrase "relevant state" to indicate that the state is not merely one that is required for knowledge-why, but is also one that is used in attempt to differentiate knowledge-why from understanding-why (e.g., in the way that Grimm uses the states of *mere assenting* and *grasping*). Thus, for the purposes of this paper, the state of *having*

state that is required for understanding-why (which according to Grimm is grasping) is that a fairly obvious way of piling up such instances does not move one from merely knowing why *P* to understanding why *P*. For instance, if a subject *S* were to merely know why *P* by merely knowing that *P* because *R*, then *S* would not move towards an understanding of why *P* by merely coming to know more propositions of the form <*P* because *R*> (e.g., by coming to know that *P* because *M*, that *P* because *Q*, that *P* because *E*, and so on). For, coming to know more and more propositions of this form might occur without *S* coming to know any of the relevant *relations* between the propositions. Thus, in this sense, Grimm is correct in saying that piling up many instances of the relevant state that is required for knowledge-why does not move one any closer to a state of understanding-why.

However, the Quantitative Account is importantly different from this approach. As illustrated above, in order for one to move from merely knowing why *P* to understanding why *P*, one must move from knowing fairly basic explanatory propositions (e.g., that *R* explains *P*) to knowing more complex explanatory propositions: namely, explanatory propositions that integrate the relations between new information and information that has already been obtained (e.g., by moving from merely knowing that <*R* explains *P*> to knowing that <*R* both explains *Q* and explains that *Q* explains *W*, and that *Q* explains that *W* explains *P*>).[12]

justification is not such a state, since neither Grimm nor I attempt to use the state to differentiate knowledge-why from understanding-why.

[12] Also note that in order for a subject *S* to know some explanatory proposition (e.g., the proposition <*R* explains *P*>), *S* must have at least some knowledge of the proposition's constituents. For instance, suppose that Mick is trying to discover why his car will not

With this in mind, consider how certain aspects regarding the relationship between understanding-why and knowledge-why can be described in terms of the Quantitative Account. According to the Quantitative Account, something is an instance of *S understanding why P* if and only if it is an instance of Σ*SP* explaining *P* beyond threshold *t2*, and something is an instance of *S knowing why P* if and only if it is an instance of Σ*SP* explaining *P* beyond threshold *t1* (where *t2* > *t1*, and Σ*SP* is the set of explanatory propositions pertaining to *P* that *S* knows). Thus, for *S* to understand why *P*, Σ*SP* must explain *P* more fully than is required for *S* to know why *P*.

In light of this, suppose that Pritchard's (2009) intuitions are correct that his son knows why, but does not understand why, the house burned down. Further suppose that *S* is Pritchard's son, and *P* is the proposition <the house burned down>. Then, Σ*SP* is the set of every proposition that both is known by Pritchard's son and provides at least some explanation of why the house burned down. Given this, the idea that Pritchard's son knows why, but does not understand why, the house burned down can be described by saying that while the explanation provided by Σ*SP* (as to why the house burned down) is good enough to establish that Pritchard's son knows why the house burned down, it is not good enough to establish that he understands why the house burned down. And in order for Pritchard's son to move from merely knowing why the house burned down to understanding why it burned down, Pritchard's son need not acquire any qualitatively new feature that is necessary for understanding-why but not for knowledge-why. Rather,

---

start, and a mechanic informs him that a break in the distributor rotor explains why the car will not start. If this interaction with the mechanic is to provide Mick with any knowledge of the explanatory proposition <a break in the distributor rotor explains why the car will not start>), then Mick must have at least some knowledge of what a distributor rotor is.

what is needed is for additional explanatory propositions to be added to Σ*SP* that would together increase how well Σ*SP* explains why the house burned down.

This suggests an interesting parallel regarding the relation between understanding-why and knowledge-why and other relations in which one state is a quantitative species of another state. Consider, for instance, the relationship between *competency* and *proficiency*. Like understanding-why and knowledge-why (but unlike surfing and moving) the difference between competency and proficiency can be described by saying that every instance of proficiency is an instance of being able to successfully engage in some action or process beyond threshold *t2*, while every instance of competency is an instance of being able to successfully engage in some action or process beyond threshold *t1* (where *t2* > *t1*). To illustrate, suppose that Nicolas is competent, but not proficient, in the Icelandic language. Then, Nicolas might be able to comprehend simple Icelandic texts, converse in colloquial Icelandic so long as people speak slowly, and distinguish grammatical and ungrammatical constructions of fairly rudimentary Icelandic sentences. However, in order for Nicolas to become proficient in the Icelandic language, he must develop a greater ability to successfully engage in relevant actions and processes pertaining to the Icelandic language; e.g., he would need to be able to comprehend more complex Icelandic texts, converse in Icelandic on more complex topics and at a more rapid pace, and distinguish between grammatical and ungrammatical constructions of fairly complex Icelandic sentences. Thus, just as a move from merely knowing why *P* to understanding why *P* can be described in terms of an increase in how well Σ*SP* explains *P*, a move from mere competency to proficiency can be described in terms of an increase in one's ability to successfully engage in relevant actions or processes.

A similar point can be made with respect to other cases in which one state is a quantitative species of another state: e.g., the move from merely desiring something to longing for it, the move from merely being tired to being exhausted, and the move from merely being large to being massive.

Moreover, the idea that an instance of *S understanding why P* requires that Σ*SP* explain *P* more fully than is required for an instance of *S knowing why P* helps to explain why it is that one can intelligibly ask, "I realize that Pritchard's son knows why the house burned down, but does he understand why it burned down?", but yet cannot intelligibly ask, "I realize that Pritchard's son understands why the house burned down, but does he know why it burned down?" A similar point can be made with respect to other cases of quantitative species. For instance, the idea that proficiency requires a greater ability to successfully engage in relevant actions or processes than is required for competency helps to explain why it is that one can intelligibly ask, "I realize that Nicolas is competent in the Icelandic language, but is he proficient in it?", but yet cannot intelligibly ask, "I realize that Nicolas is proficient in the Icelandic language, but is he competent in it?"

Lastly, consider comparative claims regarding understanding-why and knowledge-why. For example, we often make claims of the form "*S* understands why *P* better than *S** does." While I remain noncommittal about the acceptability of such claims for the purposes of this paper, it is nonetheless worthwhile to consider how the Quantitative Account might handle such claims. In brief, according to the Quantitative Account, comparative claims regarding knowledge-why and understanding-why can be described in terms of how well ΣSP explains P, as follows: for any two subjects, S and S^*, S understands (alternatively: knows) why P better than S^* does just in case the set ΣSP explains P better than ΣS^*P does. Something similar can be said for other cases of quantitative species. For instance, comparative claims regarding competency and proficiency can be described in terms of one's ability to successfully engage in relevant actions or processes, as follows: for any two subjects, S and S^*, S is more proficient (alternatively: more competent) at x than is S^* just in case S's ability to successfully engage in relevant actions or processes (pertaining to x) is greater than S^*'s ability to do so.

6. Conclusion

To summarize, I looked at three prominent proposals that have been used in attempt to differentiate understanding from knowledge. I argued that once relevant distinctions are made, the purported differences between understanding and knowledge disappear (Sections 1 – 3). Then, in Section 4, I considered a phenomenon that *does* seem to suggest a difference between understanding and knowledge, and I used this phenomenon to motivate a new account of the relationship between understanding-why and knowledge-why. Finally, in Section 5, I argued that understanding-why is a quantitative species of knowledge-why, and I examined how this proposal serves to explain certain phenomena regarding the relationship between understanding-why and knowledge-why.

This chapter paves the way for a closer examination of several questions. How does the Quantitative Account apply to other forms of knowledge and understanding;

e.g., can a similar account be constructed for the relationship between objectual knowledge and objectual understanding? What exactly is required for a proposition to be an *explanatory* proposition; or, in other words, what exactly is required for a proposition to be a member of a set such as Σ*SP*? Are all explanatory propositions equal regarding their capacity to move a subject from mere knowledge-why to understanding-why? It is to such questions that I will turn in subsequent work.

CHAPTER 2

Understanding-Why and Explanation

Overview

As many authors have suggested, there seems to be a close relationship between understanding why some fact is so and having an explanation of why it is so. In this chapter, I develop a novel account of understanding-why that is a refinement of this basic idea. To do this, I look largely to the literature on dispositions and higher-level properties, since problems discussed in that literature can be used to motivate such an account. I begin by providing a broad overview of several closely related notions, including dispositions, multiple realizability, causation, and explanation. I then introduce a notion that I call "*causal guiding explanation*," and I contrast it with the notion of *objective causal explanation*. Next, I argue that a person *S* has at least some understanding of why *P* if and only if *S* has a causal guiding explanation of why *P*. Finally, I provide a broad outline of how one might extend the proposed account of what it is to have at least some understanding-why (given in terms of causal guiding explanations) to further account for the *gradability* of understanding-why.

1. Introduction

One of the central aims of human inquiry is to understand why certain facts are so. We seek to understand why planets move in ellipses, why caffeine increases alertness, why worker bees engage in self-sacrifice, and why humans wage wars, to give just a few examples.

In this chapter, I focus on understanding why *P* is true, where *P* is a proposition that represents an *event*. Contemporary discussions on the nature of understanding-why are often embedded within discussions on the nature of explanation (Hempel 1965; Achinstein 1983; Kitcher 1989; Salmon 1998; Woodward 2003). This is perhaps not surprising, since there seems to be a close connection between understanding why some empirical proposition is so and being able to explain why the proposition is so. Indeed, as a first approximation towards an account of understanding-why, the following view, APPROX, might seem like a plausible place to start:

> APPROX For any subject *S*, and any empirical proposition *P*, *S* has at least some understanding of why *P* if and only if there is some proposition (or set of propositions) *Q*, such that *S* understands that *Q* at least partially explains why *P*.

This view on its own, however, leaves much to be said about the nature of understanding-why. For instance, the claim that some set *E* explains why *P* is so might be spelled out in a variety of ways depending on one's view of explanation.

Given this, how might one proceed in constructing a substantive and systematic account of the nature of understanding-why? The strategy that I employ begins by taking APPROX as a provisional starting point towards an account of understanding-why, and then proceeds to consider how different types of explanations might play different roles in understanding-why. While there is a significant body of literature that is relevant to determining whether something like APPROX is correct (Trout 2002, 2007; Lipton 2009; Khalifa 2012, 2013; Strevens 2013), far less has been written on how APPROX might be further refined in terms of specific types of explanations. Thus, for the purposes of this paper, I look largely to the literature on dispositions and higher-level properties, since problems discussed in that literature can be used to motivate a new positive account of understanding-why that is a refinement of APPROX.

I begin, in Section 2, by providing a broad overview of several closely related notions, including dispositions, multiple realizability, causation, and explanation. In Section 3, I then introduce a notion that I call "*causal guiding explanation*," and I contrast it with the notion of *objective causal explanation*. Given this distinction, I proceed to argue that a person *S* has at least some understanding of why *P* (where *P* is an empirical proposition) if and only if *S* has a causal guiding explanation of why *P*. Finally, in Section 4, I provide a broad outline of how one might extend the proposed account of what it is to have at least some understanding-why (given in terms of causal guiding explanations) to further account for the *gradability* of understanding-why.

2. Dispositions, multiple realizability, causation, and explanation

To begin, consider how a person might acquire at least some understanding of why a given proposition is so. In many cases, a person might do so by understanding that the fact that the proposition is so is a manifestation of some disposition. For instance, one might have at least some understanding of why a particular vase broke by understanding that its breaking was a manifestation of its *fragility*. Similarly, one might have at least some understanding of why a particular sailing vessel capsized by understanding that its capsizing was a manifestation of its *instability*. What exactly is it about certain dispositions that makes them relevant to understanding-why? To answer this, consider more closely the nature of dispositions.

Dispositions are commonly characterized in terms of their manifestations and triggers: e.g., in the way that *fragility* is typically characterized in terms of the disposition to *break* (i.e., the manifestation) when *struck* (i.e., the trigger) (Lewis 1997; Molnar 2003; Bird 2005; Cross 2005; Hawthorne & Manley 2005; Schrenk 2010). While there is still significant debate over the precise way in which dispositions should be analyzed (and even whether they can be successfully analyzed)[13], it is widely accepted that dispositional ascriptions at least entail certain counterfactual conditionals. For instance, Armstrong writes:

> Let us consider a statement such as "This glass is brittle" said truly of an unstruck glass. It is uncontroversial that this statement entails a counterfactual statement along these general lines: If this glass had been suitably struck, then this striking would have caused the glass to shatter. We need not worry about the *detail* of this statement. To do so would deflect attention from the more fundamental ontological issues that we wish to consider (Armstrong 1996, p. 15).

With this in mind, Armstrong goes on to argue for what he and others term "the truthmaker principle": i.e., the widely held principle that for every (contingent) statement that is true in our world, there is something in our world that makes the statement true (or that grounds its truth). Employing this principle, Armstrong writes, "We now apply the truthmaker principle to the counterfactual truth about the glass. What makes this truth true, what is its ontological ground?" (1996, p. 15). In response, Armstrong suggests that we should look among the *properties* of the unstruck glass to find what it is that makes the counterfactual statement true. This raises the question: How does the property (or property-complex) that makes the counterfactual associated with, say, brittleness true relate to the property of brittleness itself? There has been much debate over this question.

---

[13] Some, for instance, consider it unlikely that there will be an analysis of dispositions that will successfully address all of the relevant counterexamples, and thus instead resort to the task of *explaining* how dispositions *ground* certain counterfactuals, as opposed to offering an *analysis* of dispositions *in terms* of counterfactuals (e.g., Bird 1998; Molnar 1999).

Armstrong (1973, p. 14-16) argues that a disposition just *is* the associated truthmaker property (i.e., the property or property-complex that makes the associated counterfactual true). Thus, the fragility of a vase just is the relevant microphysical property or property-complex had by the vase. By contrast, Prior et al. (1982) propose that dispositions are to be defined functionally. On this view, a disposition is a particular type of second-order property, i.e., a property of having some first-order property or other that plays the causal role characterizing the disposition.

As David Lewis notes, the former view (that a disposition just is the associated truthmaker property) has the benefit of providing "a straightforward account of the role of dispositions in causal explanation," but it has the downside that "what we would offhand think was one disposition, fragility, turns out to be different properties in different possible cases – and, very likely, in different actual cases" (1997, p. 151). The latter view, on the other hand, has the benefit of providing an account on which "fragility is indeed a single property common to all fragile things," but it has a downside in that "if fragility is the second-order property, it is far from clear how it plays a role in causal explanation" (1997, p. 151). Lewis continues:

> When the struck glass breaks, do we want to say that the breaking is caused *both* by the second-order property which is the fragility *and* by whatever first-order property is the causal basis for the fragility in that particular case? It is not a case of overdetermination, after all! But neither should we want to say, as Prior *et al.* do, that fragility is causally impotent (1997, pp. 151-152).

So, we have a dilemma:

- On the one hand, we can adopt the view that a disposition is identical to the microphysical property or property-complex that realizes the disposition. This would allow us to account for the (seemingly plausible) idea that dispositions can be causally efficacious; but it would give us reason to reject the (also seemingly plausible) idea that a disposition, say, fragility, is a single property that is had by all fragile things.

- On the other hand, we can adopt the view that a disposition is a particular type of second-order property, i.e., a property of having some first-order property or other that plays the causal role characterizing the disposition. This would allow us to account for the (seemingly plausible) idea that a disposition, say, fragility, is a single property that is had by all fragile things; but it would give us reason to reject the (also seemingly plausible) idea that dispositions can be causally efficacious.

A similar dilemma arises with respect to multiply-realizable properties in general; e.g., properties such as *being in pain* and *being afraid*. For, if we adopt the (seemingly plausible) view that properties such as *being in pain* and *being afraid* can be realized by a variety of distinct neurophysiological properties, then it is hard to avoid the (seemingly implausible) idea that the causal efficacy of such multiply-realizable properties is excluded by the causal roles played by their neurophysiological realizers (e.g., see Kim 1998; Bontly 2001; Block 2003; Walter 2008). The question therefore becomes: Which of the following two (seemingly plausible) claims, i.e., (d1) or (d2), should we reject?

(d1) Dispositional and multiply-realizable properties (e.g., fragility, flammability, elasticity, pain, fear, et cetera) can be causally efficacious.

(d2) Dispositional and multiply-realizable properties are distinct from their lower-level realizers (e.g., in the way that *fragility* is distinct from both the property of *having weak intermolecular bonding* and the property of *having a nonuniform atomic structure*, and yet is realizable by either; or in the way that *being in pain* is distinct from the various neurophysiological properties that are each sufficient to realize *being in pain*).

I submit that we should remain committed to (d2) and reject (d1). For, the seeming plausibility of (d1) can be accounted for in terms of a closely related claim that does not stand in tension with (d2): namely, the claim that

(d1´) dispositional and multiply-realizable properties are often *relevant to obtaining an objective causal explanation of why P* even if the properties are not themselves *causally efficacious*.

Given the prima facie close relationship between properties that are causally efficacious and properties that are relevant to obtaining a causal explanation, claim (d1´) can help account for our initial hesitation to deny causal efficacy to dispositional and multiply-realizable properties, and can do so without ruling out (d2).[14] So, the question arises:

[14] With respect to dispositions, Lewis also favors the view that a disposition, say fragility, is a second-order property that is had by all fragile things. He thus writes, "If forced to

Exactly how might a dispositional or multiply-realizable property be *relevant to obtaining an objective causal explanation of why P* even if the property is not itself *causally efficacious*? I now turn to address this question in detail.

3. Causal guiding explanations

Suppose that Jamie purchases a small statue while on vacation, and when she arrives home, she opens her suitcase and finds that the statue is broken. Jamie has no idea why the statue broke, so she asks her roommate (who happens to be an avid collector of statues) why the statue broke. After examining the statue, Jamie's roommate replies that it broke because it was made of a very fragile material. Given this, it seems plausible to say that while Jamie initially had no understanding of why the statute broke (prior to talking with her roommate), Jamie acquired at least some understanding of why the statue broke by acquiring an understanding (via her roommate) that the statue broke because it was fragile. But if the fragility of the statue does not figure into the full objective causal nexus that was responsible for the fact that the statue broke, then in what sense can we say that Jamie has at least some understanding of why the statue broke by understanding that it broke *because* it was fragile?

One way that we might make sense of this, even while granting that the lower-level (e.g., microphysical) properties that realize the statue's fragility (but not the fragility

choose, I would side with Prior against Armstrong; and I would dodge the overdetermination-or-impotence issue by appeal to some fancy and contentious metaphysics" (1997: 152). My strategy, however, is not to appeal to any fancy or contentious metaphysics, but instead to show that (d1) can be rejected in place of a closely related claim, i.e., (d1′), since (d1′) can still account for the underlying intuitions that result in our initial hesitation to reject (d1).

itself) figure into the full causal nexus that is responsible for the fact the statue broke, can be provided by way of a distinction between the notion of *objective causal explanation* and a notion that I call "*causal guiding explanation*." Suppose, for instance, that once all of the causal roles of the relevant structural and microphysical properties of the statue are provided, any appeal to the statue's fragility is causally superfluous, and thus can be left out when describing the full causal nexus that is responsible for the fact that the statue broke. Nevertheless, there remains a sense in which the statue's fragility might still be relevant to *guiding* a person to an objective causal explanation of why the statue broke. In particular, by acquiring an understanding that the statue broke because it was fragile, a person might acquire the ability to identify some set of propositions Φ, such that some member or other of Φ is at least partially causally responsible for the fact that the statue broke.[15] To illustrate, Φ might contain every proposition of the form, <The statue has property *x*>, where *x* is a potential realizer of fragility. For, by understanding that the statue broke because it was fragile, a person might acquire the ability to understand that *the statue's possession of some property (or other) that realizes fragility* figures into the full causal nexus that was responsible for the fact that the statue broke. Similarly, Φ could also (or instead) contain every proposition of the form, <*x* triggered the fragility of the statue>, where *x* is a potential trigger of the statue's fragility (e.g., the striking, twisting, or dropping of the statue). For, by understanding that the statue broke because it

[15] Or, more precisely, a person might acquire the ability to identify some set of propositions Φ, such that for at least some member *M* of Φ, the fact that *M* is so is at least partially causally responsible for the fact that the statue broke. For ease of reading, I often use phrases of the form, "*M* is at least partially causally responsible for *P*," as shortened versions of, "the fact that *M* is so is at least partially causally responsible for the fact that *P*."

was fragile, a person might acquire the ability to understand that *some fact (or other) that triggered the statue's fragility* figures into the full causal nexus that was responsible for the fact that the statue broke.

In saying that the proposition <The statue broke because it was fragile> might enable a person to identify some such set of relevant propositions, I mean only that the truth of the proposition, <The statue broke because it was fragile>, but not the truth of the proposition <The statue broke>, *guarantees* that there is such a set of propositions. Thus, the idea that a given proposition provides a person with a *causal guiding explanation* can be described more generally as follows (call the proposal "CGE"):

CGE For any propositions *C* and *P*, and any subject *S*, *C* provides *S* with a *causal guiding explanation* of why *P* if and only if both:

(e1) there is some set of propositions Φ, such that the truth of *C* (but not the truth of *P*) guarantees that for some member *M* of Φ, the fact that *M* is so is at least partially causally responsible for the fact that *P*;[16] and

[16] Note that condition (e1) does *not* say that there is a particular member of Φ, such that the truth of *C* guarantees that *that member* is at least partially causally responsible for the fact that *P*. To see how this differs from what (e1) actually says, consider how it might be the case that for some proposition *Q* and some set Δ, the truth of *Q* guarantees that something holds of some member or other of Δ, even if there is no particular member of Δ, such that the truth of *Q* guarantees that the thing holds of that particular member. For instance, the truth of the proposition, <Some person is courageous>, guarantees that the set of all people has at least some courageous person, even though there is no particular person in the set, such that the truth of the proposition guarantees that that particular person is courageous. Similarly, the proposition, <The statue broke because it was fragile>, guarantees that the set containing every proposition of the form, <The statue bears property *x*> (where *x* is a potential realizer of fragility), has some member *M* that is

(e2) *S* understands that *C*.[17]

Consider first the parenthetical phrase in condition (e1), i.e., the phrase, "but not the truth of *P*." This phrase is important, since, without it, a person could have a causal guiding explanation of *why P* merely by understanding *that P*. To illustrate, suppose that *P* is the proposition, <Jonathan is having chest pains>. Then, without the parenthetical phrase in (e1), a person could have a causal guiding explanation of *why* Jonathan is having chest pains just by understanding *that* Jonathan is having chest pains. For, there is some set Φ (i.e., the set containing every proposition *Q*, where the fact that *Q* is so might be at least causally responsible for the fact that Jonathan is having chest pains), such that the truth of the proposition, <Jonathan is having chest pains>, guarantees that some member of Φ is at least partially causally responsible for the fact that Jonathan is having chest pains. Thus, the parenthetical phrase in (e1) serves to rule out such cases and ensure that one *cannot* have a causal guiding explanation of *why P* merely by understanding *that P*.[18]

at least partially causally responsible for the fact that the statue broke, even though there is no member *M** of the set, such that the truth of the proposition, <The statue broke because it was fragile>, guarantees that *M** is at least partially causally responsible for the fact that the statue broke.

[17] I follow the widely held view that understanding-that is just knowledge-that (e.g., see Elgin 2007; Pritchard 2009; Grimm 2010). On this view, condition (e2) of CGE, could be equivalently formulated as, "*S* knows that *C*." However, CGE is also compatible with other views on the nature of understanding-that; e.g., the view that understanding-that is a proper species of knowledge-that.

[18] One might worry that there could be cases in which we want to say that a person understands why *P* by understanding that *C*, where *C* is distinct from *P* and does the relevant guaranteeing for some Φ, but *P* also does the relevant guaranteeing for Φ. Such cases need not create worry, however. For, they are in fact compatible with CGE. According to CGE, it need only be the case that for any such proposition *P*, and any such

Consider one further example that is relevant to satisfying conditions (e1) and (e2) of CGE. Suppose that for some lower-level property π, Samantha understands that

> *A1.* the fact that Jeremy's brain bears property π is *causally responsible* for the fact that Jeremy went to the hospital.

Given this, Samantha has a causal guiding explanation of why Jeremy went to the hospital. For, there is a set of propositions Φ (i.e., the singleton set containing the proposition, <Jeremy's brain bears property π>), such that (trivially) the truth of *A1* (but not the truth of the proposition, <Jeremy went to the hospital>) guarantees that some member of Φ is at least partially causally responsible for the fact that Jeremy went to the hospital. This satisfies condition (e1) of CGE. Moreover, condition (e2) of CGE is also satisfied, since Samantha understands that *A1*. This example illustrates that understanding some *objective causal explanation* (such as *A1*) is just a special case of having a *causal guiding explanation*.

At this point, one might wonder how CGE differs from the following two views of explanation:

- the view put forth by David Lewis (1986) according to which explanation consists in the provision of causal information; and
- the view put forth by Jackson and Pettit (1990) according to which the realization of a higher-order property *programs for* the appearance of some lower-order efficacious property, in the sense that the

---

causal guiding explanation *C*, there is *at least some* Φ^* (which, in the proposed example, would differ from Φ), such that *C* but not *P* does the relevant guaranteeing for Φ^*.

> instantiation of the higher-order property "ensured that there was an efficacious property in the offing" (p. 114).

To see how CGE differs from these two views, consider again the fragility of Jamie's statue. On Lewis's view, the proposition <The statue was fragile> can serve to explain why the statue broke, since the proposition provides information about the causal history of the breaking of the statue. Regardless of the merits of this view of explanation, the notion of a causal guiding explanation differs in an important way from this view. To illustrate, according to CGE, the proposition <The statue was fragile> cannot provide one with a causal guiding explanation of why the statue broke. In order for a proposition *C* to provide one with a causal explanation of why *P*, there must be some set of propositions Φ, such that the truth of *C* (but not the truth of *P*) guarantees that for some member *M* of Φ, the fact that *M* is so is at least partially causally responsible for the fact that *P*. Therefore, if *C* is to provide one with a causal explanation of why *P*, then the truth of *C* must guarantee the truth of *P*. However, the truth of the proposition <The statue was fragile> does *not* guarantee the truth of the proposition <The statue broke>, since many fragile things do *not* break. This result of CGE is in fact desirable. For, CGE is designed with the purpose of constructing an account of understanding-why. And just as the proposition <The statue was fragile> cannot provide one with a causal guiding explanation of why the statue broke, having an understanding that the statue was fragile is not sufficient for having an understanding of why the statue broke (since one might understand that the statue is fragile and yet think that its breaking had nothing to do with its fragility). By contrast, the proposition <The statue broke because it was fragile> may well provide one with a causal guiding explanation of why the statue broke, and accordingly, it is also the case that understanding that the statue broke because it was fragile is sufficient for having at least some understanding of why the statue broke. To reiterate, this is not a criticism of Lewis's view of explanation. It is only meant to show that Lewis's view of explanation differs from the notion of causal guiding explanation. Unlike Lewis's view, the notion of causal guiding explanation is not meant to capture all types of explanation. Rather, it is only meant to capture the types of explanation that are relevant to understanding-why.

Now, consider the view put forth by Jackson and Pettit (1990). According to their view, the realization of a higher-order property can help explain the occurrence of some event when the instantiation of the higher-order property ensures that there is a property that was causally efficacious in bringing about the event. For instance, the realization of fragility can help explain why the statue broke in that the instantiation of fragility ensures that there is some property that was causally efficacious in bringing about the breaking of the statue. The difference between this account and CGE is similar to the difference between Lewis's account and CGE. In short, while the instantiation of fragility might ensure that there is some property that has the causal efficacy to bring about the breaking of the statue, the instantiation of fragility does *not* guarantee the breaking of the statue, and thus it cannot provide one with a causal guiding explanation of why the stature broke. Again, this result is desirable, since the notion of causal guiding explanation is intended to capture only those types of explanation that are relevant to understanding-why (and for the reasons noted above, merely having an understanding that the statue was fragile is not sufficient for having an understanding of why the statue broke).

In light of the above discussion, I now offer the following account of what it is to have at least some understanding of why *P*, where *P* is a proposition that represents an event (call the account "UWHY"):

> **UWHY** For any subject *S*, and any event proposition *P*, *S* has at least some understanding of why *P* if and only if:
>
> > (f1) there is some proposition *C* that provides *S* with a causal guiding explanation of why *P*.

Thus, according to UWHY, Jamie has at least some understanding of why the statue broke by understanding that it broke because it was fragile, since the proposition <The statue broke because it was fragile> provides Jamie with a causal guiding explanation of why the statue broke. That is, there is some set of propositions Φ, such that the truth of the proposition <The statue broke because it was fragile> (but not the truth of the proposition <The statue broke>) guarantees that for some member *M* of Φ, the fact that *M* is so is at least partially causally responsible for the fact that the statue broke.

Similarly, according to UWHY, Samantha has at least some understanding of why Jeremy went to the hospital by understanding that *A1*, since *A1* provides Samantha with a causal guiding explanation of why Jeremy went to the hospital.[19]

---

[19] The proposed account, UWHY, also suggests a potential way forward in debates over the causal relevance of higher-level properties (e.g., see Kim 1998; Bontly 2001; Block 2003; Walter 2008). For instance, it allows one to hold that while the causal efficacy of some higher-level property π (say, the property of *having a headache*) might be excluded by the causal roles that are played by the lower-level property that realizes π (say, the property of *having such-and-such a brain state*), the property π may nonetheless have considerable importance in understanding-why (e.g., in a person's understanding of why William walked to the medicine cabinet). And, granting that there is close prima facie relationship between understanding-why and causal explanation, this importance can help account for our initial hesitation to deny causal efficacy to π without further qualification.

Finally, consider how UWHY relates to accounts of explanation that make no reference to causation, such as the unificationist account. Roughly, the unificationist account maintains that explanation is a matter of unifying a wide range of phenomena under a generalization and thereby showing how the phenomenon in question is an instance of a general pattern (e.g., see Friedman 1974; Kitcher 1989). To illustrate, suppose that the acceleration of some object *o* doubles. What might explain why *o*'s acceleration has doubled? One possible explanation is as follows (call the explanation "*D*"):

> *D*. The acceleration of *o* has doubled because: (i) the net force acting on *o* has doubled while *o*'s mass has remained the same, and (ii) the acceleration of any object is equal to the net force acting on the object divided by the object's mass.

Assuming, for the sake of simplicity, that the domain of discourse includes only medium-sized objects moving significantly slower than the speed of light, proposition *D* might be deemed as explanatory by a proponent of the unificationist account, since the argument pattern exemplified by *D* can be employed to unify a diverse range of phenomena regarding the movements of all objects, and thus it reveals how the doubling of *o*'s acceleration is an instance of a more general pattern.

With respect to understanding-why, it also seems plausible to say that a person could have at least some understanding of why *o*'s acceleration has doubled by understanding that *D*. This accords with UWHY. For, even though *D* makes no reference to causes, understanding that *D* is so is nonetheless an instance of having a causal guiding explanation of why *o*'s acceleration has doubled. Namely, there is some set of propositions Φ, such that the truth of *D* (but not the truth of the proposition <*o*'s acceleration has doubled>) guarantees that some member of Φ is at least partially causally responsible for the fact that *o*'s acceleration has doubled. For instance, Φ could be the set of all propositions of the form, <The net force acting on *o* has increased from *y* to *z*> (where the value of *z* is twice that of *y*). This reveals that explanatory propositions that might be taken to exemplify some type of non-causal explanation can still be sufficient to provide one with a causal guiding explanation, and thus, according to UWHY, may be sufficient to provide one with understanding-why.

4. Increases in understanding-why

How might the above account of partial understanding-why be extended to account for the *gradability* of understanding-why? While it is widely acknowledged that understanding-why is gradable (Baumberger 2011; Strevens 2013; Grimm forthcoming), surprisingly little has been said about how a person's degree of understanding-why might be determined. In this section, I focus on *increases* in understanding-why, and I provide a rough sketch of how such increases can be spelled out in terms of *causal guiding explanations.*

Suppose that Mick's friends are tampering with his car, and they accidentally drop a box of screws in the engine, one of which falls in the distributor hole. As they are attempting to retrieve the screw, they unintentionally jab part of Mick's engine with a sharp tool. Consequently, Mick's car will not start. Given this, how might one move from having at least some understanding of why Mick's car will not start to having an even better understanding of why it will not start?

First, consider an analogy. Suppose that King Albert sends Jeremiah to a nearby island in order to locate several treasure chests that have been buried there; though neither King Albert nor Jeremiah knows exactly how many chests there are. Moreover, suppose that unbeknownst to Jeremiah, there are several maps located around the island that indicate, to varying degrees of specificity, the areas in which the chests are buried. Given this, consider how Jeremiah might *progress in his search for the chests*. As an initial step, Jeremiah might discover one of the maps (call it "*m1*") that indicates some general area in which one of the chests (call it "*t1*") is buried; e.g., *m1* might provide a set of coordinates that together indicate some general location in which *t1* is buried, or *m1* might reveal some landmark (e.g., the tallest tree on the island) that indicates some general area in which *t1* is buried. Then, there are at least two ways in which Jeremiah might further progress in his search for the chests. First, Jeremiah might discover another map, *m1**, that provides a more precise indication (relative to that provided by *m1*) of where *t1* is buried. Second, Jeremiah might discover some map, *m2*, that indicates the relative positions between two general areas: i.e., a general area in which *t1* is buried and a general area in which one of the other chests (call it "*t2*") is buried.

If either of these two possibilities were to obtain, then Jeremiah could progress still further in his search for the chests in similar ways. For instance, suppose that after discovering *m1*, Jeremiah discovers *m2*. Then, Jeremiah could further progress in his search for the chests by either: discovering some map, *m1**, that provides a more precise indication (relative to that provided by *m1*) of where *t1* is buried; discovering some map, *m2**, that provides a more precise indication (relative to that provided by *m2*) of the relative positions between *t1* and *t2*; or discovering another map, *m3*, that indicates the relative positions between three general areas: i.e., a general area in which *t1* is buried, a general area in which *t2* is buried, and a general area in which one of the other chests (call it "*t3*") is buried. These potential *ways of progressing in one's search for the chests* can be described more generally as follows (call the proposal "PRO-CHST"):

PRO-CHST For any subject *S*, *S* progresses in her search for the chests if:

(g1) *S* obtains a map that better indicates, relative to every other map that *S* already has, the area in which one of the chests is buried.

With this in mind, consider how a similar account might be given with respect to *ways of progressing in one's understanding of why P*, as follows (call the account "PRO-UND"):

PRO-UND For any subject *S*, *S* progresses in her understanding of why *P* if:

(h1) *S* obtains a causal guiding explanation that better indicates, relative to every other causal guiding explanation that *S* already has, some part of the full causal nexus that is responsible for the fact that *P*.

To return to the example with Mick, let *NOSTART* be the proposition <Mick's car will not start>, and suppose that Mick moves from having no understanding of why *NOSTART* to having at least some understanding of why *NOSTART* by understanding that his car will not start because some part of the engine is broken. Then, according to PRO-UND, Mick progresses in his understanding of why *NOSTART,* since Mick moves from having no causal guiding explanation of why *NOSTART* to having at least some causal guiding explanation of why *NOSTART* (and thus, trivially, Mick obtains a causal guiding explanation that better indicates, relative to every other causal guiding explanation that he has, some part of the full causal nexus that is responsible for the fact that *NOSTART*).

How else might a causal guiding explanation *better indicate*, relative to every other causal guiding explanation that a person has, some part of the full causal nexus that is responsible for the fact that *P*? While a complete answer to this question would require more space than is available here, it is nonetheless instructive to consider some of the general features that might be involved in such an answer. Accordingly, I offer the following proposal (call the proposal "INDICATE"):

INDICATE For any proposition *P*, and any two causal guiding explanations *C* and *C** (as to why *P*), *C** better indicates (relative to *C*) some part of the full causal nexus that is responsible for the fact that *P* if either:

(i1) *C** has greater *explanatory precision* than *C*;

(i2) *C** has greater *explanatory length* than *C*;

(i3) *C** has greater *explanatory density* than *C*; or

(i4) *C** has greater *explanatory breadth* than *C*.

Thus, consider exactly what it is for a given causal explanation to differ from another causal guiding explanation in terms of *explanatory precision*, *explanatory length*, *explanatory density*, and *explanatory breadth*.

First, consider the notion of *explanatory precision* by comparing the following two propositions:

C1. The car will not start because some part of the engine is broken.

C2. The car will not start because one of the parts of the engine that is broken is the distributor rotor.

When comparing *C1* and *C2*, there seems to be a sense in which *C2* is more *precise* than *C1*. But exactly how can we spell out this notion of explanatory precision with respect to causal guiding explanations? To answer this, consider first some new terminology. Recall that a proposition *C* is a *causal guiding explanation* of why *P* only if there is some set of propositions Φ, such that the truth of *C* (but not the truth of *P*) guarantees that some member of Φ is at least partially causally responsible for the fact that *P*. With this in mind, call any such set Φ "*a causal map that is given by C with respect to P*." And for any *M* that is a member of some such set Φ and is at least partially causally responsible for *P*, call *M* "*a causal element that is carried by C with respect to P*." Given this, we can now define the notion of *having greater explanatory precision than*, as it pertains to causal guiding explanations, as follows (call the definition "PRECISE"):

PRECISE For any proposition *P*, and any two causal guiding explanations *C* and *C** (as to why *P*), *C** has greater explanatory precision than *C* if and only if:

(j1) every causal map that is given by *C* (with respect to *P*) is also given by *C**, and some causal map that is given by *C** (with respect to *P*) is a proper subset of some causal map that is given by *C*.

According to PRECISE, *C2* has greater explanatory precision than *C1*. To see why, consider the following two causal maps that are given by *C1* with respect to *NOSTART*:

Φ_a the set containing every proposition of the form, <The *engine* bears property *x*>, where *x* is a potential realizer of the engine's brokenness;

Φ_b the set containing every proposition of the form, <Some part of the *car* bears property *x*>, where *x* is a potential realizer of the engine's brokenness.

Both Φ_a and Φ_b are causal maps that are given by *C1* with respect to *NOSTART*, since for each set, the truth of *C1* (but not the truth of *NOSTART*) guarantees that there is at least some member of the set that is at least partially causally responsible for *NOSTART*. Moreover, since the distributor rotor is part of the engine, and the engine is part of the car, both Φ_a and Φ_b are also causal maps that are given by *C2* with respect to *NOSTART*; that is, for each set, Φ_a and Φ_b, the truth of *C2* (but not the truth of *NOSTART*) guarantees that there is at least some member of the set that is at least partially causally responsible for *NOSTART*. A similar argument can be made for any causal map that is given by *C1* with respect to *NOSTART*. Thus, in accordance with the first conjunct of (j1) from PRECISE, every causal map that is given by *C1* with respect to *NOSTART* is also a causal map that is given by *C2*. Next, consider the second conjunct of (j1), which states, "some causal map that is given by *C** (with respect to *P*) is a proper subset of some causal map that is given by *C*." In accordance with the second conjunct of (j1), the set Φ_a is a causal map that is given by *C1* with respect to *NOSTART*, and there is some causal map that is given by *C2* that is a proper subset of Φ_a. For instance, consider the following set:

Φ_a' the set containing every proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the *distributor rotor's* brokenness.

Since the car will not start because the distributor rotor is broken, and the distributor rotor is part of the engine, Φ_a' is a causal map that is given by *C2* with respect to *NOSTART*, and Φ_a' a proper subset of Φ_a. Hence, according to PRECISE, *C2* has greater *explanatory precision* than *C1*, and thus, according to PRO-UND and INDICATE, a move from merely understanding that *C1* to also understanding that *C2* is one in which a person progresses in understanding why *NOSTART* (that is, it is one in which a person obtains an *increase* in understanding why *NOSTART*).

Second, consider how a given causal guiding explanation might differ from another causal guiding explanation in terms of *explanatory length*. Compare the following two propositions:

C1. The car will not start because some part of the engine is broken.

C3. The car will not start because some part of the engine is broken, due to the fact that some part of the engine was jabbed with a sharp tool.

When comparing *C1* and *C3*, there seems to be a sense in which *C3* has greater *explanatory length* than *C1*. In particular, while each of *C1* and *C3* provides a partial explanation of why *NOSTART* via some explanans *M*, *C3* also provides a partial explanation of why *M* via some explanans *M**. Thus, relative to *C1*, *C3* represents a *longer* segment of some explanatory chain that leads to the fact that *NOSTART*. Accordingly, in terms of causal guiding explanations, we can define the notion of *having greater explanatory length than*, as follows (call the definition “LENGTH”):

LENGTH For any proposition *P*, and any two causal guiding explanations *C* and *C** (as to why *P*), *C** has greater explanatory length than *C* if and only if:

(k1) the truth of *C** (but not the truth of *C*) guarantees that there is some causal element *M* that is carried by both *C* and *C** (with respect to *P*), and some causal element *M** that is carried by *C** (with respect to *P*), such that *M** is at least partially causally responsible for *M*.

To see why (k1) is satisfied with respect to *C1* (as *C*), *C3* (as *C**), and *NOSTART* (as *P*), consider again the following causal map that is given by *C1* with respect to *NOSTART*:

Φ_a the set containing every proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the engine's brokenness.

Since Φ_a is causal map that is given by *C1* with respect to *NOSTART*, Φ_a contains (among other things) some member *M* that is at least partially causally responsible for the fact that *NOSTART*; and thus, *M* is a causal element that is carried by *C1* with respect to *NOSTART*. Moreover, since any causal map that is given by *C1* is given by *C3* (since *C3* entails *C1*), it follows that *M* is also a casual element that is carried by *C3* with respect to *NOSTART*.

With this in mind, consider next a causal map that is given by *C3* but not by *C1* (call the set "Φ_c"):

Φ_c the set containing every proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the engine's being jabbed with a sharp tool.

Given the two aforementioned sets, i.e., Φ_a and Φ_c, it can be seen that the truth of *C3* (but not the truth of *C1*) guarantees that there is some member *M* of Φ_a (where *M* is a causal element that is carried by both *C1* and *C3* with respect to *NOSTART*), and some member *M** of Φ_c (where *M** is a causal element that is carried by *C3* with respect to *NOSTART*), such that *M** is at least partially causally responsible for *M*. To illustrate, recall that *C3* is the proposition, <The car will not start because some part of the engine is broken, due to the fact that some part of the engine was jabbed with a sharp tool>; and *C1* is the proposition, <The car will not start because some part of the engine is broken>. Thus, the truth of *C3*, unlike the truth of *C1*, guarantees that there is some member of Φ_c (i.e., some proposition of the form, <The engine bears property *y*>, where *y* is a potential realizer of the engine's being jabbed with a sharp tool) that is at least partially causally responsible for both the fact that *NOSTART* and the fact that *M* (where *M* is a member of Φ_a and is thus a proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the engine's brokenness). So, in sum, condition (k1) of LENGTH is satisfied with respect to *C1* (as *C*), *C3* (as *C**), and *NOSTART* (as *P*). Thus, *C3* has greater *explanatory length* than *C1*, and hence, according to PRO-UND and INDICATE, a move from merely understanding that *C1* to also understanding that *C3* is one in which a person progresses in understanding why *NOSTART* (or, in other words, it is one in which a person obtains an *increase* in understanding why *NOSTART*).

Next, consider how a given causal guiding explanations might differ from another causal guiding explanation in terms of *explanatory density*. Compare the following two propositions:

> *C0*. The car will not start because some part of the engine was jabbed with a sharp tool.
>
> *C3*. The car will not start because some part of the engine is broken, due to the fact that some part of the engine was jabbed with a sharp tool.

In contrast to the relation between *C1* and *C3* (where *C3* has greater *explanatory length* than that had by *C1*), the relation between *C0* and *C3* can be better described as one in which *C3* has greater *explanatory density* than *C0*; i.e., in that while each of *C0* and *C3* provides a partial explanation of why *NOSTART* via some explanans *M*, *C3* further provides an explanatory intermediate (call it "*M**") between the fact that *NOSTART* and the fact that *M*. That is, although *C0* and *C3* represent the same segment of some explanatory chain that leads to the fact that *NOSTART*, *C3* provides a *denser* (i.e., more populated) representation of this particular segment of the explanatory chain. To express the idea in terms of causal guiding explanations, we can define the notion of *having greater explanatory density than*, as follows (call the definition "DENSITY"):

> **DENSITY** For any proposition *P*, and any two causal guiding explanations *C* and *C** (as to why *P*), *C** has greater *explanatory density* than *C* if and only if:
>
> (m1) the truth of *C** (but not the truth of *C*) guarantees that there is some causal element *M* that is carried by both *C* and *C** (with respect to *P*), and some causal element

M^* that is carried by C^* (with respect to P), such that M is at least partially causally responsible for M^*.

Consider how (m1) is satisfied with respect to *C0* (as *C*), *C3* (as *C**), and *NOSTART* (as *P*) by considering the following causal map that is given by *C0* with respect to *NOSTART*:

Φ_c the set containing every proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the engine's being jabbed with a sharp tool.

Since Φ_c is a causal map that is given by *C0* with respect to *NOSTART*, Φ_c contains some member *M* that is at least partially causally responsible for the fact that *NOSTART* (and thus *M* is a causal element that is carried by *C0* with respect to *NOSTART*). Furthermore, given that every causal map that is given by *C0* is a causal map that is given by *C3* (since *C3* entails *C0*), *M* is also a casual element that is carried by *C3* with respect to *NOSTART*.

In light of this, consider next a causal map that is given by *C3* but not by *C0*:

Φ_a the set containing every proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the engine's brokenness.

How do Φ_c and Φ_a relate to condition (k1) of DENSITY? Given that *C3* is the proposition, <The car will not start because some part of the engine is broken, due to the fact that some part of the engine was jabbed with a sharp tool>, the truth of *C3*, like the truth of *C0*, guarantees that there is at least some member *M* of Φ_c (i.e., some proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the engine's being jabbed with a sharp tool), such that *M* is at least partially causally responsible for the fact that *NOSTART* (and thus, *M* is a causal element that is carried by both *C0* and *C3* with respect to *NOSTART*) . However, unlike the truth of *C0*, *C3* further guarantees that there is some member *M** of Φ_a (i.e., some proposition of the form, <The engine bears property *x*>, where *x* is a potential realizer of the engine's brokenness), such that *M** is a causal intermediate between the fact that *NOSTART* and the fact that *M*. Or, in other words, the truth of *C3* (but not the truth of *C0*) guarantees that for some proposition *M* (where *M* is a causal element that is carried by both *C0* and *C3* with respect to *NOSTART*), and for some proposition *M** (where *M** is a causal element that is carried by *C3* with respect to *NOSTART*), *M* is at least partially causally responsible for *M**. Thus, according to DENSITY, *C3* has greater *explanatory density* than *C0*, and hence, according to PRO-UND and INDICATE, a move from merely understanding that *C0* to also understanding that *C3* is one in which a person progresses in understanding why *NOSTART* (i.e., it is one in which a person obtains an *increase* in understanding why *NOSTART*).

Finally, consider how two causal guiding explanations might differ from each other in terms of *explanatory breadth* by comparing the following two propositions:

C2. The car will not start because one of the parts of the engine that is broken is the distributor rotor.

C4. The car will not start because both the distributor rotor and the backup distributor rotor are broken.

In contrast to the above cases regarding *explanatory precision*, *explanatory length*, and *explanatory density*, the relation between *C2* and *C4* might be best described as one in which *C4* has greater *explanatory breadth* than that had by *C2*; i.e., in that while *C2* and

C4 each provides a partial explanation of why *NOSTART* via some explanans *M*, *C4* provides a further explanation of why *NOSTART* via both *M* and some additional explanans *M** (where the fact that *M** is so neither partially explains nor is partially explained by the fact that *M*). For instance, we can suppose that while neither *the fact that the distributor rotor is broken* nor *the fact that the backup distributor rotor is broken* is sufficient to result in the fact that *NOSTART*, *the fact that both the distributor rotor and the backup distributor rotor are broken* is sufficient to result in the fact that *NOSTART*. Given such a case, a move from merely understanding that *C2* to also understanding that *C4* seems to be a move in which a person obtains a causal guiding explanation that better indicates (relative to any other causal guiding explanation that the person has) some part of the full causal nexus that is responsible for the fact that *NOSTART*. Accordingly, we can define the notion of *having greater explanatory breadth than* in terms of causal guiding explanations, as follows (call the definition "BREADTH"):

BREADTH For any proposition *P*, and any two causal guiding explanations *C* and *C** (as to why *P*), *C** has greater *explanatory breadth* than *C* if and only if:

(n1) the truth of *C** (but not the truth of *C*) guarantees that there is some causal element *M* that is carried by *C* (with respect to *P*), and some causal element *M** that is carried by *C** (with respect to *P*), such that *M* and *M** are jointly causally responsible for *P* (even if *M* and *M** are causally independent of each other).

Given this, *C4* has greater *explanatory breadth* than *C2*, and thus, according to PRO-UND and INDICATE, a move from merely understanding that *C2* to also understanding that *C4* is one in which a person progresses in understanding why *NOSTART* (that is, it is one in which a person obtains an *increase* in understanding why *NOSTART*).

Thus, in sum, there are at least five ways by which a person *S* might obtain an increase in understanding why *P*, namely:

(w1) *S* might move from having no causal guiding explanation of why *P* to having at least some causal guiding explanation of why *P*; or

(w2 – w5) *S* might move from having at least some causal guiding explanation *C* (of why *P*) to having a further causal guiding explanation *C** (of why *P*), where *C** has either

- greater *explanatory precision*,
- greater *explanatory length*,
- greater *explanatory density*, or
- greater *explanatory breadth*

than every other causal guiding explanation that *S* has as to why *P*.

While the above notions regarding causal guiding explanations provide a preliminary framework with which to capture the notion of increases in understanding-why, there is still much to be said. For instance, of particular importance, the question arises: Are the five aforementioned ways by which a person might obtain an increase in understanding why *P*, i.e., (w1) – (w5), sufficient to eventually bring a person to a *full understanding* of why *P*? And, if they are not sufficient, what further ways of obtaining an increase in understanding why *P* might be added to (w1) – (w5) to make the ways sufficient to eventually bring a person to a *full understanding* of why *P*? I plan to address such questions in a subsequent paper. In particular, I plan to provide an account of what it is for something to be *the full causal nexus* that is responsible for a given fact. Then, I plan to argue that for any person *S*, and any proposition *P*, *S has a full understanding of why P* if and only if for every proposition *Q* (where *Q* figures into the full causal nexus that is responsible for the fact that *P*), *S* understands that *Q*. Finally, I plan to show how any continued sequence of increases in understanding-why, as spelled out above in terms of causal guiding explanations, is guaranteed to eventually lead to full understanding-why.

5. Conclusion

To conclude, I began this chapter by discussing the relationship between dispositions, multiple realizability, causation, and explanation (Section 2). Then, in Section 3, I introduced the notion of *causal guiding explanation*, and I argued that a person *S* has at least some understanding of why *P* (where *P* is a proposition that represents an event) if and only if *S* has a causal guiding explanation of why *P*. Lastly, in Section 4, I provided a sketch of how one might extend the proposed account of what it is to have at least some understanding-why to further account for the gradability of understanding-why.

This chapter thus contributes to the literature in at least three ways. First, it offers a novel account of understanding-why and its relation to explanation. Second, it suggests a new way to think about the causal relevance of higher-level properties (i.e., by showing how higher-level properties can be relevant to obtaining an objective causal explanation of why *P* even if higher-level properties are not themselves causally efficacious). Third, it provides a provisional framework with which to start thinking further about gradations of understanding-why (something that, to my knowledge, has not yet been offered in the literature).

This chapter also raises a number of questions. How might the provisional framework offered for gradations of understanding-why be more fully developed? How does understanding-why (as it pertains to events) relate to other forms of understanding-why; e.g., understanding why a given act is morally wrong, or understanding why every triangle has an area that is necessarily greater than the area of its incircle? How does the notion of understanding-why relate to other types of understanding; e.g., understanding quantum mechanics, or understanding the movements of planets? I plan to address such questions in subsequent work.

CHAPTER 3

An Account of Logical Grounding and its Relevance to Understanding-Why

Overview

It is widely held that the notion of *grounding* is primitive; i.e., unanalyzable in terms of other independently intelligible notions. In this chapter, I argue that at least one prominent type of grounding (i.e., *logical grounding*) can indeed be analyzed in terms of other independently intelligible notions. I begin by considering various types of grounding in order to situate the relation of logical grounding among these various types. Then, I examine the relationship between the notion of logical grounding and a number of other notions that I construct solely on the basis of concepts that are commonly employed in logic and set theory. With these notions in hand, I proceed to develop an account of logical grounding as it pertains to logically contingent facts. Then, I extend this account so that it also applies to logically necessary facts, and I examine a few final instances of logical grounding and show how the final account applies to such cases. Finally, I considered how the account offered here pertains to the notion of understanding-why.

Introduction

What is it for a given fact to be *grounded* in some set of facts? The grounding relation is often associated with phrases such as, "metaphysical priority," "non-causal determination," and "non-causal dependence". Despite the technical language used in such phrases, many authors hold that the notions at play here (i.e., metaphysical priority, non-causal determination, and non-causal dependence) constitute a single notion that is both *quotidian* (i.e., commonly used in ordinary, day-to-day thinking) and *primitive* (i.e., unanalyzable in terms of other independently intelligible notions). Regarding its quotidian nature, for instance, Berker (2015) writes:

> "[G]rounds"-talk lends itself to the surprisingly common misconception that "grounds" is a technical term referring to a wholly new relation that was invented by Fine in 2001. But nothing could be further from the truth,

> and the intended equivalence of "grounds"-talk to "in virtue of"- and "because"-talk helps us see that. Although the use of the word "grounds" in roughly Fine's sense is a relatively recent phenomenon, "in virtue of"- and (non-causal) "because"-talk have been with us from the very beginning, [at least as far back as Plato].

To illustrate, in Plato's *Euthyphro*, Socrates argues for the position that an act is loved by the gods *in virtue of* the fact that it is pious, while arguing against the position that an act is pious in virtue of the fact that it is loved by the gods. The in-virtue-of relation here seems to be one of non-causal dependence; i.e., one of *grounding*. Accordingly, Socrates' position might be reworded as follows: The fact that an act is pious *grounds* the fact that the act is loved by the gods, but not vice versa. To use another example, consider the view that certain counterfactual facts (e.g., the fact that if the vase were struck, then it would break) hold *in virtue of* certain dispositional facts (e.g., the fact that the vase is fragile). Here, again, the in-virtue-of relation seems to be one of grounding; i.e., certain counterfactual facts *non-causally depend* on certain dispositional facts. Lastly, consider an example from logic. It is commonly held that the truth of a conjunction holds *in virtue of* the truth of its conjuncts. Or, in other words, the truth of a conjunction is *grounded* in the truth of its conjuncts. For example, the fact that 5 is both congruent and prime is grounded in the set {the fact that 5 is congruent, the fact that 5 is prime}.

In light of such examples, one might wonder whether an *analysis* of the notion of grounding can be given. While there is still some debate over this issue, it is predominantly held that grounding is *primitive*, i.e., unanalyzable in terms of other independently intelligible notions (Schaffer 2009; Rosen 2010; Audi 2012).[20] This view is not only held by several grounding theorists, but also by some grounding skeptics. For instance, when arguing against the intelligibility of the term "grounding," Chris Daly (2012, p. 91) writes:

> One way to understand a term is to trace its analytic connections with terms that are already understood. The difficulties facing grounding theorists' attempts to use this measure is that *either* the terms appealed to are such close cognates to "grounding" that they are as obscure as it is, *or* the terms appealed to have are sufficiently and independently clear but their connection to "grounding" is questionable.

While Daly appeals to the claim that the notion of grounding is indefinable in attempt to cast doubt on the intelligibility of the term "grounding", many grounding theorists appeal to the same claim to argue that the notion of grounding is primitive. In contrast to both of these approaches, I argue that at least one of the most prominent types of grounding (i.e., *logical grounding*) can indeed be defined in terms of other independently intelligible notions (i.e., notions that are commonly employed in logic and set theory).

In Section 1, I begin by considering various types of grounding in order to situate the relation of logical grounding among these various types. In Section 2, I examine the

[20] This is not to deny, however, that the notion can be illuminated in various ways. For instance, along with providing specific examples of grounding (like Socrates' views on piety, the conterfactual-dispositional dependence view, and the relation between the truth of a conjunction and the truth of its conjuncts), one might also indicate certain properties that hold of grounding, such as being hyperintensional (i.e., being such that F might be grounded in P but not in Q, even if P is logically equivalent to Q), being factive, being transitive, being irreflexive, and being distinct from mere entailment (Schaffer 2009; Rosen 2010; Audi 2012; Correia & Schnieder 2012).

relationship between *logical grounding* and *minimal entailment*, and I introduce a notion that I call "*sentential representation*", which serves to connect *sets of facts* to *sets of sentences representing those facts*. In Section 3, I consider how certain sets of sentences are *more basic* than others, and I proceed in Section 4 to explore how this idea might be employed in constructing an account of logical grounding with respect to logically contingent facts. Then, in Section 5, I extend the account that is presented in Section 4 so that it also applies to logically necessary facts, and I examine a few final instances of logical grounding and show how the final account applies to such cases. Finally, in Section 6, I considered how the account pertains to the notion of understanding-why.

1. Various types of grounding

With the notion of grounding in mind, consider two ways in which some set of facts Φ might ground a further fact *F*. First, Φ might *fully ground F*. To illustrate, let *PRIMATE* be the fact that both monkeys and apes are primates. Then, the following set *fully grounds PRIMATE*:

$$\Gamma = \{\text{the fact that monkeys are primates, the fact that apes are primates}\}.$$

Second, a set of facts might *merely partially ground* some fact; e.g., in the way that the following set *merely partially grounds PRIMATE*:

$$\Gamma' = \{\text{the fact that monkeys are primates}\}.$$

How might we further characterize this difference between the way in which Γ grounds *PRIMATE* and the way in which Γ′ grounds *PRIMATE*? One possibility is to appeal to the claim that there is a certain type of necessity that is carried by the notion of full grounding but not by the notion of mere partial grounding. Many, for instance, endorse something like the following principle (call it "NECESS"):[21]

NECESS For any fact F, and any set of facts Φ, Φ *fully grounds F* only if as a matter of metaphysical necessity, F obtains if every member of Φ obtains.

As some authors note, however, one might construct different variations of NECESS that correspond to different types of grounding (Correia 2005, 2008; Correia and Schnieder 2012; Fine 2012). For instance, it might be thought that NECESS corresponds to the notion of *full metaphysical grounding*, while the principle NECESS-C (which replaces the phrase "metaphysical necessity" with "conceptual necessity") corresponds to the notion of *full conceptual grounding*, and NECESS-L (which replaces the phrase "metaphysical necessity" with "logical necessity") corresponds to the notion of *full logical grounding*. On this view, all instances of full logical grounding are instances of full conceptual grounding, and all instances of full conceptual grounding are instances of full metaphysical grounding. But the reverse does not hold. That is, there are instances of full metaphysical grounding that are not instances of full conceptual grounding, and there are instances of full conceptual grounding that are not instances of full logical grounding. To illustrate, let *SIXEVN* be the fact that 6 is either an even number or a prime number. Then, the following set of facts fully conceptually grounds, but does not fully logically ground, *SIXEVN*:

Δ = {the fact that 6 is an integer that is divisible by 2}.

[21] E.g., see Correia 2005; deRosset 2010; Rosen 2010; Audi 2012; Correia and Schnieder 2012; Trogdon 2013; Dasgupta 2014.

Now contrast Δ with a set of facts that fully logically grounds (and thus also fully conceptually grounds) *SIXEVN*; namely, the following set:

$\Delta' = \{$the fact that 6 is an even number$\}$.

Given the distinction between full logical grounding, full conceptual grounding, and full metaphysical grounding, how might we understand the corresponding notions of *mere partial logical grounding*, *mere partial conceptual grounding*, and *mere partial metaphysical grounding*? To answer this, consider first the notion of mere partial logical grounding. One way to characterize this notion is to take the notion of full logical grounding as the more basic notion and then define the notion of mere partial logical grounding in terms of full logical grounding, as follows (call the definition "PARTIAL"):[22]

PARTIAL For any fact F, and any set of facts Φ, Φ *merely partially logically grounds F* if and only if Φ does not fully logically ground F, but some proper superset of Φ fully logically grounds F.

A similar definition can be given for the notion of *mere partial conceptual grounding* by replacing all occurrences of the term "logically" in PARTIAL with the term "conceptually"; and similarly, a definition for the notion of *mere partial metaphysical grounding* can be given by replacing all occurrences of the term "logically" in PARTIAL with the term "metaphysically".

[22] One reason to think that the notion of full grounding cannot be defined in terms of partial grounding is provided by Kit Fine (2012). Fine notes that the partial grounds are the same for both (Q & R) and ($Q \vee R$), but the full grounds for the two facts differ from each other. Thus, Fine (2012) asks, "[H]ow are we to distinguish between the full grounds [for the two facts] if appeal is only made to their partial grounds?" (p. 50).

Thus, there are at least two dimensions along which the grounding relation can be divided. The first dimension pertains to whether a particular instance of grounding is an instance of *partial grounding* or *full grounding*, and the second dimension pertains to the type of necessity (e.g., logical, conceptual, metaphysical) that is carried by the notion of full grounding. As noted in the Introduction, the primary focus of this chapter is on the notion of *logical grounding*, and a particular emphasis is placed on the notion of *full logical grounding* (since the notion of *mere partial logical grounding* can be defined in terms of full logical grounding, as is done above in PARTIAL). Therefore, let us now turn to consider the notion of full logical grounding in more detail.

2. Logical grounding, minimal entailment, and sentential representations

Recall again the definition of *mere partial grounding* that is provided above in PARTIAL. It states that a set of facts Φ *merely partially logically grounds* some fact *F* if and only if both: (i) Φ does not fully logically ground *F*, and (ii) some proper superset of Φ fully logically grounds *F*. This portrayal of the relation between partial and full logical grounding raises an important question. Namely, are there cases in which a proper superset of some full logical ground for a fact is itself a full logical ground for the fact? It seems that there are indeed such cases. For instance, let *ORBITS* be the fact that either Venus or Mars orbits the Sun. Then, it seems that *ORBITS* is fully logically grounded in each of the following three sets:

N = {the fact that Venus orbits the Sun}

M = {the fact that Mars orbits the Sun}

Ω = {the fact that Venus orbits the Sun, the fact that Mars orbits the Sun}.

Given that Ω is a proper superset of N (as well as a proper superset of M), this seems to be a case in which a proper superset of some full logical ground for a fact is itself a full logical ground for the fact.

However, this is not to say that *every* proper superset of a full logical ground for a fact is itself a full logical ground for the fact. As many have noted, it seems that grounding is *non-monotonic*; i.e., that there are cases in which a proper superset of some ground for a fact is *not* itself a ground for the fact (e.g., Rosen 2010; Audi 2012; Raven 2012). To see how this applies to the notion of full logical grounding, let *PRIME* be the fact that either 7 or 12 is a prime number. Then, consider the relationship between *PRIME* and the following two sets:

Π = {the fact that 7 is a prime number}

Π* = {the fact that 7 is a prime number, the fact that some people enjoy the sound of thunder, the fact that Madison is the capital of Wisconsin}.

While it seems plausible to say that Π fully logically grounds *PRIME*, it seems implausible to say that Π* fully logically grounds *PRIME*. Why is this? The reason, I submit, is that the notion of grounding is closely tied to the notion of *relevance*. To illustrate, while the obtainment of every fact in Π is relevant to the obtainment of *PRIME*, it is not the case that the obtainment of every fact in Π* is relevant to the obtainment of *PRIME*. Thus, any account of full logical grounding should have the resources to differentiate sets containing irrelevancies with respect to the grounded fact from sets without such irrelevancies.

Secondly, any account of grounding also faces a particular complication regarding the individuation of facts. For instance, consider the following set:

Π' = {the fact that both 6 is an even number and Jupiter is identical to itself}.

Does Π' fully logically ground *PRIME*? The answer seems to depend on one's view regarding the individuation of facts. Given a particularly coarse-grained view of facts (i.e., one on which two facts are identical if and only if they are logically equivalent), it follows that Π is identical Π'; and thus, since Π fully logically grounds *PRIME*, it must be the case that Π' fully logically grounds *PRIME*. However, on a more fine-grained view of facts (i.e., one on which Π is distinct from Π'), one would likely want to say that Π, but not Π', fully logically grounds *PRIME* (since the obtainment of the fact that Jupiter is self-identical seems irrelevant to the obtainment of *PRIME*).

With the aforementioned issues in mind, I aim to provide an account of full logical grounding (henceforth, just "full grounding") that does the following:
(i) defines the notion of full grounding in terms of other independently intelligible notions (i.e., notions that are commonly employed in logic and set theory), (ii) avoids problems pertaining to cases in which the proposed ground for some fact contains information that is irrelevant to the grounded fact, and (iii) remains neutral on how facts are individuated.

To do this, I begin by defining the notion of *minimal entailment* as a relation that holds between a *sentence* and a *set of sentences*, as follows:

MINENT For any sentence T, and any set of sentences Σ, Σ *minimally entails* T if and only if both Σ entails T, and no proper subset of Σ entails T.

To illustrate, consider the following two claims:

(A1) the set {"Venus orbits the Sun", "Mars orbits the Sun"} minimally entails the sentence, "Both Venus and Mars orbit the Sun";

(A2) the set {"Venus orbits the Sun", "3 is a prime number"} does *not* minimally entail the sentence, "Venus orbits the Sun".

With respect to (A1), the set {"Venus orbits the Sun", "Mars orbits the Sun"} minimally entails the sentence, "Both Venus and Mars orbit the Sun", since the set entails the sentence, and no proper subset of the set entails the sentence.

With respect to (A2), by contrast, the set {"Venus orbits the Sun", "3 is a prime number"} does not minimally entail the sentence, "Venus orbits the Sun", since there is a proper subset of the set that entails the sentence; i.e., the set {"Venus orbits the Sun"}.

Given that the relation of full grounding is generally taken to hold between a fact and set of facts, how might we use the notion of *minimal entailment* (which is defined in terms of sentences) to define full grounding? We can do so, I submit, by introducing a two-place relation that connects *sets of facts* to *sets of sentences that represent those facts*. Call this two-place relation, "the relation of being a *sentential representation* of," and let it be defined as follows:

SENTR For any set of facts Φ, and any set of sentences Σ, Σ is a sentential representation of Φ if and only if both: (a) for each fact in Φ, there is exactly one sentence in Σ that represents that fact; and (b) for each sentence in Σ, there is exactly one fact in Φ that the sentence represents.

To illustrate, consider the set {the fact that either Venus or Venus orbits the Sun, the fact that both Jupiter and Mars orbit the Sun}. The members of this set can be represented in a variety of ways. For instance, suppose that:

(i) *"VN"* is a sentence representing the fact that Venus orbits the Sun,

(ii) *"JP"* is a sentence representing the fact that Jupiter orbits the Sun,

(iii) *"MR"* is a sentence representing the fact that Mars orbits the Sun,

(iv) *"JAM"* is a sentence representing the fact that both Jupiter and Mars orbit the Sun,

(v) *"MAJ"* is a sentence representing the fact that both Mars and Jupiter orbit the Sun, and

(vi) *"VOV"* is a sentence representing the fact that either Venus or Venus orbits the Sun.

Then, regardless of how fine-grained one's view is on the individuation of facts, it follows from SENTR that each of the following four sets is a sentential representation of the set {the fact that either Venus or Venus orbits the Sun, the fact that both Jupiter and Mars orbit the Sun}:

Γ_1 = {"Either *VN* or *VN*", "Both *JP* and *MR*"}

Γ_2 = {"*VOV*", "Both *JP* and *MR*"}

Γ_3 = {"Either *VN* or *VN*", *"JAM"*}

Γ_4 = {"*VOV*", *"JAM"*}.

Next, consider how one's view on the individuation of facts might commit one to further claims about which sets count as sentential representations of the set {the fact that either Venus or Venus orbits the Sun, the fact that both Jupiter and Mars orbit the Sun}. Suppose, for instance, that *the fact that Venus orbits the Sun* is the same fact as *the fact that either Venus or Venus orbits the Sun*. Then, the following two sets are also sentential representations of the set {the fact that either Venus or Venus orbits the Sun, the fact that both Jupiter and Mars orbit the Sun}:

{"*VN*", "Both *JP* and *MR*"}

{"*VN*", "*JAM*"}.

By contrast, if *the fact that Venus orbits the Sun* is not the same fact as *the fact that either Venus or Venus orbits the Sun*, and moreover, we suppose that no sentence represents more than one fact, then neither {"*VN*", "Both *JP* and *MR*"} nor {"*VN*", "*JAM*"} is a sentential representation of the set {the fact that either Venus or Venus orbits the Sun, the fact that both Jupiter and Mars orbit the Sun}.

Similarly, suppose that *the fact that both Jupiter and Mars orbit the Sun* is the same fact as *the fact that both Mars and Jupiter orbit the Sun*. Then, in addition to $\Gamma_1 - \Gamma_4$, the following two sets are also sentential representations of the set {the fact that either Venus or Venus orbits the Sun, the fact that both Jupiter and Mars orbit the Sun}:

{"*VOV*", "*MAJ*"}

{"*VOV*", "Both *MR* and *JP*"}.

By contrast, if *the fact that both Jupiter and Mars orbit the Sun* is not the same fact as *the fact that both Mars and Jupiter orbit the Sun*, and moreover, we again suppose that no sentence represents more than one fact, then neither {"*VOV*", "*MAJ*"} nor

{"*VOV*", "Both *MR* and *JP*"} is a sentential representation of the set {the fact that either Venus or Venus orbits the Sun, the fact that both Jupiter and Mars orbit the Sun}.

In short, given the assumption that no sentence represents more than one fact, anything that is a sentential representation of Φ on a fine-grained view of facts is also a sentential representation of Φ on a coarse-grained view of facts. However, the reverse does not hold. As shown above, there may be a sentential representation of Φ on a coarse-grained view of facts that is not a sentential representation of Φ on a fine-grained view of facts.

Now, consider further why the notion of minimal entailment is useful in developing an account of full grounding. Suppose that Σ is a sentential representation of some set of facts Φ, and that *T* is a sentence that represents some further fact *F*. In such cases, it is often true that if Σ minimally entails *T*, then Φ fully grounds *F*. For instance, just as the set {"Venus orbits the Sun"} and the set {"Mars orbits the Sun"} each minimally entails the sentence, "Either Venus or Mars orbits the Sun", the set {the fact that Venus orbits the Sun} and the set {the fact that Mars orbits the Sun} each fully grounds the fact that either Venus or Mars orbits the Sun.

Accordingly, the utility of the notion of minimal entailment arises from the manner in which it helps ensure that a given set of sentences (and thus a given set of facts that are represented by the sentences) does not contain certain types of irrelevancies that would undermine the grounding relation. For instance, if some set Σ minimally entails a given sentence *T*, then Σ contains no members that are wholly irrelevant to the truth of *T*. To illustrate, consider the following two sets:

A = {the fact that 7 is a prime number, the fact that some birds are nocturnal}

A′ = {"7 is a prime number", "Some birds are nocturnal"}.

Just as the set A is not a full ground for the fact that either 7 or 12 is a prime number, the set A′ does not minimally entail the sentence, "Either 7 or 12 is a prime number"; for, there is a proper subset of A′ that entails the sentence, "Either 7 or 12 is a prime number".

Nevertheless, despite its utility in developing an account of full grounding, minimal entailment is neither necessary nor sufficient for full grounding. First, to see why it is not necessary, consider the following two sets:

Ω = {the fact that Venus orbits the Sun, the fact that Mars orbits the Sun}

Ω′ = {"Venus orbits the Sun", "Mars orbits the Sun"}.

While the set Ω seems to fully ground the fact that either Venus or Mars orbits the Sun, the set Ω′ does not minimally entail the sentence, "Either Venus or Mars orbits the Sun". For, there is a proper subset of Ω′ that entails the sentence, "Either Venus or Mars orbits the Sun"; e.g., the set {"Venus orbits the Sun"}.

Second, to see why minimal entailment is not *sufficient* for full grounding, consider the relation between *PRIME* (i.e., the fact that either 7 or 12 is a prime number) and the following set:

Σ′ = {the fact that 7 is a prime number if the Earth is a cube, the fact that 7 is a prime number if the Earth is not a cube}.

While it seems that Σ' does not fully ground *PRIME*, there are many sentential representations of Σ' that minimally entail the sentence, “Either 7 or 12 is a prime number”. For instance, consider the following sentential representations of Σ':

Σ'' = {“7 is a prime number if the Earth is a cube”, “7 is a prime number if the Earth is not a cube”}.

The set Σ'' minimally entails the sentence, “Either 7 or 12 is a prime number”, since Σ'' but no proper subset of Σ'' entails the sentence.

Thus, in order to capture the notion of full grounding, we need something more. In particular, we need an account that rules out sets containing *any* irrelevancy with respect to the obtainment of a fact (not just sets containing members that are wholly irrelevant to the obtainment of the fact). For example, we need an account according to which *PRIME* (i.e., the fact that either 7 or 12 is a prime number) is not fully grounded in the set {the fact that 7 is a prime number if the Earth is a cube, the fact that 7 is a prime number if the Earth is not a cube}, even though many sentential representations of this set minimally entail the sentence, “Either 7 or 12 is a prime number". Moreover, we also need an account according to which a fact such as *PRIME* is fully grounded in the set {the fact that 7 is a prime number, the fact that 12 is a prime number}, even though the sentence, “Either 7 or 12 is a prime number", is not minimally entailed by the set {“7 is a prime number”, “12 is a prime number”}.

3. Reflections and basesets

In order to use the notions of *minimal entailment* and *sentential representation* to develop such an account, I first introduce two further notions: i.e., *reflections* and *basesets*. I define the notion of a *reflection* as a relation that holds between two sets of sentences, as follows:

REFLEC For any two sets of sentences, Σ and Σ^*, Σ^* is a *reflection* of Σ if and only if:

(r1) Σ^* but no proper subset of Σ^* entails every member of Σ, and Σ but no proper subset of Σ entails every member of Σ^*.

Consider, for instance, the following seven sets of sentences, each of which is a reflection of the others (where the quotation marks are used to distinguish strings of symbols as sentences rather than facts):

$\Phi_1 = \{\text{“}N\text{”}, \text{“}R\text{”}\}$

$\Phi_2 = \{\text{“}R \rightarrow N\text{”}, \text{“}R\text{”}\}$

$\Phi_3 = \{\text{“}N \leftrightarrow R\text{”}, \text{“}R\text{”}\}$

$\Phi_4 = \{\text{“}N \,\&\, R\text{”}\}$

$\Phi_5 = \{\text{“}Q \rightarrow N\text{”}, \text{“}{\sim}Q \rightarrow N\text{”}, \text{“}R\text{”}\}$

$\Phi_6 = \{\text{“}N \rightarrow (N \,\&\, R)\text{”}, \text{“}N\text{”}\}$

$\Phi_7 = \{\text{“}N \,\&\, N\text{”}, \text{“}R \,\&\, R\text{”}\}$.

All of these sets are reflections of each other. That is, for any two of these sets, each set is a reflection of the other set. Moreover, there seems to be a sense in which Φ_1 is more basic than any of $\Phi_2 - \Phi_7$. And, when considering whether a given relation is an instance of full grounding, it seems that we only need to be concerned with the issue of whether these more basic reflections minimally entail a given sentence (not whether *every* reflection minimally entails the sentence). So, how might we capture this idea of a more basic reflection? To answer this, I now introduce the notion of an *operational value*, which can be applied to both sentences and sets of sentences. With respect to sentences, I offer the following definition (call it "OPSENT"):

OPSENT For any sentence *T*, *the operational value of T* is the total number of operations under which the atomic sentences in *T* fall.

To illustrate, consider the following two sentences:

"$(H \vee F)$ & M"

"$(N$ & $L) \vee \sim(E$ & $\sim K)$".

With respect to the first sentence, there is a total of 5 operations under which the atomic sentences fall. In particular, "*H*" falls under 2 operations (i.e., those corresponding to ∨ and &), "*F*" falls under 2 operations (i.e., those corresponding to ∨ and &), and "*M*" falls under 1 operation (i.e., that corresponding to &), resulting in a total of 5. With respect to the second sentence, there is a total of 11 operations under which the atomic sentences fall. Namely, "*N*" falls under 2 operations (i.e., those corresponding to & and ∨), "*L*" falls under 2 operations (i.e., those corresponding to & and ∨), "*E*" falls under 3 operations (i.e., those corresponding to ∨, ~, and &), and "*K*" falls under 4 operations (i.e., those corresponding to ∨, ~, &, and ~), resulting in a total of 11.

To extend the notion of an operational value to predicate logic, I treat any formula that is within the scope of a universal quantifier as an infinite conjunction, and I treat any formula that is within the scope of an existential quantifier as an infinite disjunction. For instance, "$\forall y(Py \rightarrow Ny)$" is treated as

"$(((Pa \rightarrow Nb)\ \&\ (Pb \rightarrow Nb))\ \&\ (Pc \rightarrow Nc))\ \&\ \ldots$"

And "$\exists$x(Px & Ny)" is treated as

"$(((Pa\ \&\ Nb) \vee (Pb\ \&\ Nb)) \vee (Pc\ \&\ Nc)) \vee \ldots$"

Thus, the operational value of any sentence containing a quantified is infinite.

Next, consider how the notion of an operational value can be applied to *sets of sentences*, as follows:

OPSET For any set of sentences Σ, *the operational value of Σ* is the arithmetical mean of all the operational values for the members of Σ.

Thus, consider the following set of sentences:

$\Theta = \{$"$\sim K \vee H$", "$\sim(N \rightarrow E)$"$\}$.

The operational value of Θ is 3.5, since the operational value of "$(\sim K \vee H)$" is 3, and the operational value of "$\sim(N \rightarrow E)$" is 4, thus resulting in a mean of 3.5.

With respect to predicate logic, consider the following three sets: {"$\forall x(\exists yLxy \rightarrow \sim Tx)$", "$\exists x \exists yLxy$"}, {"$\exists xMx$", "$Ma$", "$\sim Cb \vee Gc$"}, and {"$\sim Cb$", "$Kcb$"}. The operational value of each of the first two sets is infinite (since each set contains members with quantifiers), while the operational value of the third set is 0.5.

Given the notion of a reflection and the notion of an operational value of a set of sentences, I now define the notion of a *baseset* as follows (call the definition "BASE"):

BASE For any set of sentences Σ, Σ is a *baseset* if and only if:

(b1) Σ has a finite operational value; and (b2) there is no reflection Γ of Σ, such that the operational value of Γ is less than the operational value of Σ.

To grasp the idea of a baseset, consider again the following seven sets of sentences, each of which is a reflection of the others:

Φ_1 = {"*N*", "*R*"}

Φ_2 = {"*R* → *N*", "*R*"}

Φ_3 = {"*N* ↔ *R*", "*R*"}

Φ_4 = {"*N* & *R*"}

Φ_5 = {"*Q* → *N*", "~*Q* → *N*", "*R*"}

Φ_6 = {"*N* → (*N* & *R*)", "*N*"}

Φ_7 = {"*N* & *N*", "*R* & *R*"}.

Now, suppose, for the sake of simplicity, that $\Phi_1 - \Phi_7$ constitute all of the reflections of {"*N*", "*R*"}. Then, given BASE, we can show that {"*N*", "*R*"} is the only baseset among $\Phi_1 - \Phi_7$.[23]

First, since no set among $\Phi_1 - \Phi_7$ contains any sentence with a quantifier, we know that the operational value for each set if finite. Thus, to show that {"*N*", "*R*"} is a baseset, we need to show that for each reflection Γ of {"*N*", "*R*"} (i.e., for each of $\Phi_1 - \Phi_7$), the operational value of Γ is *not* less than the operational value of {"*N*", "*R*"}. This is fairly easy to show, since both "*N*" and "*R*" are atomic sentences; thus, there is no set that has an operational value that is less than that of {"*N*", "*R*"}, since the operational value of {"*N*", "*R*"} is 0.

[23] Note that Φ_1's being a baseset depends, in part, on the fact that "*N*", "*R*", and "*Q*" are all distinct atomic sentences. For instance, if "*R*" were not an atomic sentence, but were instead the conjunction "(*K* & *W*)", then Φ_1 would not be a baseset.

Next, consider why none of the other six sets, i.e., none of $\Phi_2 - \Phi_7$, is a baseset. The reason is that for every Φ in $\Phi_2 - \Phi_7$, there is some reflection of Φ (namely, Φ_1), such that the operational value of Φ_1 (which is 0) is less than the operational value of Φ (since for each Φ in $\Phi_2 - \Phi_7$, the operational value of Φ is greater than 0). Thus, it follows from BASE that none of $\Phi_2 - \Phi_7$ is a baseset.

Now, as suggested above, if we have some fact *F*, and we want to determine whether Θ fully grounds *F*, we should seek to determine not only whether some sentential representation of Θ minimally entails some sentence that represents *F*, but also whether Θ contains any irrelevancy with respect to the obtainment of *F*. For, if Θ does contain some irrelevancy with respect to the obtainment of *F*, then it does not fully ground *F*. One way to determine whether Θ contains any such irrelevancy is to employ the notion of a baseset. Thus, I now turn to show how the notion of a baseset can be used to construct an account of full grounding.

4. Full grounds for logically contingent facts

With the notion of a baseset in mind, I now proceed to define the notion of full grounding by way of two stages. In the first stage, I focus on the grounds for logically contingent facts, and I draw a distinction between the notion of a *strict full ground* and the more general notion of a *full ground*. Then, in the second stage, I extend the accounts of strict full grounds and full grounds so that they apply also to logically necessary facts.

First, regarding the distinction between a *strict full ground* and a *full ground*, consider the following two sets:

Z = {the fact that Venus orbits the Sun}

Z* = {the fact that Venus orbits the Sun, the fact that Mars orbits the Sun}.

While each of Z and Z* is a full ground for the fact that either Venus or Mars orbits the Sun, the set Z but not the set Z* is a *strict* full ground for the fact that either Venus or Mars orbits the Sun. More generally, a full ground for some fact F is just the union of strict full grounds for F. Thus, in order to define the notion of a full ground for a logically contingent fact, I first define the notion of a strict full ground for a logically contingent fact, as follows:

STRICT-C For any logically contingent fact F, and any set of facts Φ that contains no member that is logically equivalent to F, Φ is a *strict full ground* for F if and only if there is some sentential representation Σ of Φ, and there is some sentence T that represents F, such that:

(g1) Σ is a baseset that minimally entails T.

To illustrate with some examples, suppose that both Na and Pa are atomic facts, and consider the following two claims:

(C1) $\{\forall x(Nx \,\&\, Bx)\}$ is not a strict full ground for Na.

(C2) $\{Pa, Pa \rightarrow Na\}$ is not a strict full ground for Na.

Regarding (C1), just as a conjunction is grounded in the set of its conjuncts (but not vice versa), it seems that a universally generalized fact is grounded in the set of its instances (but not vice versa). Given this, consider why according to STRICT-C, the set $\{\forall x(Nx \,\&\, Bx)\}$ is not a strict full ground for Na. The reason is that for any two sentences, *"X"* and *"Y"* (where *"X"* represents $\forall x(Nx \,\&\, Bx)$, and "*Y*" represents Na), the set {*"X"*} minimally entails *"Y"* only if {*"X"*} is not a baseset. For instance, while the set {*"*$\forall x$*(Nx & Bx)"*} minimally entails "*Na*" (for, the set entails "*Na*", and no proper subset of the set entails "*Na*"), this set is not a baseset, since its operational value is infinite.

Next, consider (C2). The reason that we do not want an account of full grounding according to which the set {*Pa*, *Pa* → *Na*} is a strict full ground for *Na* can be seen by supposing that: (i) *Pa* is the fact that 7 is a prime number that is not made of copper; (ii) (*Pa* → *Na*) is the fact that 7 is a prime number that is not made of copper only if 7 is an integer; and (iii) *Na* is the fact that 7 is an integer. Given this, it seems incorrect to say that {*Pa*, *Pa* → *Na*} is a strict full ground for *Na*. This accords with STRICT-C, since for any three sentences, *"W"*, "*V*", and *"Z"* (where *"W"* represents *Pa*, *"V"* represents "(*Pa* → *Na*)", and *"Z"* represents "*Na*"), the set {*"W"*, "*V"*} minimally entails *"Z"* only if {*"W"*, "*V"*} is not a baseset. For instance, while the set {"*Pa"*, "*Pa* → *Na"*} minimally entails "*Na*", this set is not a baseset, since there is some reflection of {"*Pa"*, "*Pa* → *Na"*} whose operational value is less than that of {"*Pa"*, "*Pa* → *Na"*}; namely, the set {"*Pa"*, "*Na"*}.

Now, to further illustrate the workings of STRICT-C, suppose that *A*, *B*, *C*, *Pa*, and *Na* are all distinct atomic facts that are represented, respectively, by the atomic sentences *"A", "B", "C", "Pa",* and *"Na"* (where the quotation marks distinguish the symbols as sentences), and consider the following proof sketches that show how, according to STRICT-C, each of the following claims holds (regardless of how fine-grained one's view is on the individuation of facts).

1. $\{A, B\}$ is a strict full ground for $(A \,\&\, B)$.

   Proof sketch: {*"A"*, *"B"*} is a baseset that minimally entails "$(A \,\&\, B)$". Thus, there is some sentential representation Σ of $\{A, B\}$, and there is some sentence T that represents $(A \,\&\, B)$, such that Σ is a baseset that minimally entails T.

2. $\{A, B\}$ is *not* a strict full ground for $((A \,\&\, B) \,\&\, C)$.

   Proof sketch: For any three sentences, *"X", "Y",* and *"Z"* (where *"X"* represents A, *"Y"* represents B, and *"Z"* represents $((A \,\&\, B) \,\&\, C)$), the set {*"X"*, *"Y"*} does *not* minimally entail *"Z"*.

3. $\{A, B \,\&\, C\}$ is *not* a strict full ground for $(A \,\&\, B)$.

   Proof sketch: For any three sentences, *"X", "Y",* and *"Z"* (where *"X"* represents A, *"Y"* represents $(B \,\&\, C)$, and *"Z"* represents $(A \,\&\, B)$), the set {*"X"*, *"Y"*} is *not* a baseset that minimally entails *"Z"*. For instance, while the set {*"A"*, *"B & C"*} minimally entails "$(A \,\&\, B)$", this set is not a baseset.

4. $\{\sim A, \sim B\}$ is a strict full ground for $\sim(A \vee B)$.

   Proof sketch: {*"~A"*, *"~B"*} is a baseset that minimally entails "$\sim(A \vee B)$".

5. $\{A, B \,\&\, C\}$ is a strict full ground for $(A \,\&\, (B \,\&\, C))$.

   Proof sketch: Let *"A"* be a sentence that represents A, and let *"G"* be a sentence that represents $(B \,\&\, C)$. Thus, {*"A"*, *"G"*} is a sentential representation of $\{A, B \,\&\, C\}$, and "$(A \,\&\, G)$" is a sentence that represents $(A \,\&\, (B \,\&\, C))$. Moreover, {*"A"*, *"G"*} is a baseset that minimally entails

"(A & G)". So, there is some sentential representation Σ of $\{A, B \& C\}$, and there is some sentence T that represents (A & (B & C)), such that Σ is a baseset that minimally entails T.

6. $\{A \leftrightarrow B, B\}$ is *not* a strict full ground for (A & B).

   Proof sketch: For any three sentences, *"X"*, *"Y"*, and *"Z"* (where *"X"* represents ($A \leftrightarrow B$), "*Y*" represents B, and *"Z"* represents (A & B)), the set {*"X"*, *"Y"*} minimally entails *"Z"* only if {*"X"*, *"Y"*} is not a baseset. For instance, while the set {*"A ↔ B"*, *"B"*} minimally entails "(A & B)", this set is not a baseset; and while the set {*"A"*, *"B"*} is a baseset, this set cannot be used as a sentential representation of $\{A \leftrightarrow B, B\}$ if it is to entail some sentence that represents (A & B).

7. $\{A \leftrightarrow B, B\}$ is a strict full ground for (($A \leftrightarrow B$) & B).

   Proof sketch: Let *"E"* be a sentence that represents ($A \leftrightarrow B$), and let *"B"* be a sentence that represents B. Thus, {*"E"*, *"B"*} is a sentential representation of $\{A \leftrightarrow B, B\}$, and "($E$ & B)" is a sentence that represents (($A \leftrightarrow B$) & B). Moreover, {*"E"*, *"B"*} is a baseset that minimally entails "(E & B)". So, there is some sentential representation Σ of $\{A \leftrightarrow B, B\}$, and there is some sentence T that represents (($A \leftrightarrow B$) & B), such that Σ is a baseset that minimally entails T.

8. $\{C \rightarrow A, \sim C \rightarrow A, B\}$ is *not* a strict full ground for (A & B).

   Proof sketch: For any four sentences, *"U"*, *"X"*, *"Y"*, and *"Z"* (where *"U"* represents ($C \rightarrow A$), *"X"* represents ($\sim C \rightarrow A$), *"Y"* represents B, and

"Z" represents (A & B)), the set {*"U"*, *"X"*, *"Y"*} minimally entails "*Z*" only if it is not a baseset.

9. $\{C \rightarrow A, \sim C \rightarrow A, B\}$ is a strict full ground for $(((C \rightarrow A) \,\&\, (\sim C \rightarrow A)) \,\&\, B)$.

   Proof sketch: Let *"H"* be a sentence that represents $(C \rightarrow A)$, and let *"K"* be a sentence that represents $(\sim C \rightarrow A)$, and let *"B"* be a sentence that represents B. Then, {*"H"*, *"K"*, *"B"*} is a sentential representation of $\{C \rightarrow A, \sim C \rightarrow A, B\}$, and "(($H$ & K) & B)" is a sentence that represents $(((C \rightarrow A) \,\&\, (\sim C \rightarrow A)) \,\&\, B)$. Moreover, {*"H"*, *"K"*, *"B"*} is a baseset that minimally entails "((H & K) & B)". So, there is some sentential representation Σ of $\{C \rightarrow A, \sim C \rightarrow A, B\}$, and there is some sentence T that represents $(((C \rightarrow A) \,\&\, (\sim C \rightarrow A)) \,\&\, B)$, such that Σ is a baseset that minimally entails T.

10. $\{A, B\}$ is a strict full ground for $((A \,\&\, B) \vee C)$.

    Proof sketch: {*"A"*, *"B"*} is a baseset that minimally entails "$((A \,\&\, B) \vee C)$".

11. $\{Pa\}$ is a strict full ground for $\exists xPx$.

    Proof sketch: {*"Pa"*} is a baseset that minimally entails "$\exists xPx$".

12. $\{\exists xPx, \forall xNx\}$ is a strict full ground for $(\exists xPx \,\&\, \forall xNx)$.

    Proof sketch: Let *"X"* be a sentence that represents $\exists xPx$, and let *"Y"* be a sentence that represents $\forall xNx$. Then, {*"X"*, *"Y"*} is a sentential representation of $\{\exists xPx, \forall x\mathrm{N}x\}$, and "($X$ & Y)" is a sentence that represents $(\exists xPx \,\&\, \forall xNx)$. Moreover, {*"X"*, *"Y"*} is a baseset that minimally entails "(X & Y)". So, there is some sentential representation Σ of {$\exists xPx$, $\forall xNx$}, and there is some sentence T that represents ($\exists xPx$ & $\forall xNx$), such that Σ is a baseset that minimally entails T.

13. {Pa, $\forall xNx$} is a strict full ground for ($\exists xPx$ & $\forall xNx$).

Proof sketch: Let *"Pa"* be a sentence that represents *Pa*, and let *"Y"* be a sentence that represents $\forall xNx$. Then, {*"Pa"*, *"Y"*} is a sentential representation of {Pa, $\forall xNx$}, and "($\exists x$Px & Y)" is a sentence that represents ($\exists x$Px & $\forall x$Nx). Moreover, {*"Pa"*, *"Y"*} is a baseset that minimally entails "($\exists x$Px & Y)". So, there is some sentential representation Σ of {Pa, $\forall xNx$}, and there is some sentence T that represents ($\exists xPx$ & $\forall xNx$), such that Σ is a baseset that minimally entails T.

Now, consider how STRICT-C can be used to construct an account of the more general notion of a full ground (in contrast to the notion of a strict full ground) as it pertains to logically contingent facts, as follows:

FULL-C For any logically contingent fact F, and any set of facts Φ that contains no member that is logically equivalent to F, Φ is a *full ground* for F if and only if Φ is the union of strict full grounds for F.

Given this definition of the notion of a full ground, we now have an account according to which a fact such as *PRIME* (i.e., the fact that either 7 or 12 is a prime number) is fully grounded in the set {the fact that 7 is a prime number, the fact that 12 is a prime number}, even though the sentence, "Either 7 or 12 is a prime number", is not minimally entailed by the set {"7 is a prime number", "12 is a prime number"}.

With this in mind, consider how differences between views on the individuation of facts might result in different claims about full grounds. There is a general consensus that any conjunction is fully grounded in its conjuncts (Correia & Schnieder 2012; Fine 2012). For instance, the fact (A & B) is fully grounded in the set {A, B}. Given this, suppose that we hold to a fairly coarse-grained view of facts on which the fact (A & A) is the same fact as A, and the fact (B & ($C \vee \sim C$)) is the same fact as B. It then clearly follows that (A & B) is fully grounded in {A & A, B & ($C \vee \sim C$)}, since on the coarse-grained view that we are considering, {A & A, B & ($C \vee \sim C$)} is identical to {A, B}. This accords well with FULL-C. To illustrate, let "A" be a sentence that represents the fact A, and let "B" be a sentence that represents the fact B. Then, on the considered coarse-grained view, the set {"A", "B"} is a sentential representation of both {A & A, B & ($C \vee \sim C$)} and {A, B}. Moreover, on the considered coarse-grained view, there is some sentence that represents (A & B), i.e., the sentence "(A & B)", such that (g1) of STRICT-C is satisfied with respect to {"A", "B"} as Σ, and "(A & B)" as T. Thus, {A & A, B & ($C \vee \sim C$)} is a strict full ground for (A & B), and thus (trivially) it is the union of strict full grounds for (A & B). So, according to FULL-C, {A & A, B & ($C \vee \sim C$)} is a full ground for (A & B).

By contrast, suppose that we adopt a more fine-grained view of facts on which the fact (A & A) is *not* the same fact as A, and the fact (B & ($C \vee \sim C$)) is *not* the same fact as B. Then, while it still seems clear that {A, B} is a strict full ground (and thus a full ground) for (A & B), it also seems that none of the following sets is either a strict full ground or a full ground for (A & B):

$\Gamma_1 = \{A \,\&\, A, B \,\&\, (C \vee \sim C)\}$

$\Gamma_2 = \{A, B \,\&\, (C \vee \sim C)\}$

$\Gamma_3 = \{A \,\&\, A, B\}$.

This likewise accords with STRICT-C and FULL-C. For, regarding the claim that on the fine-grained view of facts being considered, the set $\{A, B\}$ is a strict full ground (and thus a full ground) for (A & B), consider the following line of reasoning. The set {"A", "B"} is a sentential representation of $\{A, B\}$, and, moreover, there is some sentence that represents (A & B), i.e., the sentence "(A & B)", such that (g1) of STRICT-C is satisfied with respect to {"A", "B"} as Σ, and "(A & B)" as T. Secondly, regarding the claim that on the considered fine-grained view, none of the sets within $\Gamma_1 - \Gamma_3$ is either a strict full ground or a full ground for (A & B), consider another line of reasoning. Namely, there is no Γ within $\Gamma_1 - \Gamma_3$, such that for some sentential representation Σ of Γ, and some sentence T that represents (A & B), the set Σ and the sentence T satisfy (g1) of STRICT-C. For instance, suppose that A is not the same fact as (A & A). Then, regardless of whether (A & B) is represented by some conjunctive sentence, such as "(A & B)", or some atomic sentence, such as "X", there is no sentential representation of any of $\Gamma_1 - \Gamma_3$ that is both a baseset and a set that minimally entails a sentence representing (A & B).

At this point, one might wonder how this approach deals with issues of language dependence. Given that the account works largely with sentences, which are taken to represent facts, and given that different languages can represent facts in different ways, might the account face a potential problem? Namely, might differences between languages result in the account giving different verdicts about what grounds what? I submit that the answer is "no." The reason is as follows. Given the above account, a set of facts Φ grounds some fact F just in case there is at least some sentential representation Σ of Φ (in at least some language), and some sentence T representing F (in at least some language), such that Σ is a baseset the minimally entails T. So, suppose that there is a language L that consists only of atomic sentences. Then, every sentential representation Σ whose members are all sentences in L must be a baseset. Now, consider the following claim (which I take to be true):

(q1) {*the fact that 5 is prime*} fully logically grounds *the fact that 5 is either even or prime*.

If we could only use sentences in L to represent the facts that are mentioned in (q1), then it would follow from the above account that (q1) is false (which would be problematic). For, we could only represent *the fact that 5 is prime* with some atomic sentence (say, "*A*"), and moreover, assuming that no single sentence represents two different facts, we would then have to represent *the fact that 5 is either even or prime* with some other atomic sentence (say, "*B*"). But it is false that {"*A*"} is a baseset that minimally entails "*B*", since {"*A*"} does not entail "*B*". Thus, if this were the only way of representing the respective facts, then (q1) would turn out to be false on the above account. However, the account does not run into this problem. For, even though there could be a language like L, there is also the language of classical sentential logic in which we can use "P" to represent *the fact that 5 is prime*, and we can use "$E \vee P$" to represent *the fact that 5 is either even or prime*. And since there is such a language, and since

{"P"} is a baseset that minimally entails "$E \vee P$", it follows from the account that (q1) is true; and this is so regardless of the fact that languages like L are possible.

5. Full grounds for logically necessary facts

Now, return again to the notion of a baseset. While this notion is useful towards ensuring that a given set does not contain any irrelevancies with respect to the obtainment of some *logically contingent fact*, something further is still needed if we want to capture certain key ideas pertaining to *non-contingent facts*; e.g., the idea that while the set $\{C\}$ does not fully ground ($B \vee \sim B$), the set $\{B\}$ does fully ground ($B \vee \sim B$). To capture this idea, I introduce the notion of *joint parenthood*, defined as follows:

PARENT For any sentences N and T, and any set of sentences Σ, N and Σ are *joint parents* of T if and only if:

(r1) either T is identical to N, or T differs from N only in that for some occurrences of some atomic formulas that are in N but not in any member of Σ, T has in their place occurrences of some arbitrary atomic formulas.

To illustrate, suppose that "A", "B", "C", and "D" are all atomic sentences (and thus all atomic formulas), and "Fx" and "Gx" are both atomic formulas. Then, consider the relationship between the set {"A", "C"} and the following two sentences:

Sentence 1: "$(A \& C) \vee ((B \vee \sim B) \vee \forall x(Fx \rightarrow Fx))$"

Sentence 2: "$(A \& C) \vee ((D \vee \sim B) \vee \forall x(Gx \rightarrow Fx))$"

Sentence 1 and the set {"A", "C"} are joint parents of Sentence 2, since Sentence 2 differs from Sentence 1 only in that for some occurrences of some atomic formulas that are in Sentence 1 but not in any member of {"A", "C"} (i.e., for the first occurrence of "B" and the first occurrence of "Fx" in Sentence 1), Sentence 2 has in the place of these occurrences some arbitrary atomic formulas (i.e., it has "D" in place of "B", and "Gx" in place of "Fx").

The following point is also worth noting, since it makes it easier to think about many cases that are relevant to the notion of full grounding (as defined below):

(p1) For any sentence T, and any set of sentences Σ, if each atomic sentence in T is also in some member (or other) of Σ, then T and Σ are joint parents of only one sentence: namely, T itself.

For instance, the set {"A", "C & D"} and the sentence "$((A \& C) \vee (A \& D))$" are joint parents of only one sentence: namely, "$((A \& C) \vee (A \& D))$" itself.

Now, in order to define the notion of a full ground so that it applies to both logically contingent facts and logically necessary facts, consider some notation. For any sentence T, and any set of sentences Σ, let $\mathrm{J}(T \mid \Sigma)$ be the set of every sentence that has T and Σ as joint parents. So, given (p1), if each atomic sentence in T is also in some member (or other) of Σ, then $\mathrm{J}(T \mid \Sigma) = \{T\}$. Now, I define the notion of a *strict full ground* (as it pertains to any fact) as follows:

STRICT For any fact F, and any set of facts Φ containing no member that is logically equivalent to F, Φ is a *strict full ground* for F if and only if there is some sentential representation Σ of Φ, and there is some sentence T that represents F, such that:

(s1) Σ is a baseset that minimally entails every member of J(T | Σ).

Note that given point (p1) from above, if we are considering a case in which each atomic sentence in T is also in Σ, then condition (s1) in STRICT can just be read as saying, "Σ is a baseset that minimally entails T."

Finally, given STRICT, the notion of a *full ground* (as it pertains to any fact) can be defined as follows:

FULL For any fact F, and any set of facts Φ containing no member that is logically equivalent to F, Φ is a *full ground* for F if and only if Φ is the union of only strict full grounds for F.

Consider how FULL applies to some specific cases. Let A, B, and C be distinct atomic facts, and then consider the following claims along with the corresponding proof sketches.

Claim 1: According to FULL, the set {A} is *not* a full ground for (B ∨ ~B).

Proof sketch:

Suppose first that "A" is an atomic sentence that represents A, and "X" is an atomic sentence that represents (B ∨ ~B). Then, J("X" | {"A"}) contains every atomic sentence. So, J("X" | {"A"}) contains "B". But {"A"} does *not* minimally entail "B". Thus, it is false that {"A"} is a baseset that minimally entails every member of J("X" | {"A"}). So, on the assumption that "X" is an atomic sentence, it follows that according to STRICT, {A} is not a strict full ground for (B ∨ ~B). And since A is the only member of $\{A\}$, it follows that, according to *FULL*, $\{A\}$ is *not* a full ground for $(B \vee {\sim}B)$.

Next, suppose instead that "X" is a non-atomic sentence of the form "$(B \vee {\sim}B)$". Then, since "B" is not in any member of {"A"}, the set J("$B \vee {\sim}B$" | {"A"}) has "$(E \vee {\sim}B)$" as a member. However, {"A"} does not minimally entail "$(E \vee {\sim}B)$". Thus, it is false that {"A"} minimally entails every member of J("$B \vee {\sim}B$" | {"A"}). So, on the assumption that "X" is of the form "$(B \vee {\sim}B)$", it again follows that according to STRICT, $\{A\}$ is not a strict full ground for $(B \vee {\sim}B)$. And since A is the only member of $\{A\}$, it follows that, according to *FULL*, $\{A\}$ is *not* a full ground for $(B \vee {\sim}B)$.

<u>Claim 2</u>: According to FULL, the set $\{B\}$ *is* a full ground for $(B \vee {\sim}B)$.

<u>Proof sketch</u>:

Suppose that "B" is a sentence that represents B, and "$(B \vee {\sim}B)$" is a sentence that represents $(B \vee {\sim}B)$. Given that "B" is a member of {"B"}, the set J("$B \vee {\sim}B$" | {"B"}) contains only one member: i.e., the sentence "$(B \vee {\sim}B)$". Moreover, since "B" is an atomic sentence, {"B"} is a baseset, and {"B"} minimally entails "$(B \vee {\sim}B)$". Thus, {"B"} is a baseset that minimally entails every member of J("$B \vee {\sim}B$" | {"B"}). So, according to STRICT, $\{B\}$ is a strict full ground for $(B \vee {\sim}B)$. And thus, according to FULL, $\{B\}$ is a full ground for $(B \vee {\sim}B)$.

Claim 3: According to FULL, the set {*A*, ~*A*} is *not* a strict full ground for *B*.

Proof sketch:

Suppose that "*A*" is a sentence that represents *A*, and "*X*" is a sentence that represents ~*A*, and "*B*" is a sentence that represents *B*. Then, the set J("*B*" | {"*A*", "*X*"}) contains every atomic sentence. So, J("*B*" | {"*A*", "*X*"}) contains "*A*". But {"*A*", "*X*"} does *not minimally* entail "*A*". Thus, it is false that {"*A*", "*X*"} is a baseset that minimally entails every member of J("*B*" | {"*A*", "*X*"}). So, according to STRICT, {*A*, ~*A*} is not a strict full ground for *B*.

Moreover, since "*X*" represents ~*A*, {"*X*"} fails to entail (and thus fails to minimally entail) the sentence "*A*". So, it is false that {"*X*"} is a baseset that minimally entails every member of J("*B*" | {"*A*", "*X*"}). So, according to STRICT, {~*A*} is not a strict full ground for *B*. Thus, neither {*A*, ~*A*} nor {~*A*} is a strict full ground for *B*. So, {*A*, ~*A*} is not the union of only strict full grounds for *B*. So, according to FULL, {*A*, ~*A*} does *not* fully ground *B*.

6. Understanding-why and logical grounding

Consider how the above account of full logical grounding pertains to the notion of understanding-why. First, suppose that we have some true event proposition (i.e., a true proposition representing an event), such as the proposition, <Alice and Micah summited Mount Kilimanjaro>. Then, consider the following claim with respect to full logical grounding:

(w1) *The fact that Alice and Micah summited Mount Kilimanjaro* is fully logically grounded in the set {*the fact that Alice summited Mount Kilimanjaro*, *the fact that Micah summited Mount Kilimanjaro*}.

While (w1) is true, it seems that understanding that (w1) is true is neither necessary nor sufficient for obtaining any understanding of why Alice and Micah summited Mount Kilimanjaro. This accords with the account of understanding-why offered in Chapter 2. In particular, according to that account, for any true event proposition *P*, a person has at least some understanding of why *P* if and only if the person has a causal guiding explanation of why *P* (where having a causal guiding explanation of why *P* requires no understanding of the logical ground for the fact corresponding to *P*).

Next, consider a true event proposition representing a fact that can be understood by way of logical grounding, such as the proposition, <Alice or Mary summited Mount Kilimanjaro>. Here, it seems that a person could have at least some understanding of why this proposition is true by understanding that the following claim holds:

(w2) *The fact that Alice or Mary summited Mount Kilimanjaro* is fully logically grounded in the set {*the fact that Alice summited Mount Kilimanjaro*}.

Thus, while understanding that (w1) is true does *not* seem to provide one with any understanding of why it is true that Alice and Micah summited Mount Kilimanjaro, understanding that (w2) is true *does* seem to provide one with at least some understanding of why it is the true that either Alice or Mary summited Mount Kilimanjaro. This likewise accords with the account of understanding-why provided in

Chapter 2. Recall that, according to Chapter 2, a person has at least some understanding of why *P* (where *P* is a true event proposition) if and only if the person has a causal guiding explanation of why *P*. Moreover, a person has a causal guiding explanation of why *P* if and only if there is some proposition *C*, such that: (i) the person understands that *C*; and (ii) the truth of *C* provides information, which is not provided by *P*, about the causal nexus that was responsible for the fact that *P*.

There are, however, also instances of understanding-why that are connected to logical grounding, and such that the relevant proposition is *not* an event proposition. Take, for instance, the following two propositions:

<Either 7 or 111 is prime>

<Both 7 and 101 are prime>

Now, consider the following two claims:

(v1) *The fact that either 7 or 111 is prime* is fully logically grounded in the set {*the fact that 7 is prime*}.

(v2) *The fact that both 7 and 101 are prime* is fully logically grounded in the set {*the fact that 7 is prime, the fact that 101 is prime*}.

While it seems that having an understanding that (v1) is true provides one with at least some understanding of why it is true that either 7 or 111 is prime, it also seems that understanding that (v2) is true does *not* provide one with any understanding of why it is true that both 7 and 101 are prime.

In light of the above considerations, I offer the following partial account of understanding why *P*, which applies regardless of whether *P* is an event proposition:

UNDWHY For any subject *S*, and any true proposition *P*, *S* has at least some understanding of why *P* if for some fact *F* that corresponds to *P*, and for some set of facts Γ that fully logically grounds *F*:

(u1) the conjunction of all the members of Γ is not identical to *F*; and

(u2) *S* understands that *F* obtains in virtue of the fact that every member of Γ obtains.

I say that the account is "partial," since it only provides a set of *sufficient* conditions for understanding why *P*. For instance, as noted with respect to causal guiding explanations, there are ways in which one might understand why *P* that do not involve any understanding of a logical grounding relation. To illustrate, one might have at least some understanding of why a particular vase broke by way of understanding that it broke because it was fragile. However, this need not involve any understanding of a logical grounding relation.

Moreover, there may well be other types of grounding (e.g., metaphysical grounding) that are likewise sufficient to provide one with at least some understanding of why *P*. For example, one might understand why a particular vase is fragile by understanding that the fact that it is fragile is metaphysically grounded in the fact that it has an irregular atomic structure. Even though this type of understanding-why does not require having a causal guiding explanation, it does not conflict with the account of understanding-why presented in Chapter 2. For, the account given in Chapter 2 only applies to cases in which the proposition in question is an event proposition, and the proposition <the vase is fragile> does not represent an event.

In light of this, how might we more fully understand the relationship between the account of understanding-why presented in Chapter 2 and that which is provided above, i.e., UNDWHY? To answer this, consider Figure 1.

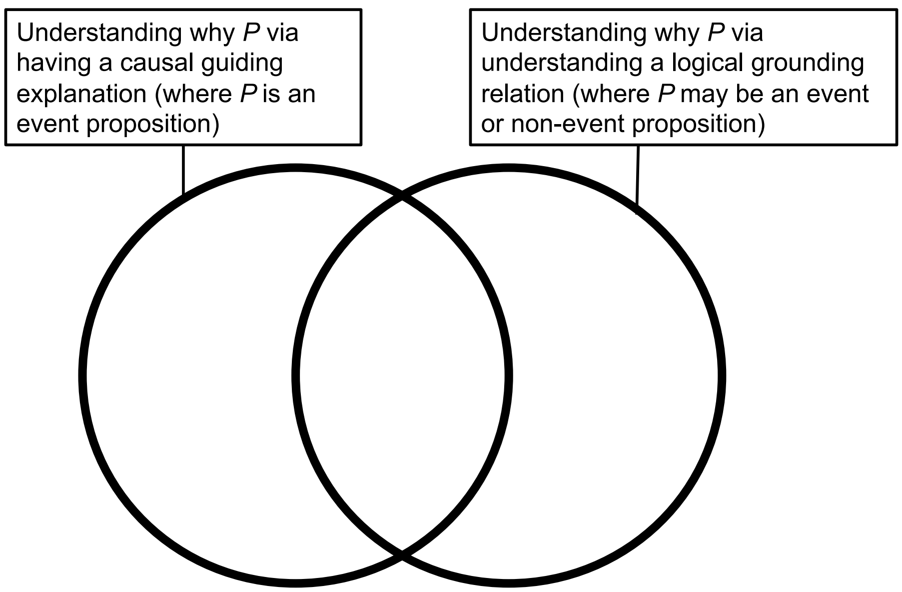

Figure 1. A representation of two types of understanding-why.

The left circle represents understanding-why by way of having a causal guiding explanation, while the right circle represents understanding-why by way of understanding a logical grounding relation.

The leftmost region represents the type of understanding-why that one has with respect to a true event proposition when one has a causal guiding explanation of why *P* but does not understand any logical grounding relation pertaining to *P*. For instance, one might understand why a vase broke by understanding that it broke because it was fragile, and this requires no understanding of any logical grounding relation.

The center region, where the two circles overlap, represents the type of understanding-why that one has with respect to a true event proposition when one does indeed understand a logical grounding relation. For instance, one might have at least some understanding of why it is the true that either Alice or Mary summited Mount Kilimanjaro by way of understanding that *the fact that either Alice or Mary summited Mount Kilimanjaro* is fully logically grounded in the set {*the fact that Alice summited Mount Kilimanjaro*}. Moreover, given the conjunction of the two accounts of understanding-why (i.e., the account presented in Chapter 2 and that presented in the present chapter), it follows that for any true event proposition *P*, understanding why *P* is true by way of understanding some logical grounding relation is an instance of understanding why *P* is true by way of having a casual guiding explanation of why *P*. Now, as noted above, according to Chapter 2, a person has at least some understanding of why *P* (where *P* is a true event proposition) if and only if the person has a causal guiding explanation of why *P*. Furthermore, a person has a causal guiding explanation of why *P* if and only if there is some proposition *C*, such that: (i) the person understands that *C*; and (ii) the truth of *C* provides information, which is not provided by *P*, about the causal nexus that was responsible for the fact that *P*. With this in mind, suppose that a person *S* has at least some understanding of why it is the true that either Alice or Mary summited Mount Kilimanjaro by way of understanding that *the fact that either Alice or Mary summited Mount Kilimanjaro* is fully logically grounded in the set {*the fact that Alice summited Mount Kilimanjaro*}. Then, it is true that *S* has a causal guiding explanation of why either Alice or Mary summited Mount Kilimanjaro. For, by understanding the logical grounding relation, *S* obtains information, which is not provided by the proposition <either Alice or Mary summited Mount Kilimanjaro>, about the causal nexus that was responsible for the fact that either Alice or Mary summited Mount Kilimanjaro; e.g., the information that there were causal properties in relation to *Alice* that were casually responsible for the fact that either Alice or Mary summited Mount Kilimanjaro.

Finally, the rightmost region of Figure 1 represents the type of understanding-why that one has with respect to a true non-event proposition by way of understanding some logical grounding relation. For instance, as mentioned above, one might understand why it is true that either 7 or 111 is prime by way of understanding that the fact that *the fact that either 7 or 111 is prime by way of understanding* is fully logically grounded in the set {*the fact that 7 is prime*}.

7. Conclusion

In this chapter, I began by considering various types of grounding in order to situate the relation of logical grounding among these various types. In Section 2, I examined the relationship between *logical grounding* and *minimal entailment*, and I introduced the notion of a *sentential representation*, which serves to connect *sets of facts* to *sets of sentences representing those facts*. In Section 3, I looked at how certain sets of sentences are *more basic* than others, and I proceeded in Section 4 to explore how this idea might be employed in constructing an account of logical grounding with respect to logically contingent facts. In Section 5, I extended the account that is presented in Section 4 so that it also applies to logically necessary facts, and I examined a few final instances of logical grounding and showed how the account offered here applies to such cases. Lastly, in Section 6, I considered how the account offered here pertains to the notion of understanding-why.

Bibliography

Achinstein, P. (1983). *The Nature of Explanation*, Oxford: Oxford University Press.

Annis, D. B. (1978). "A Contextualist Theory of Epistemic Justification", *American Philosophical Quarterly*, 15: 213–219.

Armstrong, D. M. (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.

Armstrong, D. M. (1996). "Dispositions as Categorical States", in D. M. Armstrong, C. B. Martin, & U. T. Place, *Dispositions: A Debate* (pp. 15-18). London: Routledge.

Audi, P. (2012). "A Clarification and Defense of the Notion of Grounding", in F. Correia and B. Schnieder (Eds.), *Metaphysical Grounding* (pp. 101-121), Cambridge: Cambridge University Press.

Baumberger, C. (2011). "Understanding and its Relation to Knowledge", in *Epistemology: Contexts, Values, Disagreement. Papers of the 34th International Wittgenstein Symposium. Kirchberg am Wechsel: Austrian Ludwig Wittgenstein Society,* eds. C. Jäger & W. Löffler, 16–18.

Beebe, J. (2013). "A Knobe effect for belief ascriptions," *The Review of Philosophy and Psychology*, 4, 235-258.

Berker, S. (2015). "The Unity of Grounding", manuscript.

Bird, A. (1998). "Dispositions and Antidotes", *The Philosophical Quarterly, 48,* 227-234.

Bird, A. (2005). "The Dispositionalist Conception of Laws", *Foundations of Science*, *10*, 353–370.

Block, N. (2003). "Do Causal Powers Drain Away?", *Philosophy and Phenomenological Research*, *67*, 133–50.

Bontly, T. D. (2001). "The Supervenience Argument Generalizes", *Philosophical Studies*, *109*, 75–96.

Brogaard, B. (2005). "I Know. Therefore, I Understand", manuscript.

Buckwalter, W., Rose, D., & Turri, J. (2013). "Belief through thick and thin", *Noûs*.

Carter J., & Gordon E. (forthcoming), "On Pritchard, Objectual Understanding and the Value Problem", *American Philosophical Quarterly*.

Cohen, S. (1987). "Knowledge, Context, and Social Standards", *Synthese*, 73: 3–26.

Cohen, S. (2005). "Knowledge, Speaker, Subject", *The Philosophical Quarterly*, 55: 199–212.

Correia, F. (2005). *Existential Dependence and Cognate Notions*, München: Philosophia Verlag.

Correia, F. (2008). "Ontological Dependence", *Philosophy Compass*, *3*, 1013-1032.

Correia, F., & Schnieder, B. (2012). "Grounding: An Opinionated Introduction", in F. Correia and B. Schnieder (Eds.), *Metaphysical Grounding* (pp. 1-36), Cambridge: Cambridge University Press.

Cross, T. (2005). "What is a Disposition?", *Synthese*, *144*, 321–341.

Daly, C. (2012). "Skepticism about Grounding", in F. Correia and B. Schnieder (Eds.), *Metaphysical Grounding* (pp. 81-100), Cambridge: Cambridge University Press.

Dasgupta, S. (2014). "On the Plurality of Grounds", *Philosophers' Imprint*, *14*, 1–28.

DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context: Vol. 1*, Oxford: Clarendon Press.

deRosset, L. (2010). "Getting Priority Straight", *Philosophical Studies*, *149*, 73–97.

Elgin, C. (2007). "Understanding and the Facts", *Philosophical Studies*, 132: 33–42.

Elgin, C. (2009). "Is Understanding Factive?", in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, 313–21, Oxford: Oxford University Press.

Fine, K. (2001). "The Question of Realism", *Philosophers' Imprint*, *1*, 1-30.

Fine, K. (2012). "Guide to Ground", in F. Correia and B. Schnieder (Eds.), *Metaphysical Grounding* (pp. 37-80), Cambridge: Cambridge University Press.

Friedman, M. (1974). "Explanation and Scientific Understanding", *Journal of Philosophy*, *71*, 5–19.

Gordon, E. (2012). "Is There Propositional Understanding?", *Logos & Episteme*, 3: 181-192.

Grimm, S. (2006). "Is Understanding a Species of Knowledge?", *British Journal of Philosophy of Science,* 57: 515-535.

Grimm, S. (2010). “Understanding”, in *Routledge Companion to Epistemology*, eds. S. Bernecker & D. Pritchard, New York: Routledge.

Grimm, S. (forthcoming). "Understanding as Knowledge of Causes", *Synthese*.

Hambourger, R. (1987). “Justified Assertion and the Relativity of Knowledge”, *Philosophical Studies*, 51: 241–269.

Hawthorne, J. (2004). *Knowledge and Lotteries*, New York and Oxford: Oxford University Press.

Hawthorne, J., & Manley, D. (2005). “Review of Stephen Mumford's Dispositions”, *Noûs*, *39*, 179–195.

Hempel, C. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Hills, A. (2010). *The Beloved Self*, Oxford: Oxford University Press.

Jackson, F., & Pettit, P. (1990). “Program Explanation: A General Perspective,” *Analysis*, *50*, 107–117.

Khalifa, K. (2012). “Inaugurating Understanding or Repackaging Explanation?” *Philosophy of Science, 79*, 15-37.

Khalifa, K. (2013). “The Role of Explanation in Understanding”, *British Journal for the Philosophy of Science, 64*, 161-187.

Kim, J. (1998). *Mind in a Physical World*, Cambridge, MA: MIT Press.

Kitcher, P. (1989). “Explanatory Unification and the Causal Structure of the World”, in P. Kitcher and W. Salmon (Eds.), *Scientific Explanation* (pp. 410–505). Minneapolis: University of Minnesota Press.

Kitcher, P. (2002). “Scientific Knowledge”, in *Oxford Handbook of Epistemology*, ed. P. Moser, Oxford University Press, Oxford.

Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding,* Cambridge: Cambridge University Press.

Lange, M. (2015). “Depth and Explanation in Mathematics,” *Philosophia Mathematica, 23,* 196-214.

Lewis, D. (1997). “Finkish Dispositions,” *The Philosophical Quarterly*, *47*, 143-158.

Lewis, D. (1986). “Causal Explanation,” in D. Lewis, *Philosophical Papers II* (pp. 214 – 240). Oxford: Oxford University Press.

Lipton, P. (2009). “Understanding without explanation”, in H. W. de Regt, S. Leonelli, and K. Eigner (Eds.), *Scientific Understanding: Philosophical Perspectives* (pp. 43-63). Pittsburgh: University of Pittsburgh Press.

Mayer, R. E. (2009). *Multimedia learning* (2nd ed.). New York: Cambridge University Press.

Molnar, G. (1999). "Are Dispositions Reducible?" *The Philosophical Quarterly*, *44,* 1-17.

Molnar, G. (2003). *Powers: A Study in Metaphysics*, S. Mumford (Ed.), Oxford: Oxford University Press.

Murray, D., Sytsma, J., & Livengood, J. (2013). "God knows (but does God believe?),” *Philosophical Studies*, 166, 83-107.

Prior, E. W., Pargetter, R., & Jackson, F. (1982). “Three Theses about Dispositions”, *American Philosophical Quarterly*, *9*, 251-253.

Pritchard, D. (2009). “Knowledge, Understanding and Epistemic Value”, in *Epistemology (Royal Institute of Philosophy Lectures),* ed. A. O'Hear, 19–43, Cambridge: Cambridge University Press.

Raven, M.J. (2012). "In Defense of Ground", *Australasian Journal of Philosophy*, *90*, 687-701.

Riggs, W. (2003). “Understanding ‘Virtue’ and the Virtue of Understanding”, in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. M. DePaul & L. Zagzebski, 203-226, Oxford: Oxford University Press.

Rosen, G. (2010). “Metaphysical Dependence: Grounding and Reduction”, in B. Hale and A. Hoffman (Eds.), *Modality: Metaphysics, Logic, and Epistemology* (pp. 109-136), Oxford: Oxford University Press.

Salmon, W. C. (1971). “Statistical Explanation”, in W. Salmon (Ed.), *Statistical Explanation and Statistical Relevance* (pp. 29–87). Pittsburgh: University of Pittsburgh Press.

Salmon, W. C. (1998). *Causality and Explanation*. New York: Oxford University Press.

Schaffer, J. (2009). “On What Grounds What”, in D. Chalmers, D. Manley, and R. Wasserman (Eds.), *Metametaphysics* (pp. 347-383), Oxford: Oxford University Press.

Schrenk, M. (2010). “Antidotes for Dispositional Essentialism”, in A. Marmodoro (Ed.), *The Metaphysics of Powers: Their Grounding and Their Manifestations*. Routledge.

Skyrms, B. (1980). *Causal Necessity*, New Haven: Yale University Press.

Stanley, J. (2005). *Knowledge and Practical Interests*, New York and Oxford: Oxford University Press.

Strevens, M. (2013). “No Understanding without Explanation”, *Studies in History and Philosophy of Science*, *44*, 510-515.

Trogdon, K., 2013, “Grounding: Necessary or Contingent?” *Pacific Philosophical Quarterly*, *94*, 465–485.

Trout, J. D. (2002). “Scientific Explanation and the Sense of Understanding”, *Philosophy of Science*, *69*, 212-233.

Trout, J. D. (2007). “The Psychology of Scientific Explanation”, *Philosophy Compass*, *2*, 564–591.

Van Fraassen, B. (1980). *The Scientific Image*. Oxford: Oxford University Press

Walter, S. (2008). “The Supervenience Argument, Overdetermination, and Causal Drainage: Assessing Kim's Master Argument”, *Philosophical Psychology*, *21*, 673–96.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.

Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.

Zagzebski, L. (2001). “Recovering Understanding”, in *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility and Virtue*, (ed.) M. Steup, Oxford University Press, Oxford.

Zagzebski, L. (2008). *On Epistemology*, Belmont, Calif.: Wadsworth.